UNIVERSITY of CALIFORNIA
Santa Barbara

Simulating Quantum Chemical Dynamics with Improved Superconducting Qubits

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

Doctor of Philosophy

in

Materials

by

Anthony E. Megrant

Committee in charge:

Professor Chris Palmstrom, Chair
Professor Andrew Cleland
Professor Arthur Gossard
Professor Chris Van De Walle

September 2016

The dissertation of Anthony E. Megrant is approved:

Professor Andrew Cleland

Professor Arthur Gossard

Professor Chris Van De Walle

Professor Chris Palmstrom, Chair

August 2016

Acknowledgements

I was fortunate enough to work closely with three great advisors during my graduate work, Chris Palmstrom in the Materials department and both Andrew Cleland and John Martinis in the Physics department. When I first joined, they were beginning a collaboration to build quantum devices with improved performance through materials research. My role was to bridge between the two fields. A common trait among all three advisors is the push to build things ourselves. I really enjoyed this, as it was so familiar to me from my previous work experiences.

When I first met Chris he joked that my future will be filled with 'Many Boring Evenings' to repair the 'Mostly Broken Equipment' on the 'Mega Bucks Evaporator', better known as a molecular beam epitaxy (MBE) system. Of course his jokes, or prophecies, came true repeatedly over the years of grad school. Chris taught me from the very start to take these set backs with a smile and a joke. Thank you for all your guidance in materials research and MBE advise over the years. I will miss the group BBQ's on the beach and at your house.

During my first meeting with Andrew he handed me a stack of the group's previously published papers and said, 'Here read these. You probably won't understand much, but keep reading and you will start to get it.' This was definitely true and I still think back to that moment when I read a paper that I don't understand. Andrew also guided me through dealing with the large scope of my task, dissecting it in to smaller chunks. He encouraged me to read over the instruction manuals of the numerous pieces of equipment which I might use, to learn the theory behind the tool and how to best use it, according to the people who made it. This seems like such common sense but after working with many other scientists, I realize how rare this trait is. Thank you for giving me the opportunity to join the group, for the great advise over the years, and for your door always being open to talk about basically anything.

John it is obvious that you are truly passionate about actually building a quantum computer, which rubs off on many of your students. This is probably best illustrated in your enthusiasm toward understanding and openly discussing all of the 'dirt physics' that needs to be solved, as opposed to just getting something to work. Thank you for teaching me the immense value of plot-ology: plotting data in the best way to make the conclusions obvious, fostering my leadership skills, and providing us with an amazing lab equipped with so many cool tools. Partnering with Google was a major milestone for you, the group, and the field and I look forward to continuing to work on this awesome project with you at Google.

I would like to thank all three of you for giving me the opportunity to work on such interesting projects with amazing equipment and for promoting such a collaborative effort from the top down.

I would like to acknowledge the other members of my thesis committee: professors Chris Van de Walle and Arthur Gossard. Thank you for the valuable advise you have given me over the years and encouragement after my exams and thesis defense.

Mike Geller, has collaborated with the group since before my time. It was great to finally get a chance to work closely together on the collision simulation. Thanks for developing such an interesting project to work on and providing the theoretical support and encouragement along the way.

I would like to acknowledge the Coherent Superconducting Qubits program under Intelligence Advanced Research Projects Activity (IARPA). This program began when I first joined the group and funded all of my research up until my final project. I would also like to thank Google for research support during my final work on quantum simulation. The transition from government to industrial funding was a stressful time for the group and ultimately led to returning equipment, that kept us from fabricating new devices for almost half a year. However, it was during this time that I seized the opportunity to work on quantum simulations, something drastically different than my previous work. While this entire process certainly prolonged my time as a graduate student, I really enjoyed tackling the simulation project.

When I joined the group, the lab already had considerable infrastructure, both physically and in software. I would like to thank all the previous graduate students and postdocs for helping build such a great lab to work in. In particular, I would like to point out the postdocs Ludwig, Alex, and Martin. Thanks for getting me started on the ins and outs of MBE systems and growing and characterizing my first rhenium thin films.

I would like to thank Charles and Pedram who helped me on the collision simulation. Charles, was my first partner in crime. Together we figured out how to reliably measure resonators and stumble our way through the massive software infrastructure of the lab. We did much of this through trial and error, leading us to use the existing infrastructure in some pretty unintended ways. Recently, we got the chance to work together again on the collision simulation. I am really grateful for your willingness to drop almost anything to lend a helping hand or to talk about random physics. Thanks for designing and building such a great device. Pedram, after you first joined the group, I quickly pegged you as an equipment hoarder. When ever anybody asked where this piece or that tool was, I would respond, “Did you check Pedram’s office?” This is usually where the search ended. I am so glad I got to know you much better while working together on the collision simulation. I thank you for your mentoring during this project and your lasting friendship. As far as your hoarding goes, I now understand that you were simply trying to add some order to the disordered lab, that we can all thank you for.

Yu Chen, we never worked closely together on a project, but you always have great advise for me. I watched you tackle some of our group’s toughest infrastructure projects even though they weren’t very flashy or high profile. The entire group is better off for the work you have done. Thank you for your work on the gmon architecture which made my collision simulation possible. Also thanks for all of the delicious authentic Chinese food in San Francisco. Rami you got me started on my first resonator measurements and provided the group with a great measurement setup, aptly named, ‘The Rami Setup’. Thanks for starting my resonator measurements off on the right foot and for developing a great qubit design with Julian.

Next I would like to thank Jimmy, Andrew, Chris, Ben and Brooks, who I worked closely with on the SiXmon project as well as other materials based projects. Jimmy, we have worked together on several projects, from airbridges to Arduinos to Xmons on silicon. You hit the ground running after joining the group and have only sped up since. It has been great working with you on all of these projects. Andrew, moving Xmons to silicon has been difficult to say the least. It was a long, hard road but in the end we figured it out and got some great results. Thanks for sticking it out with me till the end, for being a great neighbor, and for teaching me that adding ‘guac’ is always the right answer. I hope the MBE treats you as well as it did me. Chris, you came to the lab with so much experience and knowledge already. It has been great watching you tackle difficult projects with ease. I’m sure if anyone can build a useful quantum annealer it will be you and Yu Chen. Thank you for all your help and keen attention to detail. Ben, we have worked together on many projects. From the very beginning you had a strong work ethic, staying late to get this measurement finished or that sample grown. Even now with your growing family, I see you tackling your flux noise measurements with a similar level of intensity. Thanks for all your hard work over the years. Brooks, its been great working with you on early on with your packaging project. I know you are close to finishing up that project and can’t wait to see how your next scaling project pushes the field forward.

Peter, Ted, Amit, Julian, Dan, it was great sharing an office with all you guys. I found that I was much more likely to laugh and gossip then get work done, but that’s because it was so easy to get along with you guys. Peter, thanks for all your patience and help with the many programming questions I had and for your work on the ADR control server. Also, thanks for the fresh veggies from your garden and for staying up way too late with me at APS. Ted, its safe to say when people walk past us having a conversation they are just as likely to hear about physics as somebody getting kicked in the face or choked out. Its been great spreading the word about MMA together and seeing people get interested in our conversations and choose their favorite fighters, even though this has led to some spoilers. Of course your work on amplifiers helped make every qubit experiment more successful, including my collision simulation. Amit, you quite literally hold the plug to the lab in your hands. Thanks for all of your work on the infrastructure from the awesome simulation server you built to securing our data integrity. Julian, during our time in the lab the group moved from the previous generation’s phase qubits to the Xmon qubits. Your hard work was obviously instrumental in making this a success and paved the way for experiments like mine. Dan, thanks for your help over the years troubleshooting electronics. The group and I have also benefited greatly from your habit of documenting calculations and derivations. I wish more of your habits would have rubbed off on me.

My original project, ‘make better materials’ was a little broad in scope, to say the least. Jim’s participation ratio simulations gave me something much more tangible to focus on. Jim, you are always willing to help out others, which I have benefited from greatly over the years. And once again, thanks for shutting off the pump cart before you leave (I’m sure I owe you at least one more). Josh, thanks for your help with MBE and ADR maintenance over the years. Your amplifier work made all of our qubit experiments better or even possible. Evan, you are a know it all. Not in the negative sense, you just seem to know it all. I strive to be able to say something intelligent on half the things you know. Austin, your work on the surface code provided the entire lab with a clear direction with realistic goals. It has been great to see your hard work in helping build the infrastructure of the lab as we continue to expand. Also, thanks for all the cakes and treats you have brought in to the lab over the years.

Mom and Dad, thanks for having me and encouraging me to go back to school and continue on this crazy journey. Nikola, you made me feel like the coolest big brother when I was in the army and the biggest nerd when I went back to school for physics. Thanks for keeping me humble little bro. Kristiann I do not possess the words to express my gratitude to you. We have made it through some very difficult times that very few couples see through. We have had two long distance relationships and we even made it without running water. What ever difficulties life throws our way, I know that I have a partner who will see it through with me. Driving across the country with you and Willow was so much fun, and I can't wait for our European road trip. Thank you for your patience, kindness, love, and your unwavering support for not only the past six years but the entire time I have known you.

Curriculum Vitæ

Anthony E. Megrant

Education

(2016) Ph.D., Materials, University of California Santa Barbara

(2009) B.S., Engineering Physics, Magna Cum Laude, Cornell University

(2004) A.S., Engineering Science, with High Honors, Jefferson Community College

(2002) A.A.S., Business Administration, with High Honors, Jefferson Community College

(1996) Larkin high school, Elgin, IL

Publications

“Planar superconducting resonators with internal quality factors above one million”, A. Megrant, C. Neill, R. Barends, B. Chiaro, Yu Chen, L. Feigl, J. Kelly, Erik Lucero, Matteo Mariantoni, P. O’Malley, D. Sank, A. Vainsencher, J. Wenner, T. C. White, Y. Yin, J. Zhao, C. J. Palmstrom, John M. Martinis, A. N. Cleland. *Appl. Phys. Lett.* 100, 113510 (2012).

“Fluctuations from edge defects in superconducting resonators”, C. Neill, A. Megrant, R. Barends, Yu Chen, B. Chiaro, J. Kelly, J. Mutus, P. OMalley, D. Sank, J. Wenner, T. C. White, Yi Yin, A. N. Cleland, John M. Martinis. *Appl. Phys. Lett.* 103, 072601 (2013).

“Coherent Josephson qubit suitable for scalable quantum integrated circuits”, R. Barends, J. Kelly, A. Megrant, D. Sank, E. Jeffrey, Y. Chen, Y. Yin, B. Chiaro, J. Mutus, C. Neill, P. OMalley, P. Roushan, J. Wenner, T. C. White, A. N. Cleland, John M. Martinis. *Appl. Phys. Lett.* 111, 080502 (2013).

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

“Room temperature deposition of sputtered TiN films for superconducting coplanar waveguide resonators”, S. Ohya, B. Chiaro, A. Megrant, C. Neill, R. Barends, Yu Chen, J. Kelly, D. Low,

J. Mutus, P. OMalley, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, Y. Yin, B. D. Schultz, C. J. Palmstrom, B. A. Mazin, A. N. Cleland, J. M. Martinis. *Supercond. Sci. Technol.* 27, 015009 (2014).

“Characterization and reduction of microfabrication-induced decoherence in superconducting quantum circuits”, C. Quintana, A. Megrant, Z. Chen, A. Dunsworth, B. Chiaro, R. Barends, B. Campbell, Yu Chen, I.-C. Hoi, E. Jeffrey, J. Kelly, J. Mutus, P. OMalley, C. Neill, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, A. N. Cleland, John M. Martinis. *Appl. Phys. Lett.* 105, 062601 (2014).

“Superconducting quantum circuits at the surface code threshold for fault tolerance”, R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O’Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, John M Martinis. *Nature* 508, 500-503 (2014).

“State preservation by repetitive error detection in a superconducting quantum circuit”, J. Kelly, R. Barends, A. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Mutus, B. Campbell, Yu Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. O’Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, John M Martinis. *Nature* 519, 66-69 (2015).

“Dielectric surface loss in superconducting resonators with flux-trapping holes”, B. Chiaro, A. Megrant, A. Dunsworth, Z. Chen, R. Barends, B. Campbell, Y. Chen, A. Fowler, I. Hoi, E. Jeffrey, J. Mutus, C. Neill, P. O’Malley, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, J. Martinis. *submitted* (2016).

Abstract

Simulating Quantum Chemical Dynamics with Improved Superconducting Qubits

by

Anthony E. Megrant

A quantum computer will potentially solve far-reaching problems which are currently intractable on any classical computer. Many technological obstacles have prevented the realization of a quantum computer, the main obstacle being decoherence, which is the loss of quantum information. Decoherence arises from the undesired interaction between qubits and their environment. Isolated qubits have better coherence but are more difficult to control. Superconducting qubits are a promising platform since their macroscopic size allows for easy control and coupling to other qubits. While the coherence of superconducting qubits has substantially improved over the past two decades, further improvements in coherence are required.

We have repeatedly and reliably increased the coherence times of superconducting qubits. Currently decoherence in these devices is dominated by coupling to material defects. These defects are present in the dielectrics used to fabricate these devices or introduced during fabrication. Using simpler resonators as a testbed, we individually isolate, characterize, and then improve each step of the more complicated fabrication of superconducting qubits. We increased the quality factor of resonators by a factor of four by first identifying the surfaces and interfaces as a major source of loss and then by optimizing the substrate preparation. Furthermore, we measure and subsequently mitigate additional defect loss, which is dependent on the position of ground plane holes used to limit the loss from magnetic vortices. Implementing these improvements led to an increase of our qubit coherence times by more than an order of magnitude.

The progress made in coherence while maintaining a high degree of connectivity and controllability has been directly used in more complex circuits. One such device is a fully connected three qubit ring with both tunable qubit frequencies and adjustable qubit-qubit couplings. The considerable level of control allows us to generate the quantum dynamics resulting from a collision between Na and He atoms by encoding the problem into a time-dependent multi-qubit interaction. We compute a scattering cross section by measuring the probability of exciting the Na atom over a wide range of collision parameters. The scattering cross-section shows that there is a velocity where the excitation probability is maximized, in agreement with scattering experiments.

Contents

Chapter 1

Introduction

Computing has led the world into the technological revolution we are experiencing today. Smaller, faster, better have all been used to describe the newest “must-have” in a series of ground breaking technological devices, many of which could not have been imagined 50 years earlier. As we continue down this road of wearable tech that we interact with at every waking moment (and some even monitor us while we sleep), we are always asking “what’s next?”, “what can’t they do?”.

These devices are the product of the continuous improvement in computing power which was predicted by Gordon Moore and is known as Moore’s law [1], which states that transistor density would double every two years and lead to a doubling of performance every eighteen months. Even though Moore’s law was originally a prediction for the subsequent decade, it has held true for nearly 50 years. However exponential growth based on the shrinking footprint of the transistor can not continue forever as these devices can only be reduced so much in size.

Moore's law also holds true for the increase in processing power of supercomputers. The processing power of supercomputers has doubled roughly every fourteen months, since the inception of the Top500 list in 1993, a biannual list which tracks the performance of the world's 500 fastest supercomputers [2]. This performance increase is due not only to improved processors but also due to the increase in the number of processors, with the current state of the art having ~10^6 cores [2].

Major advancements in nuclear test simulations, weather forecasting, and molecular dynamics simulations are due to the incredible increase in processing power. However there are many problems that computer scientists think may never be solvable on any system based on this so-called “classical” computing. The issue here is in the nature of the solution to these problems.

Until recently, every computer sold performed fundamental computations the same way, by flipping switches or bits between ”on” and ”off” (or ”1” and ”0” respectively) to perform calculations. Based on the binary number system, these switching elements allow both arithmetic and logic computations. There have been many physical constructs used for the switching element, but the breakthrough came from the invention of transistors and their eventual use in integrated circuits. Modern computers can be used to evaluate any computable function which will produce the correct output by performing a well defined and finite procedure on a given input from its domain, although the time and space resources may grow exponentially with the problem size.

In contrast to the above classical computers, a computer which utilizes the unique properties of quantum mechanics such as quantum entanglement and quantum superposition will behave quite differently. With a quantum computer the fundamental operations are performed with “qubits”. Qubits are analogous to the classical bit mentioned above, however the qubit can be in a quantum superposition of both ”1” and ”0” or ”on” and ”off” simultaneously. Additionally, entanglement between two individual qubits form a new system which can not be described by the individual qubit states, instead the new system must be described as a whole. The entangled system has 2^2 states and for n qubits this leads to 2^n possible states, where the system can be in a superposition of these states.

Quantum superposition and entanglement provide quantum computers with tremendous computational potential. Underlying this potential is quantum interference. Quantum interference of different computational paths leads the initial state to evolve in such a way that correct outcomes constructively interfere while other possible solutions destructively interfere [3]. Naively this can be considered as a form of parallel computation, since in a sense different solutions are being considered simultaneously. However it is not analogous to a multi-core processor doing multiple calculations at once.

This is not to say that if a quantum computer is built, classical computers will soon after be relics. In fact, classical computing is an integral part of all potential quantum computing architectures. It may be better to think of quantum computers as a performance booster for certain applications. Several main categories of possible applications include: simulating other quantum systems, optimization problems, and quantum circuit based algorithms.

These problems can grow exponentially in computational complexity as the size of the problem increases. The exponential growth of both the processing power and computational complexity of certain problems leads to only a slow growth in problem size capabilities on a classical computer, assuming Moore's law is able to continue. However if the number of qubits in a quantum computer are able to follow a similar trend as Moore's law, classically intractable problems at larger systems sizes may be solvable on a quantum computer.

1.1 Quantum computing: potential applications

1.1.1 Simulating quantum systems

The most intuitive application for a quantum computer is simulating other similar quantum systems [4]. In fact some potential candidates for qubits are actually simulating the behavior of an ideal qubit. However an individual qubit is easy to simulate exactly on a classical computer, so a useful quantum computer should be able to simulate much larger systems that involve many electrons and nuclei, such as chemical reactions with catalysts or high-T_c superconductors.

Catalysts provide a lower activation energy through an alternative reaction pathway. They are used in most industrial manufacturing processing to reduce cost, pollution, and energy use. The catalytic converter in your car's exhaust system converts the incomplete combustion and side products to less harmful molecules such as N_2, O_2, and CO_2. Quantum simulation could allow us to engineer catalysts to perform these functions better or find catalysts for new functions such as removing CO_2 and greenhouse gases from the atmosphere.

High temperature superconductors can potentially reduce loss in transmission lines or may be used in magnetically-levitated trains to reduce rolling friction allowing faster more efficient mass transit. However research is still on going to explain exactly how superconductivity arises in the high-temperature systems. Quantum simulations can be used to study the systems already known to elucidate this mechanism as well as possibly finding higher T_c superconducting systems that are easy to manufacture and utilize.

Currently research conducted in these areas use approximation methods and supercomputers. However current supercomputers limit the exact simulation of interacting quantum spins at around 30-50, depending on the type of problem considered [5, 6, 7]. It is difficult to utilize the full processing power of today's supercomputers to solve these problems using exact diagonalization due to the limited bandwidth between the many nodes of the supercomputer and the limited local memory at each node [7]. Relevant simulations can involve 100's of spins. In chapter 7 I discuss progress toward quantum simulators using current state of the art superconducting qubits.

1.1.2 Combinatorial optimization

Consider a delivery truck driver who has four stops to make across town before heading home for the day. If they know the area well including construction closures, traffic, etc they can likely determine a near optimal route, and if there is a better route available it is likely not much better. Now consider a large multinational delivery company, which delivers to more than six million customers per day across 220 countries and territories. Given the enormous number of possible routes and constraints such as traffic conditions, equipment, fuel, weather, employees, subcontractors, package size, etc., finding a near optimal solution or even a better than current solution is a very difficult task worth many millions of dollars.

There are many similar real world problems facing industry today. Collectively known as combinatorial optimization these problems consist of finding the optimal solution given a very large but finite number of possible solutions. The example above is known as the vehicle routing problem [8] and is a generalization of the Traveling Salesman Problem, which is the most well-known example of combinatorial optimization and has algorithms for both exact and approximate solutions [9, 10, 11].

Combinatorial optimization also has applications in artificial intelligence and machine learning [12]. These related fields use powerful computers to identify known or hidden patterns is very large sets of data. There are many applications for machine learning such as improved voice and optical character recognition, data compression, and even self-driving cars to name a few.

1.1.3 Quantum circuit model

The quantum circuit model decomposes quantum algorithms into a series of quantum gates. The sequence of quantum gates perform a unitary evolution on n-qubit quantum circuits, which began in a prepared initial state. These unitary operations are reversible and ensure that the total probability of all outcomes remains 1. With a universal set of quantum gates one can approximate any other quantum circuit. The previous applications can be solved using the quantum circuit model as well as some other well known quantum algorithms such as Groover's search and Shor's factoring algorithms.

Grover's algorithm searches through possible function inputs and returns the single input where the function is true [13]. It utilizes a technique called amplitude amplification which can lead to a quadratic speedup over the best possible classical algorithm.

Shor's algorithm is probably the most well-known quantum algorithm [14]. It is used to factor an integer into its prime factors and to compute the discrete logarithm. This is in a group of algorithms based on the Quantum Fourier Transform (QFT). Other algorithms using the QFT include the similar quantum phase estimation algorithm and the more general hidden subgroup problem.

1.2 Quantum computing: potential architectures

1.2.1 Universal fault-tolerant quantum computer

In practice the quantum circuit model can require many gates to perform a given computation and any errors which occur during the computation would likely render the results meaningless. These errors can be represented by undesired random gates appearing during the course of the computation. Therefore a physical universal quantum computer must satisfy more than having a universal set of quantum gates.

The five necessary requirements for a physical realization of a universal quantum computer are known as the DiVincenzo criteria [15].

1. a scalable physical system with well characterized qubits

2. the ability to initialize the state of the qubits to a simple fiducial state, such as $|000\ldots\rangle$
3. a qubit-specific measurement capability
4. a universal set of quantum gates
5. long relevant decoherence times, much longer than the gate operation time

The first criterion states that we need a collection of qubit-like quantum elements whose physical parameters are well understood, such as the energy eigenstates of the qubit, the interactions with other possible states outside of the qubit subspace as well as interactions with other qubits, and couplings to external fields to control the state of the qubit. The second criterion states that we should begin a computation in a known, easy to prepare state. The third criterion ensures that we can read out the result of a computation by measuring specific qubits. The forth criterion allows efficient approximation of any quantum circuit. The last criterion states that we need the computer to retain its "quantumness" long enough to use it for computation.

Quantum decoherence is the loss of information to its surrounding environment and is thought to be one of the biggest obstacles to the realization of a useful quantum computer. Qubits that comprise a quantum computer need to be extremely well isolated from unwanted interactions while still remaining highly controllable and allowed to interact with other qubits. These two qualities are typically in direct opposition to one another. Additionally, even if unwanted interactions could be suppressed, errors can arise from imperfect control pulses and measurement errors. Meeting all of these requirements with current realizations of physical qubits is well out of reach. Error correction will almost certainly be requirement for a universal fault-tolerant quantum computer.

Quantum error correction

Redundancy is typically the approach to classical error correction. This works because errors are assumed to be unlikely and a majority vote will provide the desired outcome. However, we are unable to copy quantum information due to the no-cloning theorem, and we are unable to measure the state during a given computation to see if an error has occurred since that would typically preclude us from continuing the computation due to the collapse of the state. So to protect quantum information from errors, we must spread the information across the entangled state of many physical qubits, creating a single logical qubit. We can then measure the parity of physical qubits to detect if an error has occurred. The parity measurement won't collapse the state of the qubit to the measurement basis so the computation can continue.

All quantum error correction schemes place a threshold on the probability of a certain type of error before it fails. Various quantum error detection protocols allow a compromise between the magnitude of the error threshold, number of physical qubits per logical qubits, as well as inter-qubit connectivity. The surface code is an example of a quantum error detection protocol that requires a planar 2-D array of physical qubits which have nearest-neighbor coupling and has a modest error threshold of 1% per operation [16]. In this protocol many thousands of physical qubits per single logical qubit will be required to suppress errors low enough for the large circuit depths required for useful quantum computation using gate based computation.

1.2.2 Quantum computing with pre-threshold qubits

Current qubit technologies have improved dramatically over the past several decades to the point where meaningful quantum computation may soon be available. However a universal fault-tolerant quantum computer is probably still decades away due to the very large number of physical qubits required to suppress errors. In contrast to the digital gate based model discussed above, the following architectures use analog control and are possible using pre-threshold qubits.

Quantum annealing

A very popular topic recently in the literature and even the mainstream news is ”quantum annealing”. This is due to the public relations success of the D-wave quantum annealer made up of ~500 superconducting qubits [17]. The D-wave quantum annealers are the first commercially sold quantum computation systems.

The quantum annealer is 2-D array of superconducting qubits with adjustable coupling between qubits [18]. This creates an artificial programmable Ising spin system believed to be very powerful for certain problems that are of great interest to the business and scientific communities. Some of these applications include condensed matter physics, protein folding, machine learning, traveling salesman, financial forecasting and risk modeling. The devices are engineered and built with these particular problems in mind and can not perform arbitrary quantum computations using them. Instead they are designed to encode and solve these particular problems using an analog control scheme.

The quantum annealer allows the problem of interest to be mapped/encoded on to network of interconnected qubits with adjustable coupling between qubits. The encoding creates a potential energy landscape similar to a landscape of mountains (energy barriers) and valleys (energy minima). The optimal solution to the original problem is contained in the deepest valley and can be found by allowing the quantum state to tunnel through the energy barriers during the annealing process and end in the ground state of the problem.

Single excitation subspace method

In contrast to fairly well established quantum annealer, I will discuss the first implementation of a single excitation subspace (SES) quantum processor in chapters 6 and 7 of this thesis. The SES quantum processor also uses superconducting qubits with adjustable coupling between qubits, however it operates quite differently than the quantum annealer. An SES quantum computer can implement any gate-based algorithm or quantum circuit and is therefore universal [19]. While the SES processor does not require the same level of coherence as a fault-tolerant quantum computer, it does require coherence times 100-1000 times longer than a quantum annealer.

The SES method utilizes the n-dimensional single excitation subspace of the 2^n-dimensional Hilbert space of an n-qubit fully coupled quantum processor (e.g. $|100\rangle$, $|010\rangle$, and $|001\rangle$ for a $n = 3$ qubit device. Both the diagonal and off-diagonal SES matrix elements are directly and independently controlled by the qubit frequencies and the qubit-qubit coupling strengths. This allows many unitary operations to be carried out in a single step.

The SES processor is capable of uniform state preparation, Grover's search algorithm, eigenvalue estimation, Schrödinger equation solver for time-independent and time-dependent Hamiltonians [19]. The SES approach is not scalable since the the operational subspace grows linearly with the number of qubits, although due to the high degree of controllability and parallelization, it is predicted to reach a break even speed with a petaflop supercomputer at 50-100 qubtis for certain time-dependent Hamiltonian emulations. This is due to the simulation time being independent of the number of qubits, n, and is instead dependent on maximum qubit-qubit coupling strengths, problem of interest, and decoherence times. See chapter 7 or Ref. [19] for a more through discussion.

1.3 Potential qubit candidates

The ideal attributes of a qubit candidate are different depending on its use in either a prethreshold analog quantum simulator or in a universal fault-tolerant digital quantum computer. Although several key features are important for both. The Divicencio criteria discussed above cover these attributes for both systems, except a universal set of gates is not required for an analog quantum simulator.

Potential qubit candidates can be effectively divided into two categories: 1) systems which display good individual qubit control and readout [trapped ions [20, 21], electronic spins (quantum dots) [22, 23], nuclear spins (NMR) [24, 25], superconducting qubits [26, 27, 28], and photons (linear optics) [29]] and 2) systems which have been shown to scale to large numbers of quantum elements [neutral atoms [30], superconducting qubits [17]]. Superconducting qubits are the only subgroup of candidates which have shown both attributes [17] as well as moderately long coherence times relative to gate lengths [27, 28] and the ability to suppress errors through detection circuits [28]. Although these key features have not been shown simultaneously, a quantum computer based on superconducting qubits currently appears to be the most promising candidate.

1.4 Conclusion and outlook

The processing power of classical transistor based computers appears to be slowing below the exponential growth predicted by Moore's law. However even with an exponential growth of processing power, certain classes of problems are believed to be intractable on any classically based computer. In the long-term, a universal fault-tolerant quantum computer provides an alternative route to solving some of these problems, but first robust error-correction is required. Before this is achieved, improvements in coherence, controllability and scaling to larger numbers of qubits provide a path for meaningful quantum simulations and solutions to difficult optimization problems.

The remainder of this thesis is focused on improving the coherence, controllability, and connectivity of superconducting devices for quantum computation and then using these improvements to simulate quantum chemical dynamics. In chapter 2, I discuss the energy-loss mechanisms of superconducting devices. The improvement of the materials used to fabricate these devices is described in chapter 3. In chapter 4, I move to a new material system to further improve coherence and look toward more advance architectures. Chapter 5 focuses on energy loss due to the motion of magnetic vortices. In chapter 6, I discuss how these improvements have directly benefited superconducting qubits by highlight several experimental results. Finally in chapter 7, I use these improved qubits to simulate the inelastic collision between Na and He atoms over a large range of energies and calculate the inelastic cross section and compare it to previous collision experiments in chapter 7.

Chapter 2

Superconducting devices and energy loss

2.1 Decoherence

The decoherence of a qubit is described with two characteristic time scales, the rate of energy decay $1/T_1$ and the rate of dephasing $1/T_2$, coming from NMR spectroscopy [31]. The energy decay time T_1 is the time it takes for the state to decay from an excited state to a lower energy state. The pure dephasing time T_ϕ is a measure of the system to lose the information stored in the phase of the quantum state. However energy dissipation causes dephasing. The total dephasing time is characterized by the relation $1/T_2 = 1/(2T_1) + 1/T_\phi$ [32]. Both $1/T_1$ and $1/T_2$ describe the rate at which information is lost from quantum systems to their environment [31, 32].

A significant limitation on the dephasing times of superconducting qubits is from frequency noise, which has a characteristic $1/f^\alpha$ power spectral density at low frequency, where $\alpha{\sim}1$ in contrast to white noise where $\alpha = 0$. The source of this noise is not well understood and is still being actively debated [33]. Several possible sources of this noise are critical current fluctuations, charge fluctuations, and magnetic flux fluctuations [34]. Critical current fluctuations are fluctuations caused by trapped electrons in a tunnel barrier. Charge fluctuations are caused by the hoping of electrons between traps on the surface of the superconductor or substrate which induces charges on the surface of a nearby superconductor. Magnetic-flux fluctuations may arise from dangling bonds at the surface of the superconductor or substrate. An alternative source of magnetic-flux fluctuations has recently been identified as weakly bound adsorbates such as O_2 [35]. More recently, we measured a previously unreported telegraph noise source [36]. This noise source dominates at short time scales, such as our single and two qubit gate times of 20 ns and 40 ns, respectively. The source of the telegraph noise is also not known. In the remainder of this chapter, I focus on sources of energy loss.

2.1.1 Quality factor and participation ratio

There are many sources of energy loss for superconducting devices. A measure of how well a device can retain energy or information is the quality factor Q. The quality factor is the total energy stored in the system E_{tot} divided by the energy lost per cycle $Q = \omega E_{tot}/P$, where ω is the angular frequency at which the total energy and power lost P are measured. Q is related to the energy decay time T_1 by $Q = T_1\omega$. A measure of the loss is also given by the loss tangent, $\tan(\delta) = 1/Q \sim \delta$, for large Q. Independent loss mechanisms add together to give the total loss of the system $\delta_{total} = \delta_{external} + \delta_{internal}$, where $\delta_{internal}$ is the sum of the losses that are present even if the device was separated from the measurement circuit and $\delta_{external}$ is the loss of the device to the measurement circuit. We are interested in measuring and lowering $\delta_{internal}$ which is typically dominated by only one or two sources in each regime of operation. In section 2.4 I discuss loss mechanisms that are currently the dominating sources of loss in at least one regime of operation.

How much does a specific loss channel contribute to the total loss of a resonator or qubit? Intuitively the answer must depend on how much of the total energy is exposed to the that particular loss mechanism. For example, the measurement circuit is a source of loss for a device. By changing the coupling strength or the amount of the device's energy stored in the measurement circuit, we effect the total measured loss of the system. The participation ratio is the energy stored in a particular loss mechanism compared to the total energy of the device. In general the total loss of a device is the some the various loss channels weighted by their participation ratio. Therefore we can reduce the loss of a given channel by either improving its loss tangent or by decreasing it participation.

2.2 Superconducting resonators

High quality factor microwave resonators provide critical elements for superconducting electromagnetic radiation detectors [37], quantum memories [38, 39], and some quantum computer architectures [40]. They are also used to test material properties [41, 42] over a wide range of conditions, where they are susceptible to many of the same decoherence mechanisms as other superconducting devices. There are many varieties of superconducting resonators, each with their own advantages and disadvantages.

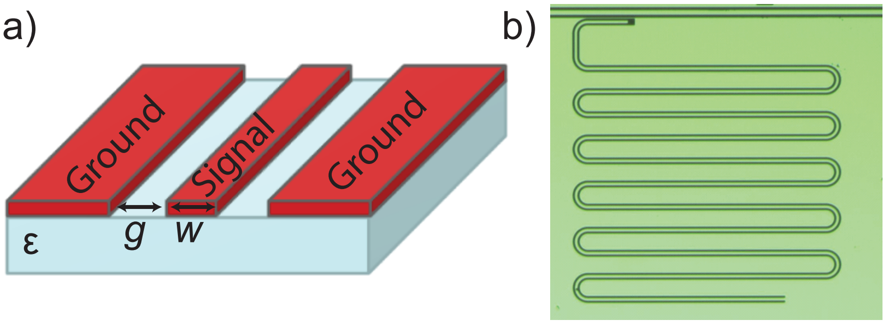

Figure 2.1: (a) Perspective drawing of a coplanar waveguide (CPW) fabricated on a dielectric substrate with permittivity ϵ. The center conducting strip has width w and is separated from the ground planes on either side by a gap g. (b) Micrograph of a single $\lambda/4$ CPW resonator with $w = 6\mu$m and $g = 2w/3$. The resonator is capacitively coupled to a CPW feedline with coupling strength $Q_c = 2 \times 10^5$.

Resonators with the largest quality factors are typically produced using very large macroscopic objects such as 3-D superconducting resonant cavities ($Q > 10^8$ [43, 44]) or whispering gallery modes of single crystal dielectrics (Q~10^8 [45]). One major drawback of these resonators is they are orders of magnitude larger than typical 2-D planar resonator designs, which hinders their use when scaling to a many resonator system such as quantum memory used with qubits [40].

A seemingly more scalable approach are resonators that can be fabricated on chip using standard semiconductor integrated circuit processing methods. Their size on chip is still an important consideration for these resonators, as well as performance in the desired operating regime and the complexity of the fabrication process.

Superconducting coplanar waveguide resonators (SCWR) provide an ideal test device in lieu of more complex superconducting devices such as qubits. SCWR only require a single lithography step in their simplest form and can be easily tested in both the low power regime (corresponding to an average photon number in the resonator of $\langle n_{photon} \rangle \sim 1$) where qubits are operated and the high power regime ($\langle n_{photon} \rangle > 10^6$) where MKID's are operated. As seen in Fig. 2.1(a) a coplanar wave guide consists of a center strip of conductor of width w, with corresponding ground planes on either side which are separated by a gap g. We typically use $\lambda/4$ SCWR which consist of an electrically open end which is capacitively coupled to a second CPW used for readout. The other end of the SCWR is electrically shorted to the ground plane [Fig. 2.1(b)]. This approach allows for coupling many resonators to a single feedline by spacing the resonators out in the frequency domain, frequency multiplexing [46], which simplifies the experimental setup. Also by measuring many resonators on a single chip we can collect a larger number of statistics in nearly identical conditions. Typically the resonators are meandered in a serpentine fashion [Fig. 2.1(b)] to reduce extraneous coupling to other resonators.

A natural question to ask is, "Can we use superconducting resonators as qubits?" We can answer this question by referring back to the DiVincenzo criteria discussed in 1. The first criterion states that we must have well defined qubits. This means that if more than two energy levels are present we must be able to individually address the transition between only these two levels without exciting any other transitions. However this is not the case with the harmonic oscillator potential of the superconducting resonators which has equal spacing between energy levels. In other words, if we desired to excite the resonator from the ground state $|0\rangle$ to the first excited state $|1\rangle$ we could apply a tone at the transition frequency and place one quantum of energy into the resonator. However the resonator would be in the coherent state $|1\rangle$. In the

Fock or number basis there would be a superposition of states $|n\rangle$ with a Poisson distribution of probability of measuring the $|n\rangle$ Fock state. By breaking the degeneracy between energy levels we could individually address the transitions and choose two energy states for the qubit subspace.

2.3 Xmon qubits

The Xmon qubit was specifically designed to meet the requirements of the surface code for quantum error detecting [16]. It is a transmon style superconducting qubit[47] which has much longer coherence times compared to the phase qubits previously used at UCSB. It is a weakly anharmonic oscillator which consists of a planar capacitor and a nonlinear inductor in the form of a Josephson junction for fixed frequency devices. We can tune the frequency of the device by substituting a superconducting quantum interference device, SQUID, for the Josephson junction. The DC SQUID consists of a superconducting loop interrupted by two Josephson junctions. By threading flux through the SQUID loop we can tune the effective critical current, and thus the inductance of the qubit.

Currently, the 3-D transmon has better coherence over planar devices [48]. It consists of a very large planar qubit placed in a high quality factor 3-D cavity similar to those discussed in the previous section. By making the qubit structure very large it reduces the participation of the lossy surfaces leading to the improved results. Another important result is the Josephson junction can be made very low loss by using small area junctions.

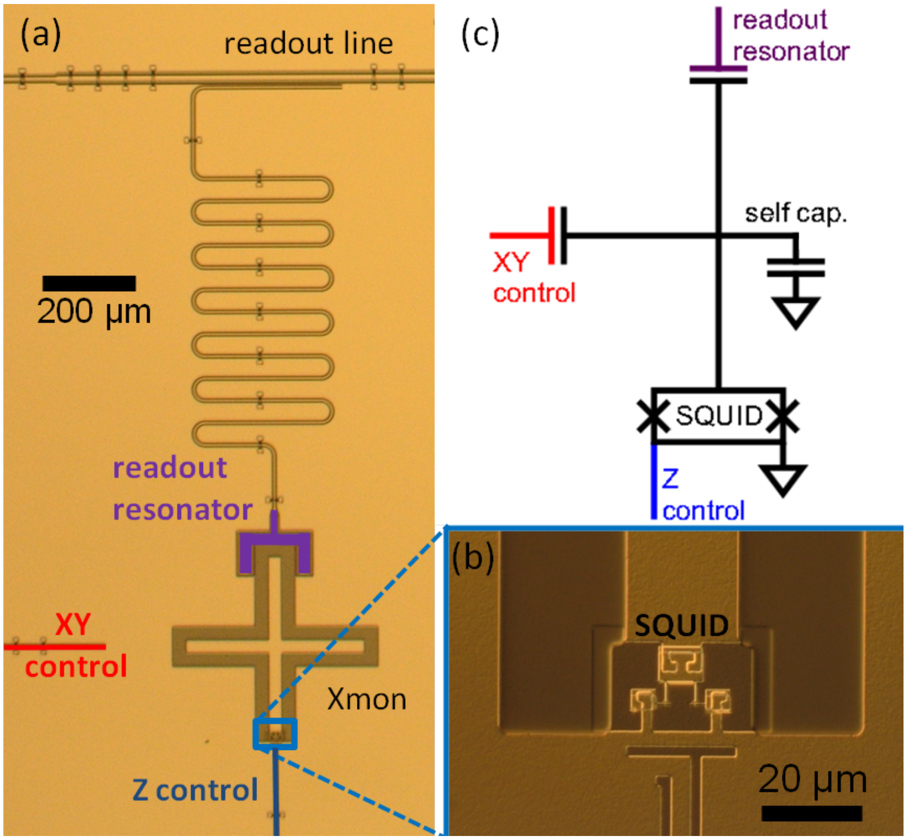

Figure 2.2: (a) Optical micrograph of an Xmon qubit using differential interference contrast. The exposed Si substrate (dark) is revealed after etching device structures into the Al film (light). The Xmon is capacitively coupled to a $\lambda/4$ readout resonator (top), a XY control line (left), and inductively coupled to a Z control line (bottom). (b) Zoom in of SQUID loop and Z control bias line. (c) Corresponding electrical circuit of qubit, control, and readout.

2.3.1 Xmon capacitance

The capacitance required (C ~80 fF) is easily achieved using reasonably-sized (250 μm x 250 μm) planar structures, as seen in Fig. 2.2a. Transmon qubits traditionally use interdigitated capacitors which are electrically isolated from the ground plane of the device. Xmon qubits instead use strips of intersecting coplanar waveguides for the capacitive element, each having a length of 250 μm. This provides several desired features such as direct comparison of the quality factor with CPW resonators as well as distributing the capacitance over a larger area. The latter allows for more room to route control wires, selective capacative coupling to neighboring qubits, and to reduce crosstalk between these elements and stray coupling to other elements in the circuit. The total CPW length of the Xmon capacitor is roughly 10 percent of the length of a $\lambda/4$ CPW resonator at 6 GHz. Many of the loss mechanisms understood and improved with resonators should directly enlighten how to build a better Xmon capacitor.

2.3.2 Josephson junctions

The majority of the Xmon inductance originates from the Josephson effect [49, 50]. The Josephson effect also provides the nonlinearity required for individually addressing the transitions between the Xmon's energy levels. A Josephson Junction (JJ) is created when two superconducting regions are separated by a 'weak link'. The weak link between the superconducting regions can be a narrow constriction between the superconductors, a thin section of a normal metal or insulating barrier. The Xmon JJ are fabricated using a thin (~1 nm thick) AlO_x insulating barrier, see appendix A for details on their fabrication.

The presence of the weak link allows the phases of the superconducting condensate in the two superconducting regions to differ by δ. The current, I, and voltage ,V, originating from Cooper pairs tunneling through the weak link of the JJ are related to phase difference, δ, according to the Josephson relations, [49, 51]

$$I = I_c \sin\delta \tag{2.1}$$

$$V = \frac{\Phi_0}{2\pi}\frac{d\delta}{dt}, \tag{2.2}$$

where I_c is the critical current of the junction, and $\Phi_0 = h/2e \approx 2.07 \times 10^{15}$ Weber is the superconducting flux quantum, with h the Planck constant, and e the charge of an electron.

The physical structure of the Xmon JJ resembles that of a capacitor rather than an inductor. However we can further manipulate Eq. 2.1 and Eq. 2.2 to show that the JJ is indeed an inductive element. We first differentiate Eq. 2.1 and substitute $d\delta/dt$ with V according to Eq. 2.2, allowing us to rewrite in the familiar form

$$V = L_J\frac{dI}{dt}, \tag{2.3}$$

with Josephson inductance L_J defined as

$$L_J = \frac{\Phi_0}{2\pi I_c \cos\delta}. \tag{2.4}$$

Eq. 2.3 not only explicitly displays the inductive nature of the JJ but it clearly reveals its non-linearity in the $1/\cos\delta$ dependence. The inductance can became very large as $\delta \to \pi/2$ and can also become negative. This non-linearity provides the key element missing from linear resonators of the previous section which allows transitions between energy levels to be individually addressable permitting a multiple energy level system to be operated as a qubit.

It is interesting to note that while the JJ behaves as an inductive element it does not have an associated magnetic field. Instead the energy is stored in the kinetic inductance of cooper pairs. Kinetic inductance arises from the inertia of charge carriers resisting the change of current flow. The effect of kinetic inducance can be as much as four orders of magnitude greater than the geometric inductance [52].

The Xmon utilizes two JJ in parallel, as seen in Fig. 2.2(b) creating a DC SQUID. By threading magnetic flux through the SQUID loop we can tune its inductance from a minimum value dependent on the critical current of the junctions to very large values. This provides the ability to tune the frequency of the qubit over a wide range of values, allowing optimal placement in frequency space and turning on or off the interactions to other circuit elements. The typical size of an Xmon JJ is 250 nm $\times 350$ nm with an inductance of L_{SQUID}~10 nH for the SQUID. The inductance of the SQUID and the capacitance of the Xmon provides a maximum transition frequency from ground to the first excited state of $f_{10} = 5.6$ GHz and an impedance of ~$350\ \Omega$.

2.3.3 Xmon in a circuit

The ground and first excited states of the Xmon serve as our qubit subspace. The next two crucial elements of a useful qubit are control and measurement. The Xmon takes a departure from the more typical transmon qubit approach that embeds the electrically floating qubit in a resonant cavity [47, 48]. The interaction between the cavity and the qubit causes a state dependent shift of the cavity frequency. The cavity also allows manipulation of the state of the typical transmon qubit.

We separate the control and readout in the Xmon design. Since the Xmon is electrically connected to ground, we physically separate the qubit and readout resonator. Additionally we use a one-to-one coupling scheme of readout resonators to Xmons, allowing for individual readout of each qubit meeting the requirements of both the DiVincenzo criterion [15] and the surface code [53, 27]. We capacitively couple the Xmon to the readout resonator. In order to achieve the large coupling required ($g/2\pi = 40$ MHz - 100 MHz) for fast readout, we use a large claw structure to surround one leg of the Xmon, as seen in Fig. 2.2(a). Xmon control is achieved using a separate CPW drive line with ~50 aF coupling capacitance. We vary the Xmon frequency using a galvanic connection between the SQUID and the DC bias line, creating a large inductive coupling of $M = 2.2$ pH. This allows rapid modulation of the qubit frequency on nanosecond timescales. Fixed capacitive coupling between neighboring Xmons of $g/2\pi = 15$ MHz allows for fast two qubit gate times (40 ns) [27]. This interaction is turned off by detuning the qubits by Δ~800 MHz. In chapter 6, I analyze the T_1 spectra of a linear chain of nine coupled Xmons.

2.3.4 gmon qubits

Building upon the successful control and coherence of the Xmon qubits, gmon qubits add electrically adjustable qubit-qubit coupling [54]. The addition of tunable inter-qubit coupling as well as the tunable qubit frequency permits many additional applications beyond those of the Xmon. Some applications include analog quantum simulation as well as the possibility of higher fidelity two qubit gates. The optical micrograph in Fig 2.3(a) shows two Xmon style qubits with an adjustable coupler between them. The coupling network is achieved by adding a linear inductor between ground and the SQUID loop, as seen in Fig. 2.3(b). This linear inductor is then coupled to a superconducting loop containing a Josephson junction. A second qubit is coupled to the same loop in a similar fashion. This allows a small fraction of the first qubit's current to be redirected through the second qubits inductor. This coupling scheme allows $g/2\pi$ to vary between $+4$ MHz to -50 MHz going smoothly through zero, seen in the spectroscopy plot in Fig. 2.3(c). The linear inductor is placed strategically at a low voltage point to reduce the participation of dielectric loss from the linear inductor as well as in the coupler structure. The voltage divider between L_g and L_J reduces the participation by a factor of 2000 in this design. In chapter 6, I discuss the T_1 spectrum of a 3-gmon quantum processor which is used to simulate the collision between Na and He discussed in chapter 7.

2.4 Energy loss mechanisms

2.4.1 Radiation and wiring loss

Radiation loss is a concern particularly for MKID's as they are operated at high powers where this may be the dominant source of loss. For resonators, choice of geometry can reduce unwanted radiation [55]. Properly designed sample mounts significantly reduce coupling to parasitic modes; suspending the chip over a cavity reduces coupling to possible parallel plate modes

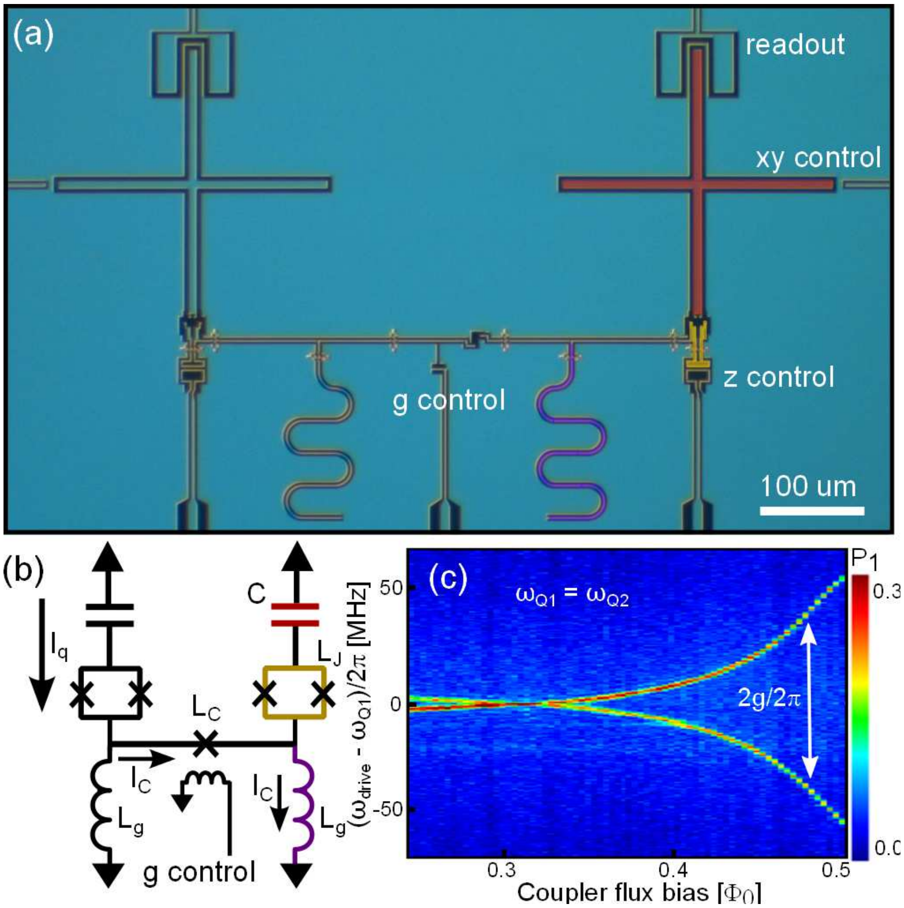

Figure 2.3: This figure reproduced from Ref. [54]. (a) Optical micrograph of two inductively coupled gmon qubits. Building upon the design of the Xmon, a linear inductor L_g is added between the qubit SQUID and ground, shown schematically in (b). This allows a small fraction of the current of the qubit excitation to be routed to the second qubit. A Josephson junction located between the qubits controls the strength of the interaction.(c) Spectroscopy showing coupling strength of two on resonance coupled gmon qubits as the coupler bias is varied.

which can occur if the chip is mounted directly onto a metal mount [56]. By choosing a sample mount which encloses the fields in a smaller volume, such as a sample box with a lid, the radiation loss is also reduced [57, 56].

Xmons and resonators can also couple to the many control lines of a superconducting circuit. Each of these lines provides a decoherence channel for the device. The coupling capacitance between these elements has traditionally been found using numerical solvers, however this can be very computationally expensive when a large chip (1 cm × 1 cm) is used. More recently a simple formula was developed which calculates the planar coupling and self-capacitance for these devices when embedded in a ground plane [58],

$$C_{12} = (\epsilon/\pi) \int \int dA_1 dA_2 / |r_1 - r_2|^3, \tag{2.5}$$

where ϵ is the average dielectric constant between the substrate and vacuum, A_1 and A_2 are the areas of the two planar objects, and $|r_1 - r_2|$ the separation between the differential elements. This formula allows for much faster calculations of the parasitic capacitance to large objects that are far away, such as bond pads.

The qubit control and readout CPW lines break up the ground plane of the chip, potentially forming slotline modes that can couple to the resonators and qubits. Suppressing these modes requires crossovers to electrically short the two ground planes together. A simple way to achieve this is using wirebonds. However wirebonds can potentially be too long when using a manual wirebonder, leading to a large impedance of 20-30 ohms at 6 GHz [59]. Since they can be much shorter, fabricated crossovers are much more effective at eliminating slotline modes.

2.4.2 Inductor loss

Magnetic vortices

For type II superconductors, one of the defining characteristics of superconductivity, the Meissner effect, does not hold above a magnetic field H_{c1}, where $H_{c1} < H_c$, with H_c the thermodynamic field. Instead magnetic flux can pierce through the superconductor creating Abrikosov vortices. Each vortex contains a quantum of flux equal to $\Phi_0 = h/2e$ and are surrounded by screening currents confining the majority of the flux within a quasinormal core with radius $r_c \sim \xi$, where ξ is the superconducting coherence length [60]. The formation of vortices allows superconductivity to remain up to $H_{c2} > H_c$. The vortices repel each other and form a lattice in ideal defect free materials, which is either square or triangular depending on the film thickness and bulk superconducting properties [61], with lattice spacing $a \sim \sqrt{\Phi_0/B}$, where B is the magnetic flux density . Typically these lattice structures are rarely observed due to pinning of the vortices at defect sites in the superconductor, causing only a short range glass-like order [62]. When a transport current passes a single vortex it feels a Lorentz force [60] $\mathbf{f} = \mathbf{J} \times \mathbf{B}/c$. If this force is larger than the pinning force, the vortex will move causing dissipation.

Quasiparticles

When Cooper pairs are broken, two quasiparticles are formed. Quasiparticles cause dissipation due to scattering, in the same manner as normal metal electrons in a transport current. There are several sources of pair breaking mechanisms, some of which include: temperature, magnetic fields, large currents, incident photons, and magnetic impurities [60]. For finite temperatures there exists a thermodynamic equilibrium quasiparticle density described microscopically by Mattis-Bardeen theory [63]. Fortunately their density decreases exponentially with decreasing temperature, so if cooled to below $T_c/10$ their contribution to the loss is largely mitigated. However even at this temperature non-equilibrium quasiparticles can be a significant source of loss. The physics behind large populations of non-equilibrium quasiparticles is not completely understood, however a recent experiment [64] has shed some light on at least one of their sources. Incident photons with energies above the superconducting gap Δ break Cooper pairs creating a large local non-equilibrium quasiparticle density, analogous to valence electrons being excited to the conduction band by photons with energies above the bandgap of a semiconductor.

Stray radiation can find its way through what may be considered "light tight" designs. Previously reported short decoherence times originally believed to be caused by material properties or design flaws may have been an unknown measurement of this radiation. The "quick and dirty" method to solve this problem is to encapsulate the entire sample box with Eccosorb, a microwave absorbing mixture [65]. However Eccosorb is filled with iron filler and thus can lead to large magnetic fields at the sample. A permanent solution to this problem is by using a "box-in-box" design where the inner surface of the outer box is covered with a granular microwave absorbing coating, which effectively traps photons reducing their number by an order of magnitude [64]. For example in our adiabatic demagnetization refrigerator experimental set-up, see appendix C, the 4 K blackbody radiation entering our light shield is reduced significantly by the absorber, while the light shield itself is only emitting 50 mK radiation. This shielding design increased the quality factor of the resonators tested from 5×10^4 to 5×10^5 and the T_1 time of a phase qubit increased from 200 ns to 600 ns which corresponds to a quasiparticle density decrease from 400 μm^{-3} to 170 μm^{-3} [64].

2.4.3 Capacitor loss

Two-level states

At very low temperatures and energies the dielectrics used in superconducting circuits begin to show much larger losses than would otherwise be expected [66, 41]. This increase in the dielectric loss is consistent with the device coupling to a bath of two-level-states (TLS) located in the dielectrics of the device [66]. A TLS is modeled as a material defect having two different energy states in a double well potential. A subset of these defects posses a dipole moment that can couple to the electric fields of superconducting devices. The interaction of the device with the TLS can lead to energy loss and frequency noise. At high powers and temperatures the TLS saturate and their energy losses go unnoticed, however they may still cause dephasing. At low energy and temperature the TLS are all in the ground state, except for when they are excited by the coupling to the device and absorb some energy.

Dielectrics that have high coordination numbers, such as single crystal sapphire and silicon, typically will have a larger splitting in energy levels and therefore higher transition frequencies as well as a larger barrier for the the transition to occur. Even when very low defect density crystalline dielectrics are used in devices, losses are still present which have the characteristic power dependence of TLS. The loss is consistent with our simulations of interfaces between both the high quality dielectrics and superconductor and the vacuum-dielectric interface are equally like to cause dissipation. The superconductor-vacuum interface has 100 times less participation.

While the previous loss mechanisms can be substantially reduced through careful circuit design and proper shielding and filters, the same can not be said about the inherent nature of TLS defects. This loss needs to be tackled head on with systematic materials characterization and improvements to reduce the density of these defects and by reducing the participation of these defects. Both approaches are explored in the following three chapters of this thesis.

Chapter 3

Planar superconducting resonators with internal quality factors above one million[1]

3.1 Introduction

High quality factor microwave resonators provide critical elements for superconducting electromagnetic radiation detectors [37], quantum memories [38, 39], and quantum computer architectures [40]. Good performance and stability can be achieved for such applications using aluminum resonators patterned on sapphire substrates. Aluminum is a favored material due to its robust oxide and reasonable transition temperature (1.2 K), and sapphire provides an excellent substrate due to its very low microwave loss tangent[45] $\delta \sim 10^{-8}$ and its chemical inertness. However, the quality factors measured in such resonators are lower than expected;

---

[1] The following was published as: "Planar superconducting resonators with internal quality factors above one million", Anthony Megrant, et al. *APL* 100, 113510 (2012).

recent simulations [67] and experiments [68] suggest that the unexplained loss arises mostly from imperfections at the metal-substrate interface. Using an experimental apparatus with minimal stray magnetic fields and infrared light at the sample [64], here we show that careful substrate preparation and cleaning yields aluminum-on-sapphire resonators with significantly higher internal quality factors Q_i. We also introduce a method for evaluating the resonator microwave response.

3.2 Deposition systems and substrate preparation

The aluminum for the resonators was deposited on c-plane sapphire substrates in one of three deposition systems: A high vacuum direct current (DC) sputter system (base pressure $P_{\text{base}} = 6 \times 10^{-8}$ Torr), a high vacuum electron-beam evaporator ($P_{\text{base}} = 5 \times 10^{-8}$ Torr) or an ultra-high vacuum (UHV) molecular beam epitaxy (MBE) system ($P_{\text{base}} = 6 \times 10^{-10}$ Torr) with electron-beam deposition. The sapphire substrates were sonicated first in a bath of acetone then isopropanol followed by a deionized water rinse. For the sputter-deposited and e-beam evaporated samples, we further cleaned the substrates prior to Al deposition by Ar ion-milling. For the MBE-deposited samples, we first cleaned the substrates with a load-lock outgassing at 200°C followed by heating to 850°C, in UHV or in ~10^{-6} Torr of activated oxygen (O_2^*) generated from a radio-frequency plasma source. After substrate treatment, ~100 nm of Al was deposited at room temperature. Specific process parameters are given in Table 3.1.

Table 3.1: Sample process information; w is the resonator center stripline width, f_0 the resonant frequency, and Q_i-H and Q_i-L the internal quality factors at high power (before over-saturation) and low power ($\langle n_{photon} \rangle \sim 1$), respectively.

Process[a]	*In vacuo* cleaning	w (μm)	f_0 (GHz)	Q_i-H $\times 10^6$	Q_i-L $\times 10^6$
A) Sputter	900 eV Ar$^+$ mill for 2 min	3	3.833	4.3	0.16
		15	6.129	4.5	0.40
B) E-beam	200 eV Ar$^+$ mill for 10 sec	3	3.810	9.9	0.66
		15	6.089	4.4	0.72
C) MBE	None	6	4.973	5.70	0.53
		15	6.120	4.33	0.76
D) MBE	LL^2 anneal	3	3.773	6.58	0.75
		15	6.125	5.38	0.80
E) MBE	LL^2 & 850°C anneal	3	3.876	10.1	1.15
		15	6.127	6.4	0.92
F) MBE	LL^2 & 850°C anneal in O$_2^*$ ^{3}	3	3.767	12.7	1.1
		15	6.121	8.5	1.72

[a] All films deposited at room temperature.

[b] 200°C anneal in load lock ($\sim 10^{-6}$ Torr).

[c] Activated oxygen (O$_2^*$) generated from a radio frequency plasma source.

3.3 Aluminum film characterization

The substrate cleaning process changes the morphology of the deposited Al film, as shown in Fig. 3.1: No cleaning yields a rough Al film and a diffuse polycrystalline reflected high-energy electron diffraction (RHEED) pattern, while a high-temperature sapphire anneal gives a much smoother film and sharp crystalline RHEED. X-ray diffraction scans of all films show strong sapphire substrate peaks as well as an Al(111) peak [Fig. 3.1(c)], with the MBE-grown samples displaying stronger Al peaks. MBE-grown samples heated to 850°C with or without O$_2^*$, and MBE samples that underwent only a 200°C anneal (not shown), had RHEED, atomic-force microscopy, and x-ray diffraction scans that were nearly indistinguishable from one another.

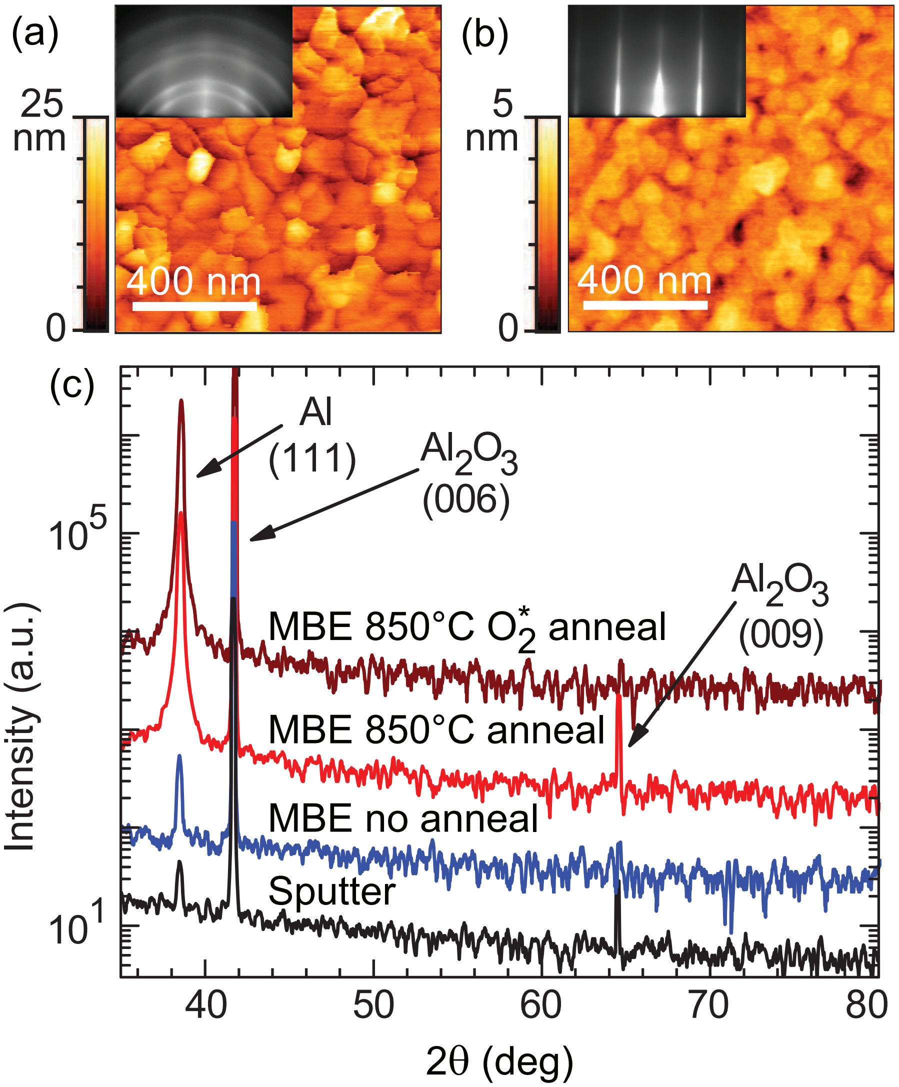

Figure 3.1: (Color online) (a) MBE-deposited Al with no pre-deposition substrate cleaning yields a rough Al surface, shown in the atomic force microscope (AFM) surface scan with ∼25 Å root mean square (RMS) roughness and (inset) a diffuse polycrystalline RHEED pattern, compared to (b) MBE-deposited Al with an 850°C anneal in O_2^*, with a much smoother film (∼4 Å RMS roughness) with sharply streaked RHEED. (c) $\theta - 2\theta$ x-ray diffraction scans of sputtered and MBE-grown films. Vertical scale corresponds to sputtered sample, with other samples plotted on same scale but vertically offset for clarity. All samples display substrate peaks and an Al(111) diffraction peak.

3.4 Fabrication and measurement setup

We patterned resonators from the Al films using an optically-patterned resist and etching in a 300 W BCl_3/Cl_2 inductively-coupled plasma. This step defined sets of 12 quarter-wave coplanar waveguide resonators capacitively-coupled to a single central transmission line [64]. Resonators were designed with one of three center trace widths: w = 3 μm, 6 μm or 15 μm. The center trace gap to the ground planes on either side was $g = 2w/3$. The resonator coupling to the central transmission line corresponded to one of the coupling strengths [46] Q_c = 5×10^4, 2×10^5, 5×10^5 or 1×10^6; resonator frequencies f_0 ranged from 4 to 8 GHz.

After etching, we stripped the resist and diamond saw-cut the wafers into 6×6 mm^2 dies, placing individual dies in an Al sample box and wiring with 25 μm diameter. Al wire-bonds. We mounted the sample box on the 50 mK stage of an adiabatic demagnetization refrigerator, enclosed in a copper light-tight shield, surrounded by a magnetic shield at 4 K, as described in Ref. [64]. We note that all screws and coaxial connectors were made of non-magnetic materials. The use of non-magnetic materials inside of the 4 K magnetic shield reduced the residual magnetic field at the sample from ~100 mG to 3 mG, which we verified using cryogenic tests similar to Ref. [39].

3.5 Extracting the internal quality factor

We measured the transmission coefficient S_{21} of the 50 mK resonators with a vector network analyzer (VNA). The cryostat microwave input cable had 40 dB attenuation at room temperature and 30 dB attenuation at the 4 K stage, followed by 50 mK microwave powder filters [64] with ~1 dB attenuation before the connection to the sample's central transmission line. The sample's output passed through a second 50 mK microwave powder filter followed by a circulator at 4 K, which protects the resonators from noise originating in the high electron mobility transistor (HEMT) amplifier (noise temperature ~4.5 K) also located at 4 K. A typical normalized transmission spectrum is shown in Fig. 3.2(a), displaying a dip in $|S_{21}|$ and a step in the phase of S_{21} near the resonance frequency f_0. Note the slight asymmetry about f_0 in both the magnitude and phase, which we attribute to a small impedance mismatch in the central transmission line on either side of the resonator[69], likely originating from the wire-bond connections [70], sample mount imperfections, or the transmission line geometry. Analyzing a circuit that includes small in-line complex impedances ΔZ_1 and ΔZ_2 on either side of the resonator [Fig. 3.2(b)], the transmission S_{21} is given by

$$S_{21} = \frac{2V_2}{V_1} = \frac{2Z_0}{Z_1 + Z_2} \frac{1}{1 + Z/2Z_r}, \tag{3.1}$$

where $Z_0 = 50\ \Omega$ is the cable impedance, $Z_1 = Z_0 + \Delta Z_1$, $Z_2 = Z_0 + \Delta Z_2$, $1/Z \equiv 1/2Z_1 + 1/2Z_2$, and the impedance Z_r of the resonator in series with its coupling capacitor is given by [46], for frequencies f near resonance,

$$Z_r = \frac{Z_0 Q_c}{2Q_i}(1 + i2Q_i\delta x), \tag{3.2}$$

with $\delta x = (f - f_0)/f_0$, and $Q_{i,c}$ are the internal and coupling quality factors, respectively. The nearly frequency-independent pre-factor in Eq. (3.1), of order unity, is absorbed in an off-resonance transmission calibration of the cabling and amplifiers, leaving us with a normalized

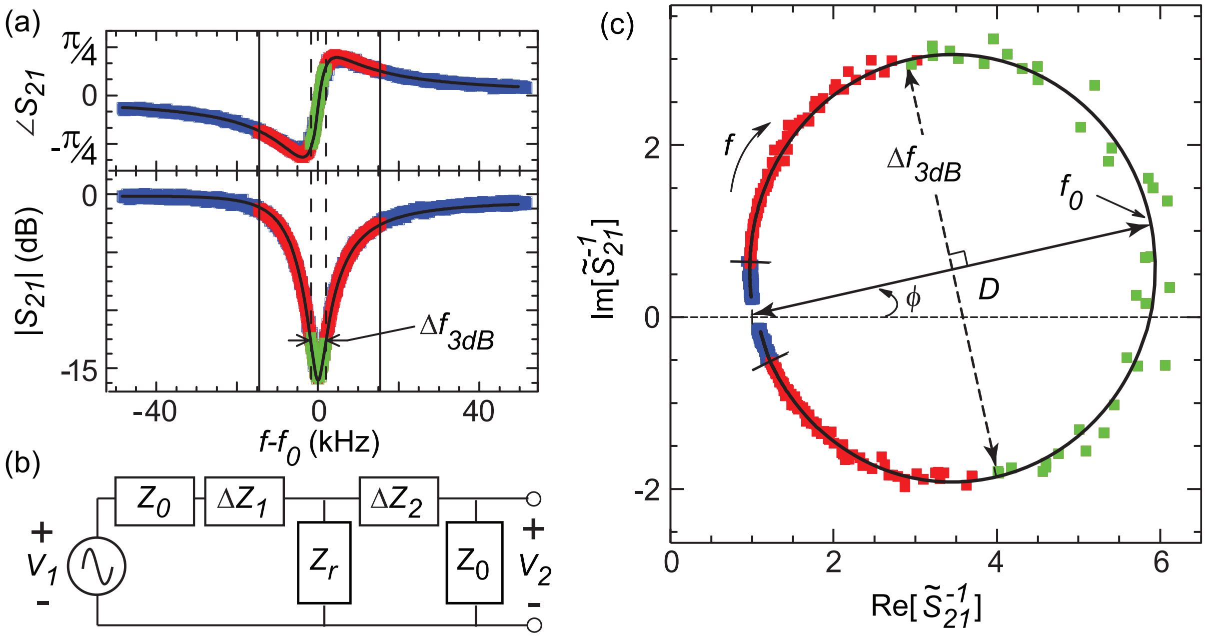

Figure 3.2: (Color online) (a) Measured normalized transmission magnitude $|S_{21}|$ and phase $\angle S_{21}$ (colored squares), with a fit to Eq. (3.3) (solid black line) which yields $Q_i = 1.7 \times 10^6$, $Q_c^* = 4 \times 10^5$, $f_0 = 6.121$ GHz and $\phi = 11.7°$. The resonator had a $w = 15\ \mu$m center stripline width and was patterned from an MBE-deposited Al film on a sapphire substrate that was annealed at 850°C in O_2^*. The data was taken at 50 mK at low power with $\langle n_{photon} \rangle \sim 1$ in the resonator. (b) Circuit diagram including mismatched complex impedances ΔZ_1 and ΔZ_2 for the transmission line input and output, Z_r includes the resonator and its coupling capacitance to the transmission line, which has characteristic impedance $Z_0 = 50\ \Omega$. (c) Parametric plot and fit of Im$[\tilde{S}_{21}^{-1}]$ vs. Re$[\tilde{S}_{21}^{-1}]$ of the same data and fit as (a). Points in the frequency range between the dashed lines (green squares) in (a) correspond to the points to the right of the $\Delta f_{3\mathrm{dB}}$ dashed line in (c).

inverse transmission $\tilde{S}_{21}^{-1}$ given by

$$\tilde{S}_{21}^{-1} = 1 + \frac{Z}{2Z_r} = 1 + \frac{Q_i}{Q_c^*} e^{i\phi} \frac{1}{1 + i2Q_i\delta x}, \tag{3.3}$$

with the magnitude of $Z = |Z|e^{i\phi}$ absorbed in the re-scaled coupling quality factor $Q_c^* \equiv (Z_0/|Z|)Q_c$.

In Fig. 3.2(c) we display the data in Fig. 3.2(a) in the inverse transmission $\tilde{S}_{21}^{-1}$ complex plane. In this representation the resonator response is a circle starting and ending at $\tilde{S}_{21}^{-1} =$

1, achieved for frequencies far from resonance. At resonance $f = f_0$, $\tilde{S}_{21}^{-1} = 1 + De^{i\phi}$ is diametrically opposite $\tilde{S}_{21}^{-1} = 1$, with diameter $D = Q_i/Q_c^*$ and the circle center rotated from the real axis by the impedance mismatch angle ϕ. The two frequency points ±90° around the circle from the resonance frequency are the $\tilde{S}_{21}^{-1}$ 3 dB points, with their frequency difference Δf_{3dB} the full width at half maximum (FWHM) of $|\tilde{S}_{21}^{-1}|$, yielding the internal quality factor $Q_i = f_0/\Delta f_{\mathrm{3dB}}$. Equation 3.3 and the plot of Fig. 3.2(c) not only allow the internal quality factor Q_i to be determined directly, these also emphasize how to properly measure and visually scrutinize the data.

3.6 Dependence on surface preparation and deposition method

The low-power Q_i's of the 15 μm center trace width resonators [Fig. 3.3] range from 4×10^5 for the sputtered film with ion mill substrate cleaning to 1.7×10^6 for the MBE-deposited resonator shown in Fig. 2. We find that as the measurement microwave power is increased [Fig. 3.3], the internal quality factors remain roughly proportional, with Q_i for most resonators increasing by about a factor of ten between high and low powers. The increase in Q_i with power is consistent with loss being dominated by interfacial and surface two-level-states (TLS)[67, 71].

We discovered that for the highest quality factor resonators at the lowest powers, the internal quality factor can fluctuate in time by as much as 30% over a period of several hours, a variation significantly larger than the statistical uncertainty of the fit Q_i. We believe this is due to fluctuations in the population of TLS [72]. The values for Q_i reported here are long-term averages rather than the highest transient values, and are representative of many different samples.

Most significantly, we see a systematic dependence of Q_i on the substrate cleaning and deposition process. We attribute this dependence to changes in the aluminum-sapphire interface, in agreement with simulations [67]. The sputtered film resonators had the lowest Q_i for all powers, and are comparable to the published literature [39, 71, 69]. We believe the lower Q_i measured here is due to the pre-deposition ion mill, which may remove most surface contaminants, but may also bury some contaminants as well as leave a damaged substrate surface [73]. The ion mill used for e-beam evaporation was lower in energy and duration than that for the sputtered films, and produced resonators with Q_i's comparable to those deposited by MBE with no substrate anneal.

We obtained better resonator quality factors with MBE-grown Al with pre-deposition annealing, consistent with better surface properties as shown by the RHEED and atomic force microscopy (AFM) [Fig. 3.1]. However, the very best quality factors were achieved with high temperature (850°C) anneals in O_2^*, yielding significantly higher Q_i even though RHEED and AFM are nearly identical to other MBE-annealed samples. These results are consistent with microscopic theories attributing TLS to surface-bound hydroxyl groups [74], which can saturate the sapphire (0001) surface [75], and remain even after annealing at 1100°C in UHV [75, 76]. Carbon and hydrogen have also been observed on sapphire after similar annealing processes [76]. The highly reactive activated oxygen used here may remove such surface-bound molecules, more effectively than a simple anneal, and with less damage than with ion milling, although we have not verified these hypotheses.

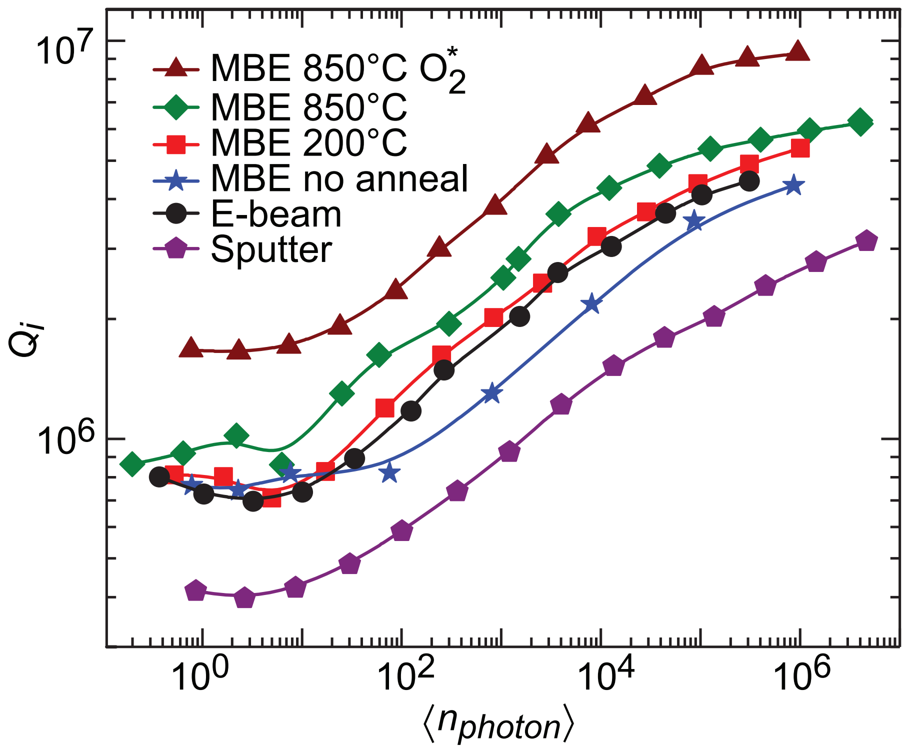

Figure 3.3: (Color online) Power dependence of the internal quality factor Q_i versus average photon number in the resonator $\langle n_{photon} \rangle$, for resonators with $w = 15\ \mu$m. Lines are guides to the eye. The typical low-power statistical error from a least-squares fit of Eq. (3.3) is ~3%, smaller than the symbol size.

3.7 Etching sapphire

Previous surface loss simulations using a finite element solver have shown that trenching down into the substrate in the CPW gap can reduce the loss by a factor of 2 or more[67]. For common superconducting CPW resonator widths, the surface loss is dominated by the surrounding 1 μm of lossy surfaces on either side of the superconducting edges [72]. By etching into the substrate in the CPW gap we reduce the divergence of the electric field at the corner of the edge and reduce the magnitude of the electric field at the substrate air interface, which has been moved further from the superconductor edge.

Trenching CPW resonators on silicon substrate has been shown experimentally to improve the quality factor by a factor of 2x consistent with the above simulations[77, 78]. There are many methods available to etch into silicon with different selectivity to various superconductors. However there is much less literature on the etching of sapphire [79, 80] and no literature on etches with exposed aluminum present. We have developed a dry etch which uses a modified version of our more standard BCl_3/Cl_2 aluminum dry etch we discussed earlier.

In order to etch the much more inert sapphire substrate while not damaging the sidewalls of the covering Al layer we need to create a more physical etch. First we include a larger ratio of BCl_3/Cl_2, making it 1:1 instead of the previous 1:2. This is because pure Cl_2 etches aluminum but doesn't etch Al_2O_3, while BCl_3 is used to remove the native oxide of Al as well as scavenge any water present on the wafer or chamber. We increase the ICP power from 300 W to 900 W and reduced the pressure in order to increase the ion content in the plasma and to create a more physical etch for the hard sapphire substrate. Finally we choose an etch time of 90 s and

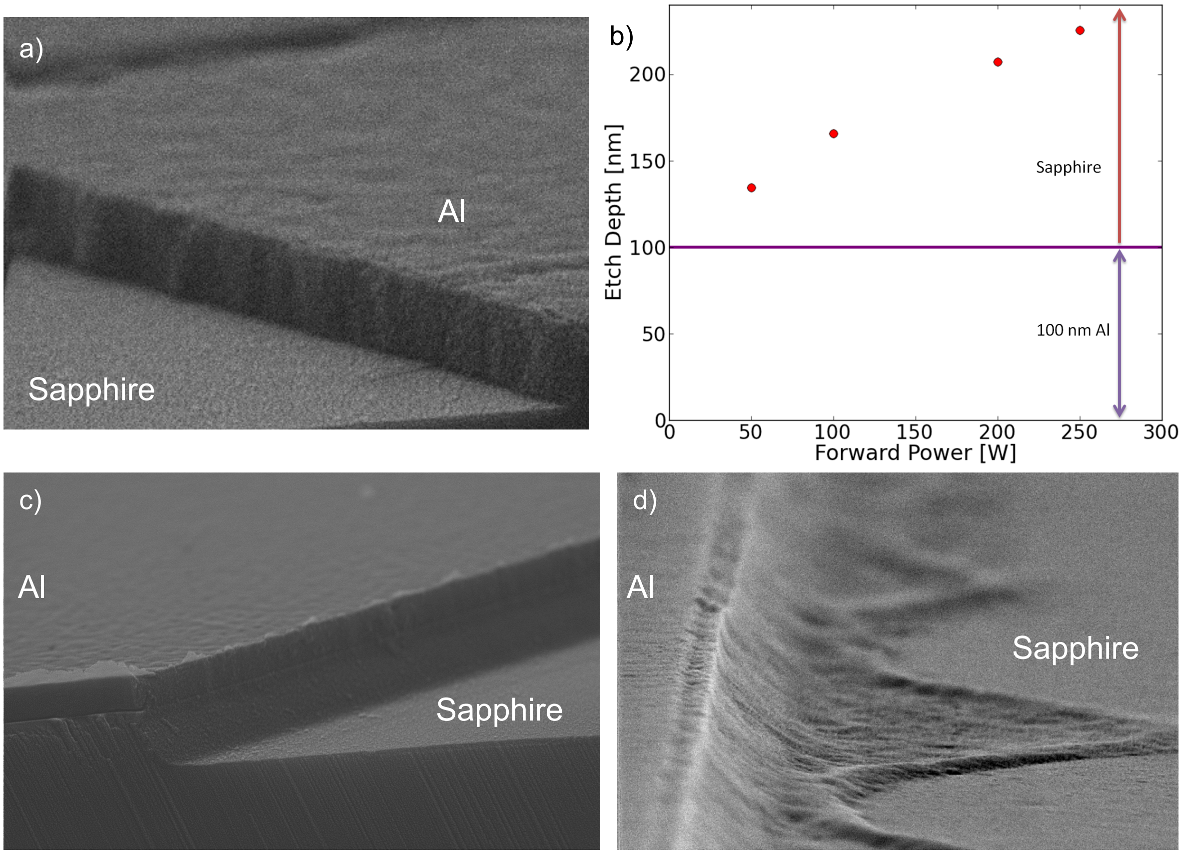

Figure 3.4: (a)Scanning electron micrograph after our standard aluminum 300 W BCl_3/Cl_2 etch using an inductively coupled plasma reactive ion etcher, see A. (b)Characterization of the new sapphire etch, showing the total etch depth as a function of substrate bias forward power for a total etch time of 90 s. Scanning electron micrographs of new sapphire etch for (c)with little resist residue on edges (d) significant reflow of resist into the gap. (d) was found to be more likely to occur for higher forward power and wider gaps.

perform a sweep in the power of the RF bias on the substrate (forward power), seen in Fig. 3.4(b). We see a slightly nonlinear relationship between the magnitude of the forward power and etch depth into the substrate. At the largest forward power we see an etch depth of 225 nm into the sapphire substrate.

Amazingly this much stronger, more physical etch doesn't appear to damage the aluminum sidewalls, as seen in Fig. 3.4(c), compared to the standard Al etch, Fig. 3.4(a) The slope of the trench sidewall depended on the substrate bias forward power. At a forward power of 50 W we find a slope of ~50 degrees, while at 250 W the sidewall is more vertical at 80 degrees. However it appears that the wider the opening in the resist is the more likely it is to develop the issue seen in Fig. 3.4(d). It appears that during the etch the photoresist heats up enough causing it to reflow into the gap, this problem becomes worse at larger forward power.

Next we fabricate CPW resonators to see if this new etch can improve the quality. We deposit 100 nm of Al using electron beam deposition and then dice the wafer into quarters and use optical lithography to expose the same resonator circuit device discussed in the preceding sections. For the first quarter we perform our standard Al etch to define the resonators. On a separate quarter we our new sapphire etch described above with a forward bias power of 100 W and perform a total etch time of 3 min, for a total trench depth of ~150 nm into the sapphire substrate. We then segement the dies, wirebond and cooldown the resonators in the same manner as described above. After reaching T ~50 mK we measure the power dependence of both devices, Fig. 3.5.

We find the typical quality factors for the resonators fabricated using the standard Al etch,

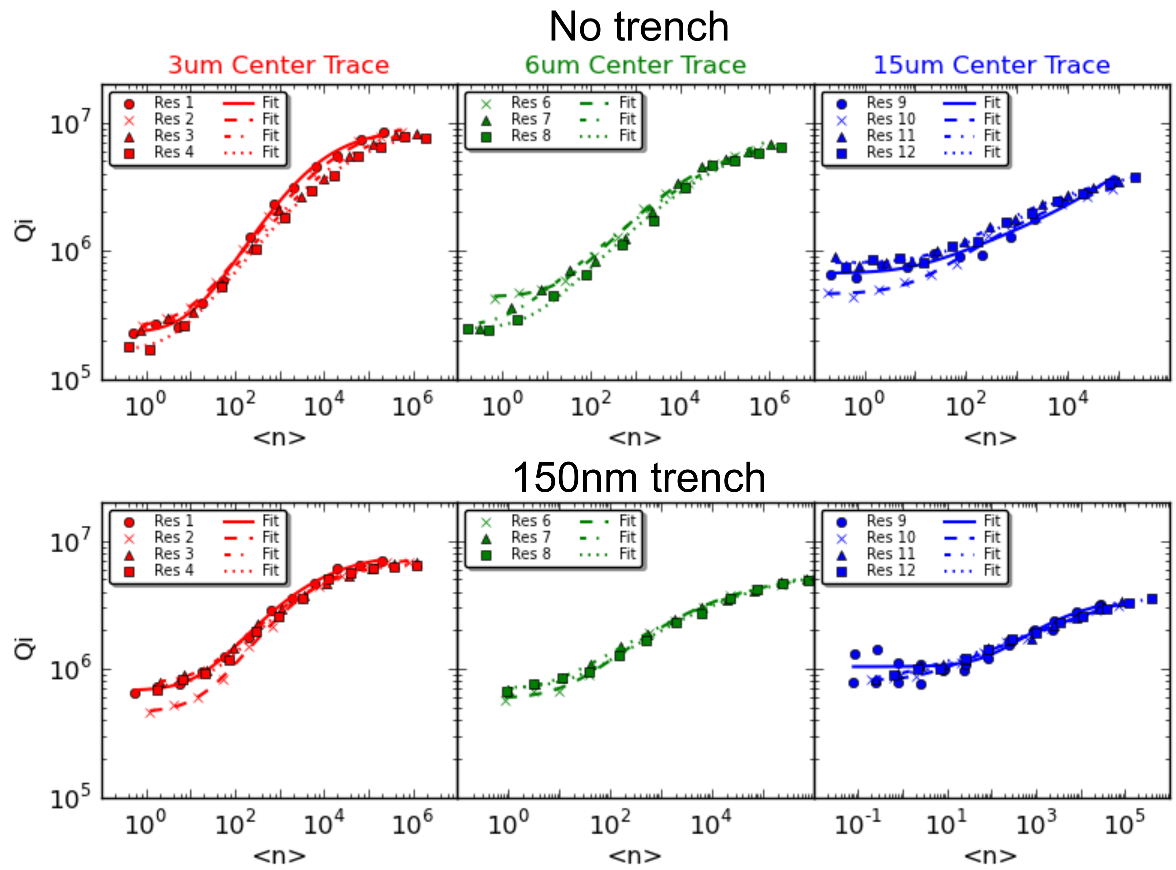

Figure 3.5: Comparison of the CPW width and power dependence of the internal quality factor for two nominally identical resonator devices except for the etch used to fabricate each. The 'No trench' device uses our standard BCl_3/Cl_2 aluminum etch, while the '150 nm trench' device uses the newly developed sapphire etch described in the text. Notice that $w = 3\ \mu$m center trace resonators see almost a factor of three improvement at low drive power while the $w = 15\ \mu$m resonators only slightly improved.

with the $w = 15$ μm and $g = 10$ μm resonators having an average low drive power quality factor of Q$_i$~7×10^5. This is slightly lower than the resonators discussed previously due to the under-performance of resonator 10. For resonators of the same width and 150nm trench we find a slightly improved quality factor of Q$_i$~9×10^5. Similarly at high drive power there appears to be little difference due to the trenching, with both devices having quality factors of Q$_i$~4×10^6. However for the narrower width resonators there is much more significant improvements.

The control resonators with $w = 3$ μm and $g = 2$ μm have low drive power quality factors of Q$_i$~2.3×10^5, while the same width trenched resonators have Q$_i$~6.6×10^5, almost a factor of three improvement. The improved quality factor is now larger than those achieved with the same size NbTiN resonators on silicon which used a much deeper trench depth of 0.9 μm and a more complex reactive sputtering deposition compared to the simple elemental superconductor deposited by electron beam evaporation done here [77]. We measure a similar improvement from trenching with the $w = 6$ μm and $g = 4$ μm resonators, although not quite as large, with the trenched resonators having a factor of 2x improvement over the controls. Both the $w = 3$ μm and $w = 6$ μm trenched resonators had slightly reduced high drive power quality factors compared with the controls.

It is unknown how the resist residue, seen in Fig. 3.4(c) & (d), plays a role in the quality factors, and may be the cause of the discrepancy between the relative improvements for the various width resonators. Improvements to this etch need to be made before more systematic studies can be performed to determine the cause. Breaking the etch into several steps did not improve the etch results significantly, however switching to a more stable resist or using a hard mask should improve the etch results.

3.8 Conclusion

In summary, we have measured the internal quality factors of superconducting aluminum resonators while varying the substrate preparation, deposition method and measurement power. The largest internal quality factors, $Q_i \sim 1.7 \times 10^6$ at single photon energies, are achieved for resonators, which have $w = 15$ μm and $f_0 \cong 6$ GHz, where we cleaned the substrate surface as thoroughly as possible without damaging the underlying crystal. The same process also produces the best high-power $Q_i \sim 12.7 \times 10^7$, but for resonators with $w = 3$ μm and $f_0 \cong 4$ GHz. Additionally we developed a dry etch capable of etching both the Al film and sapphire substrate while not damaging the exposed Al sidewalls. This improved the quality factor by almost a factor of three compared to the control resonators for the $w = 3$ μm resonators.

Chapter 4

Anomalous temperature dependence of resonator internal quality factor

4.1 Introduction

In the previous chapter, I showed that the combination of Aluminum superconductor and sapphire dielectric substrate was an excellent combination of materials useful for high coherence superconducting circuits. The highest quality factors are achieved when proper care is taken to ensure that the interfaces are as clean and ordered as possible. However, the sapphire substrates may limit the possibility of making more complex and sophisticated devices due to its chemical inertness. While these properties allow relatively straightforward processing procedures to achieve high quality factor devices, they limit the use of relatively easy fabrication techniques such as trenching and undercutting into the substrate or more complex techniques such as through substrate vias (TSV).

In this chapter we will explore the use of the most prevalent substrate available, silicon. A superconducting qubit based quantum computer would ideally be based on the silicon substrate in order to expedite its creation by utilizing the vast knowledge and experience of the integrate circuit community. This is apparent in the literature available on the use of silicon including advanced processing techniques [81, 82], chemical interactions [83, 84], surface studies [85, 86], and has already been shown to be a high quality substrate material with other superconductors such as TiN [87] and NbTiN [77] with internal quality factors. However previous experiments with aluminum on silicon have not produced resonators with comparable quality factors [88, 71]. Here we report Al-on-silicon resonators with $\mathrm{Q}_i > 5 \times 10^6$ at low drive powers where, on average, a single photon is circulating in the resonator. Additionally, we measure a significant increase in the quality factor as the temperature is decreased below T $\lesssim$ 60 mK.

4.2 Al deposition and substrate preparation

The aluminum for the CPW resonators was deposited on high resistivity Si (100)((111)) with $\rho > 10$ kΩ-cm($\rho >$ 5kΩ-cm) in an ultra-high vacuum (UHV) molecular beam epitaxy (MBE) system ($P_{\text{base}} < 5 \times 10^{-11}$ Torr) using electron-beam deposition. Before depositing the aluminum metallization, we thoroughly cleaned the substrate surface, which has been shown to substantially improve Q_i of Al-on-sapphire coplanar waveguide resonators (CPW) [42]. The substrates were sonicated first in a bath of acetone then isopropanol followed by a deionized water (DI) rinse. Next the wafers were soaked in heated Nanostrip for 10 min at 70° C, followed by DI water rinse. We then strip the native oxide on the wafers by soaking in buffered oxide etch (BOE) for 1 min, followed by a final DI water rinse before blow drying with dry N_2. After loading the substrate into the MBE we outgass the samples at 200° C in the load-lock for 30 min ($P_{\text{base}} < 1 \times 10^{-8}$ Torr), followed by annealing in UHV for 5-10 min at 400°C - 600°C.

The buffered oxide etch (BOE) removes the native oxide covering the silicon substrate leaving a hydrogen-terminated H-Si(100)((111)) surface [89]. The H-Si surfaces are relatively stable in atmosphere [90] and have previously been shown to be a lower loss interface than native or thermal oxides [77, 41]. Upon heating in UHV, the H-Si surfaces begin to reconstruct at $400°$ C $\leq$ T$_{sub}$ $\leq$ $550°$ C, consistent with literature [91]. The reconstruction of the Si surface ensures that the majority of possible contaminates have been removed leaving an ultra-clean and an atomically smooth surface [92]. After the substrate treatment, 100 nm of Al is deposited at room temperature at a rate of 0.2 Å/s - 0.5 Å/s using electron beam evaporation. The background pressure during deposition would increase to $P \sim 10^{-8}$ Torr.

4.3 Surface reconstruction and thin film characterization

We continuously measure the Si surface during the UHV anneal using a Reflective High Energy Electron Diffraction (RHEED) system. As T$_{sub}$ increases from room temperature up to the reconstructing temperature, T$_R$, the RHEED pattern along the Si<110> direction sharpens and decreases intensity, Fig. 4.1(a) & (d) for Si(100) & Si(111) surfaces respectively.

For the Si(100) surface, when T$_{sub}$ $\geq$ T$_R$, we see the formation of additional diffraction peaks centered between the original primary peaks (Fig. 4.1(b)), indicative of a (2x1) sur-

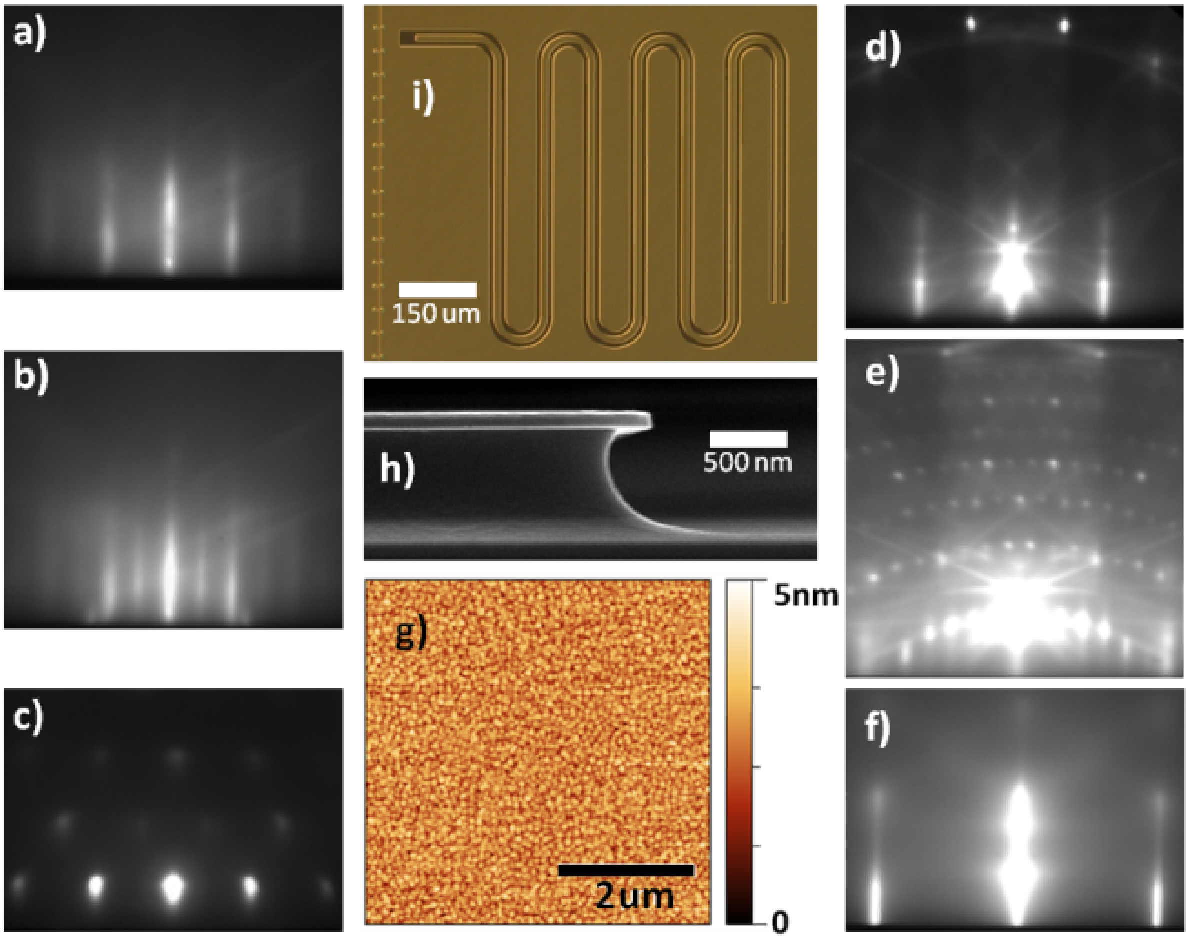

Figure 4.1: Reflection High Energy Electron Diffraction (RHEED) patterns along <110> after a 200° C load-lock outgassing of hydrogen terminated (a) H-Si(100) and (d) H-Si(111). (b) and (e) RHEED along (a) and (d), respectively, after 600° C UHV anneal for 1 minute. (c) and (f) 100 nm Al film deposited at room temperature on (b) and (e) respectively. As the Si temperature is increased and $T_{sub} < T_R$ the intensity of (a) and (d) decreases and the RHEED streaks become sharper, the later likely due to further outgassing of adsorbed contaminants. When $T_{sub} = T_R$, additional diffraction peaks appear between the original primary peaks due to the Si surface reconstructing to (b) Si(100)-(2x1) and e) Si(111)-(7x7) surfaces. g) Atomic force micrograph of the Al film deposited on e). The film has an root mean square roughness of 4 Å. h) A cross-sectional scanning electron micrograph of the 100 nm Al film on Si(100) where the Si has been etched using BCl_3/Cl_2 followed by a SF_6 etch forming a trench and undercut in the Si. i) Differential interference contrast micrograph of the final Al-on-Si CPW $\lambda/4$ resonator capacitively coupled to CPW feedline for readout.

face reconstruction. The (2x1) reconstruction involves neighboring rows of Si atoms along the Si<110> direction forming dimer bonds to minimize the surface energy. We observe a double (2x1) reconstruction where reconstruction occurs along the two orthogonal Si<110> directions.

For the Si(111) surface, we see surface reconstructions over a similar temperature range as with Si(100). Above the reconstruction temperature we measure the well studied (7x7) surface reconstruction, Fig. 4.1(e) [91]. Both Si reconstruction RHEED patterns increase in intensity as the sample cools back down to room temperature where both reconstructions remain stable independent of cooling rate.

The reconstructed Si surface significantly changes the morphology of the room temperature deposited 100 nm thick Al film, as seen in Fig. 4.1(c) & (f) for deposition on Si(100) & Si(111) surfaces respectively. The Al film on the Si(100) reconstructed surface exhibits a textured growth RHEED pattern Fig. 4.1(c) while on the Si(111) it shows a twinned epitaxial film of Al(111), similar to the Al-on-sapphire deposition discussed in chapter 3. Aluminum films deposited on non-reconstructed Si(100) surfaces at room temperature have RHEED patterns consisting of polycrystalline rings with a more diffuse background (not shown). After the deposition, we perform atomic force microscopy (AFM) on the surfaces of the Al films. For the 100 nm Al films deposited on reconstructed Si(111) surface, Fig. 4.1(g), we measure very smooth films with a root mean square (RMS) roughness of 4 Å. For the textured Al films deposited on reconstructed Si(100) surfaces (AFM not shown), we find an order of magnitude greater RMS roughness of 5 nm.

4.4 Resonator fabrication

In Fig. 4.1(i), we see an optical micrograph of a $\lambda/4$ CPW resonator that we fabricated from the Al films using an optically patterned resist followed by etching in a 300 W BCl_3/Cl_2 inductively coupled plasma reactive ion etcher. This process selectively etches Al more than 3 times faster than Si and produces nearly vertical Al sidewalls with ~80° Si sidewalls. On some devices we include an additional 700 W SF_6 etch, immediately after the BCl_3/Cl_2 etch without breaking vacuum, in order to further trench and undercut the Si substrate. The scanning electron micrograph, Fig. 4.1(h), shows the center trace of a CPW resonator in cross section after the SF_6 etch, which produced a ~600 nm trench into the Si substrate and ~350 nm undercut beneath the Al film. The amount of undercut relative to the trench depth can be tuned by varying the RF bias on the substrate with little change to the trenching etch rate.

These processes define quarter-wavelength CPW resonators capacitively coupled to a single CPW for readout. The individual dies were segmented, mounted into an Al sample box, and measured in an adiabatic demagnetization refrigerator, see Ref. [42] for a more detailed description of these. We measure the temperature of the device using a ruthenium-oxide (RuOx) resister thermometer mounted to the cryostat cold plate. The RuOx thermometer's resistance as a function of temperature is calibrated from room temp down to 50 mK. For lower temperatures we extrapolate the temperature-resistance curve from the provided calibration data using a variable-range hopping model [93] down to T~35 mK, the base temperature of the ADR.

4.5 Extracting the internal quality factor

In order to extract the internal quality factor of resonators we measure the complex transmission coefficient, S_{21}, near resonance with a vector network analyzer (VNA) using two methods. We first characterize the resonators using the more standard S_{21} as a function of frequency and then fit the resonance circle using Eq. 4.1, see chapter 3 or Ref. [42] for more detailed discussion.

$$S_{21}^{-1} = 1 + \frac{Q_i}{Q_c^*}\frac{e^{i\phi}}{1 + i2Q_i\delta x}, \tag{4.1}$$

where $S_{21}^{-1} = 1/S_{21}$, Q_i is the internal quality factor, Q_c^* is the renormalized coupling quality factor to the feedline, ϕ is the rotation angle of the circle center relative to the real axis due to an impedance mismatch across the resonator, and δx is the fractional frequency difference between the drive and resonance frequency, $\delta x = \left(f - f_0\right)/f_0$.

The raw S_{21} data needs to be calibrated first before processing. For this we normalize the raw data using S_{21} measured far off-resonance (typically ± 10 MHz). For measurements at low drive powers where the SNR is worse, we average consecutive scans until the statistical error of the quality factor from a least squares fit is below 5 percent.

After verifying the coupling quality factor and rotation angle are constant in time and for various drive powers, as one would expect, we can rewrite Eq. 4.1 to solve for Q_i and δx independently.

$$\frac{1}{Q_i} = \mathrm{Re}\left\{\frac{e^{i\phi}}{Q_c^*}\frac{1}{S_{21}^{-1} - 1}\right\}, \tag{4.2}$$

$$\delta x = \frac{1}{2}\mathrm{Im}\left\{\frac{e^{i\phi}}{Q_c^*}\frac{1}{S_{21}^{-1} - 1}\right\}. \tag{4.3}$$

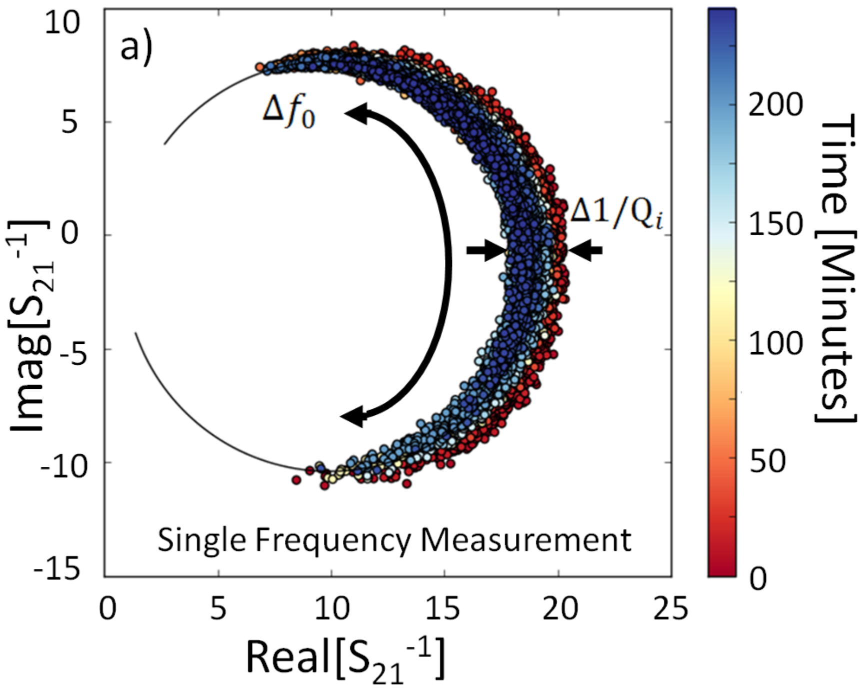

Figure 4.2: This figure reproduced from Ref. [72]. S_{21} measured at a single frequency measurement and plotted on the complex $S_{21}^{-}1$ plane. The time over which this data was measured is shown in false color. For comparison the solid black line is the resonance curve measured as a function of frequency. Due to fluctuations in the resonance frequency the measurements of S_{21} over time will trace out the same resonance curve that was measured as a function of frequency.

Using Eq. 4.2 and Eq. 4.3 we can measure $S_{21}(t, f \simeq f_0)$ as a function of time and at a single drive frequency, near resonance, and map this to the resonance frequency and dissipation. We originally developed this method to measure the power spectral density of the dissipation and resonance frequency noise [72]. Due to the fluctuations in the resonance frequency, measuring S_{21} at a single frequency as a function of time will trace over the same resonance circle measured using the more typical $S_{21}(f)$ method, as seen in Fig. 4.2. Here we plot the resonance curve in the complex S_{21}^{-1} plane measured as a function of frequency (solid black line), and overlay that with S_{21} measured at a single frequency for 4 hours (false color). This provides an excellent sanity check that the time domain measurement and analysis are functioning properly.

4.6 Results

4.6.1 Scaling of quality factor with resonator width

We measured the power dependence of the internal quality factor of resonators with center trace widths $w = \{3\ \mu\text{m}, 6\ \mu\text{m}, 15\ \mu\text{m}, 30\ \mu\text{m}, \text{and } 60\ \mu\text{m}\}$, with a gap, $g = 2/3w$ and also a resonator with $w = 50\ \mu$m, with $w = g$. The design resonance frequencies range from 4.0 GHz to 6.5 GHz. All of the resonators are located on a single chip except for the $w = 60\ \mu$m and $w = 50\ \mu$m resonators which are located on a separate chip containing four resonators of each. Due to the large CPW widths of some of the resonators we performed a field cool measurement to determine and then actively zero the residual magnetic field of B$_{residual}$~6.5 mG (see chapter 5 for more discussion of field cool measurements and loss due to magnetic

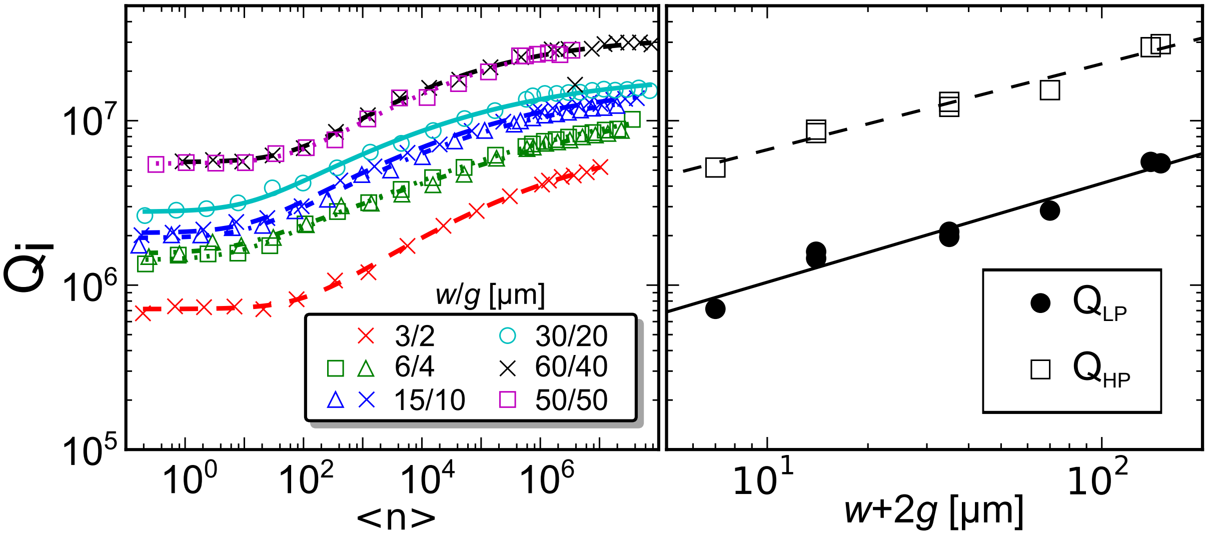

Figure 4.3: Power Dependence for various CPW widths w/g. Here we plot only the resonators with coupling quality factors closest to the internal quality factor of the resonator at low drive power. Note how consistent the data is for multiple resonators of the same or very similar size. The right plot is the quality factor at $\langle n \rangle = 1$, Q_i^{LP}, and at $\langle n \rangle = 10^7$, Q_i^{LP}, as a function of total CPW width $w + 2g$. Both Q_i^{LP} and Q_i^{HP} show a power law scaling with CPW width of $w^{0.60}$ (solid line) and $w^{0.52}$ (dashed line), respectively. Interestingly we see that the quality factor at low and high drive powers scale at nearly the same rate.

vortices).

In Fig. 4.3(a), Q_i is plotted as a function of the average population of photons in the resonator, $\langle n \rangle$. When $\langle n \rangle \geq 10^6$ photons, Q_i saturates to a constant value which is dependent on resonator width. When $\langle n \rangle < 10^6$ photons we measure a decline in Q_i with decreasing $\langle n \rangle$. When $\langle n \rangle < 10$ photons, Q_i plateaus to the zero drive power limit, consistent with the standard tunneling model of TLS [74]. It is interesting to note that over 6 decades of $\langle n \rangle$ the ratio of quality factors between the various CPW widths remains constant, except for $w = 3$ μm. This is similar to the results of chapter 3, were the improvement in low power quality factor due to surface cleans also included a corresponding increase in the high drive power quality factor.

The measurements at low drive powers, with $\langle n \rangle \sim 1$, is the regime most applicable to quantum computing, where a single excitation in qubits are by definition a necessity. In this regime, we find that the internal quality factor of the resonators is strongly dependent on CPW resonator width. For example, the $w = 3$ μm resonator has a low drive power quality factor of Q$_i = 7 \times 10^5$ while the $w = 6\mu$m resonator has Q$_i = 1.5 \times 10^6$. This near one-to-one scaling of quality factor and CPW width is present as the CPW width increases to $w = 60$ μm, with Q$_i = 5.6 \times 10^6$, which if converted to T$_1 = 190$ μs at its resonance frequency of 4.7 GHz. The scaling of the low drive power Q$_i$ with CPW width is consistent with the dominant source of loss arising from lossy surfaces which contain TLS.

A vertical slice through Fig. 4.3(a) at $\langle n \rangle = 10^7$, Q_i^{HP}, and $n = 1$, Q_i^{LP}, as a function of total CPW width, $w + 2g$, is plotted in Fig. 4.3(b). Both Q_i^{LP} & Q_i^{HP} exhibit a power law dependence on CPW width of $Q_i^{LP} \propto w^{0.60}$ and $Q_i^{HP} \propto w^{0.52}$, Fig. 4.3(b) solid and dashed lines, respectively . Surface loss simulations of the CPW cross section predict $Q_{TLS} \propto w^{0.86}$ in good qualitative agreement to what is measured here. However the power law of the CPW width dependence of both Q_i^{LP} & Q_i^{HP} are likely lower bounds on the actual scaling. For Q_i^{HP} it is quite difficult to perfectly zero the magnetic field in a cryostat which uses large magnetic fields as part of its cooling strategy, as is the case with an ADR, so the large CPW width resonators are likely still suffering from non-negligible loss from magnetic vortices. We find Q_i^{LP} to be significantly dependent on temperature. Since these measurements are performed in an ADR which uses a single shot cooling cycle, the temperature of the fridge is continuously, although very slowly, warming up. The data in Fig. 4.3 was taken in the temperature range of 36 mK to 41 mK over a several hour long period.

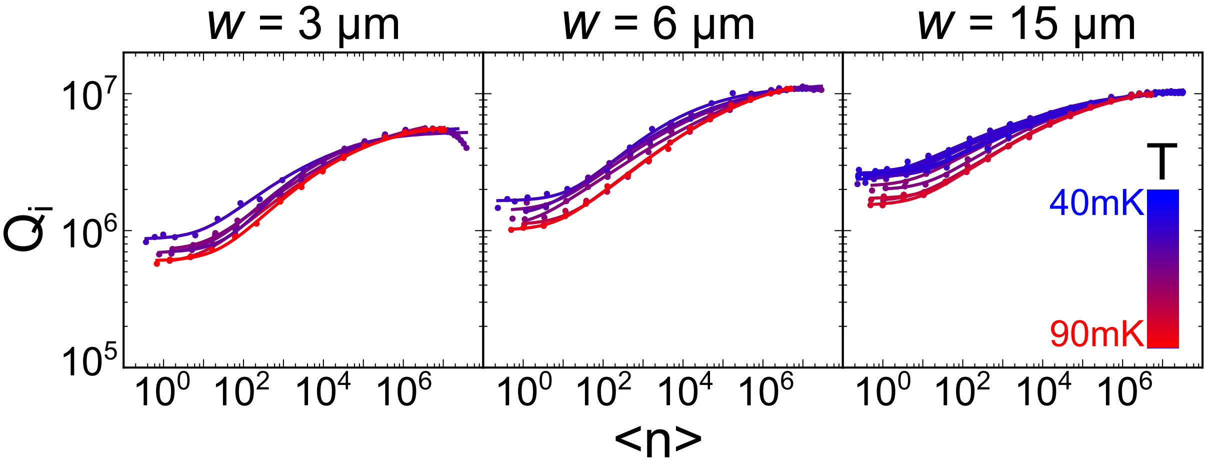

Figure 4.4: Power dependence of internal quality factor for $w = 3$ μm, 6 μm, and 15 μm resonators over the temperature range 40 mK to 90 mK, shown as false color. Note how the quality factor at low photon number increases significantly as the temperature is reduced while the quality factor at high photon number remains unchanged in this temperature range.

In order to further explore this anomalous temperature dependence we repetitively measure the power dependence of $w = 3$ μm, 6 μm, and 15 μm resonators from a different wafer, located in a separate but nominally identical ADR cryostat as the temperature of the cryostat slowly warms from 40 mK to 90 mK (false color) over a ~24 hour period, Fig. 4.4. The $w = 15$ μm resonator has a high power $Q_i = 8 \times 10^6$ independent of temperature, T. As the drive power is reduced over 6 decades, Q_i decreases to 2.7×10^6 at 40 mK (1.7×10^6 at 90 mK), a 60% improvement at the coldest temperatures measured. There is a clear trend of the Q_i^{LP} decreasing with increasing temperature while Q_i^{HP} remains unchanged. A similar temperature dependence is measured for $w = 3$ μm and $w = 6$ μm resonators. Note that $Q_i = 1.7 \times 10^6$ at 90 mK is identical to the quality factor measured for the same size resonator using MBE Al-on-sapphire, while at the colder temperatures, resonators on Si outperform those on sapphire, since this anomalous temperature dependence does not appear in resonators fabricated on sapphire.

4.6.2 Temperature dependence of the quality factor at low drive power

We better characterize the Q_i^{LP} temperature dependence by continuously measuring $S_{21}(t, f \simeq f_0)$ in time and then map to f_0 and Q$_i$ using Eqs. 4.2 and 4.3. We measure S_{21} at a sampling rate of 50 Hz and at a drive power which produces an average population of $\langle n \rangle$ ~1 at the largest internal quality factors measured. As the quality factor decreases from the maximum value, the average photon population also decreases [94]. After continuously measuring for ~5 min we then update f_0 with the average value of the previous scan for the next ~5 min scan. The data for up to five resonators are measured during the same warming cycle in an interleaved fashion. We continue the repetitive interleaved measurements while the ADR slowly warms from T~40 mK to T=300 mK over ~60 hours.

In Fig. 4.5, we plot Q_i^{LP} as a function of temperature for various sizes of resonators from $w = 3$ μm, $g = 2$ μm up to $w = 50$ μm, $g = 50$ μm. We plot data of resonators fabricated on both Si(100) and Si(111) substrates for comparison. The data in Fig. 4.5 from resonators on Si(100) are the same as those used in Fig. 4.3, so these plots can be directly compared. For all of the resonators we measure short time-scale variations consistent with our previous measurements of dissipation in Al resonators on sapphire substrates arising from surface TLS [42, 72]. These short time scale variations are uncorrelated between the resonators. We also measure long time scale trends in Q_i^{LP}, which are dependent on the temperature of the resonator. These slow temperature dependent variations are well correlated between the various resonators and CPW widths. We find at the lowest temperatures, Q_i^{LP} decreases as the temperature increases. The quality factor remains roughly constant in the temperature range 55 mK $\lesssim$ T $\lesssim$ 100 mK. In

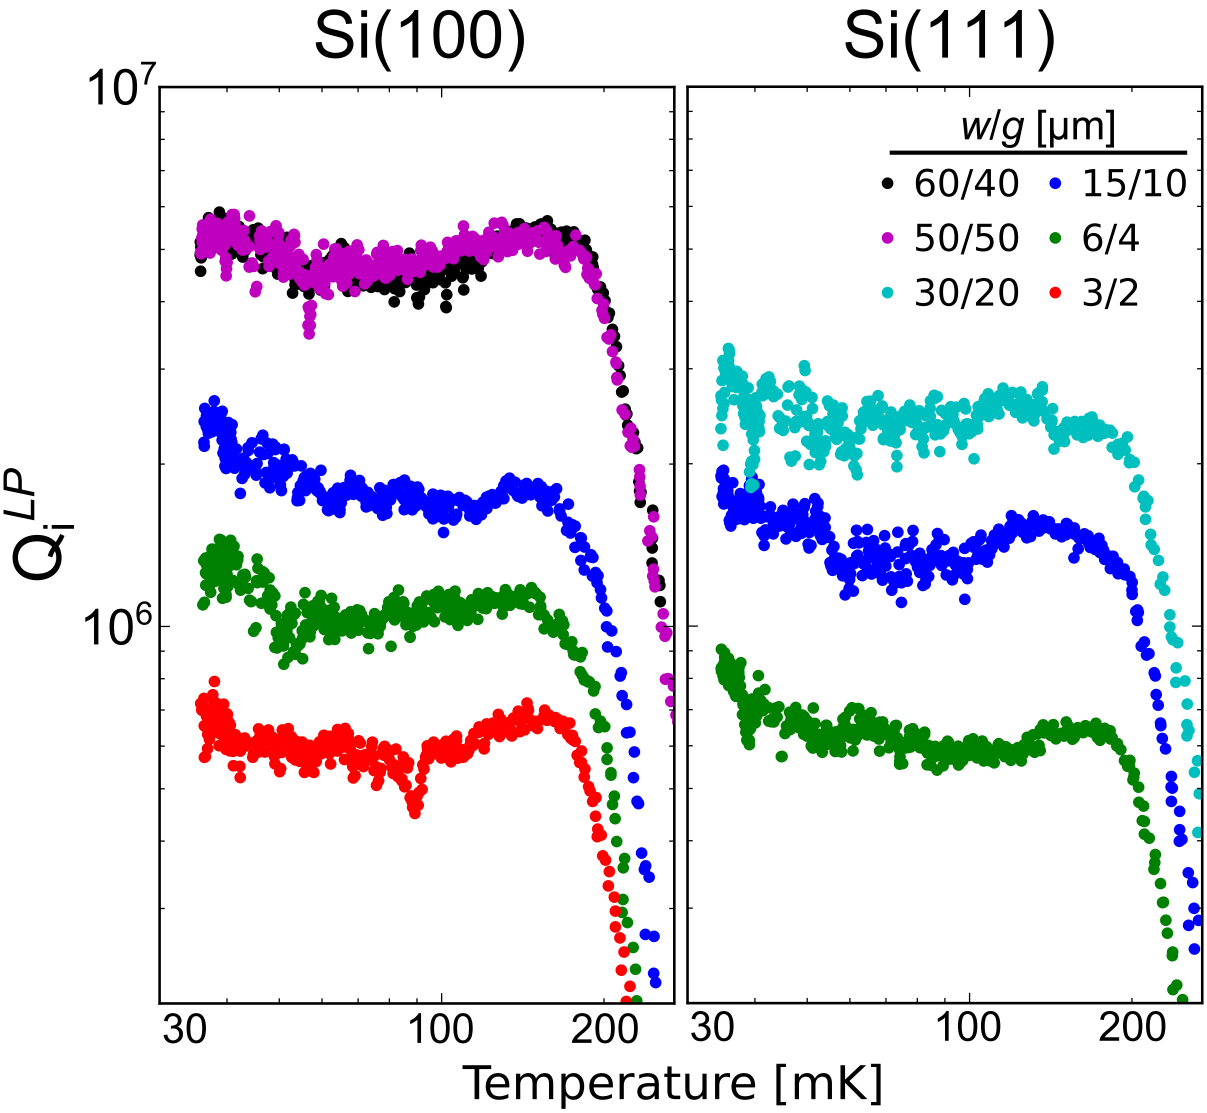

Figure 4.5: Internal quality factor Qi measured at low drive power $< n > \lesssim 1$ photons circulating in the resonators for a range of CPW widths. The left plot is for resonators on Si(100) whereas the right plot shows resonators on Si(111). The data is measured as the cryostat slowly warms from T~36 mK to 300 mK over a 60 hour period. Note how all of the resonators show an unexpected increase in the quality as the temperature decreases below T~55 mK.

the range 100 mK $\lesssim$ T $\lesssim$ 170 mK, we find Q_i^{LP} increases with increasing temperature. Above T=170 mK, Q_i^{LP} sharply decreases with increasing temperature.

The temperature dependence of the quality factor of superconducting resonators is best understood at higher temperatures, so I will begin the discussion there. At T~300 mK, Q_i^{LP}~1×10^4 (Q_i^{LP}~3.5×10^5) for the $w = 3$ μm ($w = 60$ μm) resonators and dramatically increases with decreasing temperature. This scaling is consistent with a thermally excited quasiparticle population proportional to $e^{-\Delta/k_BT}$, with Δ the superconducting energy gap, and is well described by the MattisBardeen theory of superconductors [63]. At T~170 mK the quality factor rolls over, with a local maximum of $Q_i^{LP} = 6.7 \times 10^5$ ($Q_i^{LP} = 5.3 \times 10^6$) and begins to decrease as the temperature further decreases. According to the Mattis-Bardeen theory the equilibrium quasiparticle population should continue to decrease as the temperature approaches zero, so the quality factor should continue to increase. Alternatively, if we have a non-equilibrium quasiparticle population (*e.g.* from stray IR radiation) this would result in a constant quality factor [64], which is not measured here.

The decrease in Q_i^{LP} with decreasing temperature is consistent with the standard tunneling model of TLS, where their thermal excitation leads to reduced loss in the resonator i.e., 1/Q$_{TLS}$ $\propto$ tanh($\hbar\omega/2k_BT$) [74, 95, 94]. The standard tunneling model would therefore predict an asymptotic approach to the zero temperature value, a behavior which the resonators follow down to T~55 mK, where $Q_i^{LP} = 5.7 \times 10^5$ ($Q_i^{LP} = 4.4 \times 10^6$) for the $w = 3\mu$m ($w = 60\mu$m) resonators. However below this temperature we consistently measure a significant increase in Q_i^{LP} not described by the standard tunneling model. At the coldest temperatures our cryostat can reach (T~36 mK), we measure $Q_i^{LP} = 7 \times 10^5$ ($Q_i^{LP} = 5.5 \times 10^6$), a ~25% improvement compared to Q_i^{LP} at 55 mK. This trend is the same seen in the power dependence measurements of Fig. 4.4.

We measure similar improvements in Q_i^{LP} at low T for all resonators. When averaged over the various CPW widths, this results in a ~28% increase in the quality factor below T = 55 mK, while the temperature decreased roughly the same amount, ~33%. We note that while the thermometer is not robustly calibrated in this temperature range, an error in our extrapolation would only affect the slope slightly, which appears to be near Q_i^{LP}(T $\lesssim$ 55mK) ~1/T dependence.

The data below T=55mK is not explained by the standard tunneling model. However the power dependence of the internal quality factor is consistent with TLS loss. Also the improvement of Q_i^{LP} with CPW width points to a surface loss model, also consistent with TLS loss. Finally, the temperature dependence of the quality factor between T=55 mK and T=170 mK fits very well to the thermal excitation of the TLS.

A possible explanation for this anomalous temperature dependence is an error or artifact in the measurements below T=55 mK. However this seems ruled due to several factors. First is the reproduciblity of this behavior on many resonators from different fabrications, at different frequencies, in different refrigerators, and measured with different VNAs. Additionally, we do not measure this effect with Al-on-sapphire resonators made from the same mask set, which results in resonators with only slightly different frequency and coupling quality factors due to the change in dielectric constant of the substrate. Further support of this anomalous temperature dependence is found in the literature, as I will discuss below.

Previous surface loss simulations have identified the thin metal-substrate and substrate-vacuum interfaces store 10 times as much energy as the metal-vacuum interface, due to the large relative dielectric constants of sapphire and silicon ($\epsilon \gtrsim 10$). We have shown that substantial improvements in the quality factor of Al-on-sapphire resonators are obtained by improving the metal-substrate interface, see chapter 3. Here we measure similar or better quality factors than those on sapphire, qualitatively implying that the metal-substrate interface is of similar or better quality than on sapphire. Indeed one would assume that the removal of the native oxide on the Si surface as well as the high temperature UHV anneal before the Al deposition would leave an ultra-thin or nonexistent lossy interface between the Al film and Si substrate. A very thin and sharp interface between Al and Si has been shown experimentally for Al films grown on Si substrates in UHV [96]. This assumption is further supported since small amounts of oxygen and Si impurities do not alter the superconductivity of aluminum appreciably [97, 98], likely leaving a sharp transition between superconductor and low loss single crystal substrate.

The most likely source of TLS in these resonators is the substrate-vacuum interface, since the exposed Si substrate forms an amorphous native oxide after exposure to atmosphere, and a-SiOx has been shown to be a lossy material [41]. Previous measurements of TLS in a-SiOx films have shown an anomalous behavior below T=100 mK similar to what we measure here [99]. They theorized that this anomalous behavior is due to the increased dipole-dipole interactions between TLS at low temperatures. The standard tunneling model ignores such TLS-TLS interactions. More recently, further experimental evidence in support of an interacting TLS model comes from noise measurements of microwave resonators [100], dissipation measurements of mechanical resonators [101], direct excitation of TLS using superconducting phase qubits [102], and T_1 measurements of a 3D-transmon [103].

In Ref. [103], they measure the T_1 of a fixed frequency 3-D transmon as a function of time over 10 hours. They find the energy decay rate, $\Gamma_1 = 1/T_1$, varies by as much as 25% during the meaurement. They explain this variation through an interacting TLS model. They model that there are one to a few TLS with frequencies near the transition frequency of the qubit, ω_{01}. These strongly coupled TLS interact with much lower frequency thermally activated fluctuators with energies, E> k_BT through either a dipole-dipole interaction or through phonons. When the low frequency fluctuators are excited or decay, their shift causes a broadening of the linewidth, γ_2, of the TLS coupled to the qubit as well as shifting it in and out of resonance with the qubit. They determine that the temperature dependence of T_1 of the qubit coupled to a single TLS and over an ensemble average of fluctuators should be,

$$\langle \Gamma_1 \rangle \propto \frac{2\gamma_2}{\gamma_2^2 + \delta\omega^2}, \tag{4.4}$$

with $\delta\omega$ the frequency differnce between the TLS and qubit. The authors as well as others [104] argue that $\gamma_2 \propto T^{1+\alpha}$ due to the interaction with nearby fluctuators. This temperature dependence has been measured experimentally for TLS coupled to phase qubits, where $\alpha = 0.24$ [102]. Combining this temperature dependence with Eq. 4.4 for strongly coupled TLS where $\delta\omega \lesssim \gamma_2$, they predict $\langle \Gamma_1 \rangle \propto T^{-1+\alpha}$. This is consistent with $Q_i \sim 1/T$ measured here.

Their 3D transmon device consists of two rectangular pads each 350 μm$\times 700$ μm, seperated by 50 μm and connected using a single Josephson junction. It has an energy gap, $\omega/2\pi =$

3.58 GHz and a T_1 of ~80 μs. If we convert this to a quality factor, it has Q $= 1.8 \times 10^6$. Both the quality factor and overall size of the device are comparable to the resonators measured here, implying their assumptions may also be consistent with this data. The anomalous temperature dependence of the resonators measured here could be further confirmation of interacting TLS at low temperatures.

4.7 Conclusion

We have shown that carefully controlling the interface between the substrate and the superconductor can dramatically improve the quality factors of planar resonators. These resonators still appear to be limited by TLS located in the Si native oxide. We have measured very high quality factors of the resonators up to $Q_i^{LP} = 5.5 \times 10^6$ at drive powers that populate the resonator with a single photon on average. These quality factors improve as the temperature is decreased with roughly a 1/T dependence. If this trend continues to lower temperatures, these quality factors may improve by another factor of two at common dilution refrigerator temperatures, T~20 mK, where superconducting qubit experiments are typically performed. The next step is fabricate Xmon capacitors using a similar process, to both improve their performance and hopefully better characterize this anomalous temperature dependence. I describe the preliminary results of this work in chapter 6.

Chapter 5

Mitigating energy loss from magnetic vortices

5.1 Formation of magnetic vortices

There are two types of superconductors, classified by the Ginzburg-Landau parameter $\kappa = \lambda/\xi = 1/\sqrt{2}$, where λ is the magnetic penetration depth, ξ is the superconducting coherence length and $\kappa_{typeI} < 1/\sqrt{2} < \kappa_{typeII}$. For bulk Aluminum $\xi = 1.6$ μm [60] and $\lambda = 50$ nm [105] so it is a strong type I superconductor. However for certain geometries and temperatures near Tc, bulk type I superconductors will behave like type II superconductors. For a thin film $\lambda_{Eff} = \lambda^2/d$ for d $\ll \xi$ where d is the thickness of the film [60]. The coherence length is limited by the mean free path $1/\xi \sim 1/\xi_0 + 1/l_e$, where $l_e \sim d$ for thin films since surface scattering typically dominates [60]. So as the film thickness decreases the coherence length

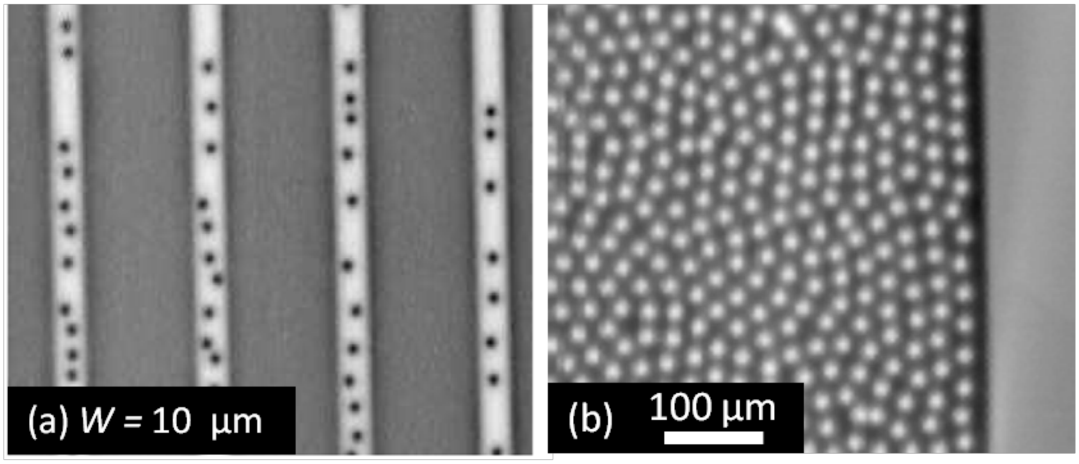

Figure 5.1: (a) This figure reproduced from Ref. [106]. Scanning Hallprobe micrograph of trapped vortices (black dots) in narrow Niobium strips (white traces) of width, $w = 10\ \mu$m . These narrow strips are comparable to the center trace of CPW resonators. (b) This figure reproduced from Ref. [62]. Scanning superconducting quantum interference device micrograph of trapped vortices (white dots) in an a-MoGe superconductor. Wide region of superconductor compares to the ground planes of CPW resonators.

decreases to $\sim d$ and the penetration depth increases, which in the very thin film limit implies that all superconductors will behave as type II.

When a strip of superconducting thin film of width w is cooled through T_c in a magnetic field, the vortices can become trapped in the superconductor. The strength of the magnetic field needed to trap vortices depends on both the width of the superconducting strip and the pinning strength of defects in the superconductor. Below this threshold magnetic field, the vortices are expelled from the superconductor [106, 107]. Above the threshold field the number of vortices grows linearly with magnetic field up until their spacing becomes small enough to fit 2 rows of vortices across the width of the strip [106].

5.2 Dissipation from vortex motion

As mentioned in chapter 2 the movement of vortices causes dissipation and is modeled as their movement through a viscous fluid [60]. However the Lorentz force needs to be greater than the pinning force to allow movement to occur. The currents associated with the modes of the SCWR overcome the pinning strength and cause vortex motion leading to dissipation. We characterize this loss using the high quality factor MBE Al-on-sapphire resonators discussed in chapter 3.

We performed field-cooled experiments on the same MBE Al on sapphire resonator devices discussed in chapter 3 using the same ADR measurement setup with the addition of a magnet located at the 4 K stage that surrounds the Al sample box, see appendix C. The protocol to perform a single field-cool is as follows. After cooling the cryostat to base temperature, we heat the sample to 1.7 K, well above aluminum's $T_c = 1.2$ K. We then apply a magnetic field perpendicular to the surface of the film. While the field is applied we cool the sample back down to ~ 50 mK, then measure S_{21} at high drive power for each resonator. We then repeat this protocol for each magnetic field. We then extract Q_i and f_0 using the fitting procedure described in chapter 3. After subtracting the loss at zero magnetic field from the other field-cool cycles, we arrive at the loss due to motion of trapped vortices $1/Q_v = 1/Q_i(B) - 1/Q_i(0)$ [Fig. 5.2(a)] and determine the fractional change in frequency $(f_0(0) - f_0(B))/f_0(0)$ [Fig. 5.2(b)].

In Fig. 5.2, notice there is no visible threshold field and near $B = 0$, the loss increases as,

$$1/Q_v = \alpha(B - B_0)^{\beta}, \tag{5.1}$$

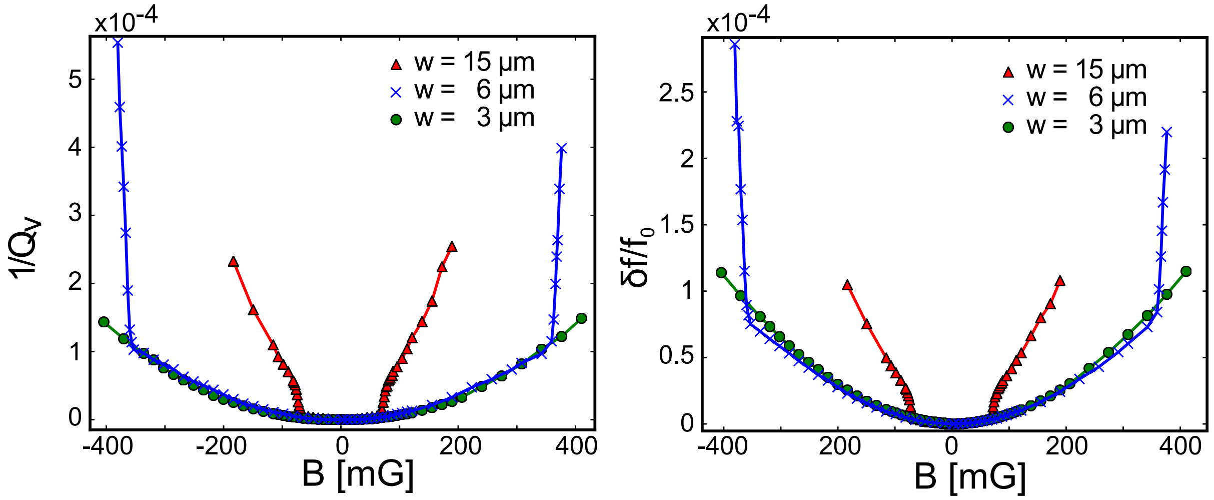

Figure 5.2: (top) Loss due to vortices $1/Q_v = 1/Q_i(B) - 1/Q_i(0)$ versus perpendicular magnetic field B which was generated using magnetic coil shown in Fig. C.1(a). (bottom) Fractional change in frequency $(f_0(0) - f_0(B))/f_0(0)$ versus perpendicular magnetic field B.

with $\beta = 2$, α encapsulates material dependent properties such as the depinning frequency and viscosity as well as the nonuniform current distribution, and $B_0 = 3$ mG is the residual field present when there is no applied field. The data in the plots has already been shifted by B_0 to be symmetric about $B = 0$.

At $B = 53$ mG a substantial increase in the loss occurs for the $w = 15$ μm resonator. This is the field at which vortices start to become trapped in the center trace of the resonator. The reason for the continuous increase in loss above 53 mG as opposed to a discontinuous jump is most likely due to a range of pinning strengths [106] associated with randomly positioned defects which trap vortices at fields below the threshold field [107, 108]. At $B = 67$ mG the loss changes back to a nearly parabolic dependence on B with $\beta \sim 1.7$. This is the threshold field associated with the center trace width. The number of vortices in the center trace now increases linearly with B up until two rows of vortices becomes energitically stable [107, 108].

At $B \sim$ 360 mG vortices begin to enter the $w = 6$ μm resonator, consistent with the scaling of the center trace width $(15$ μm/6 μm$)^2 \times 53$ mG $= 330$ mG. The largest measured loss from vortices of the $w = 6$ μm resonator has an internal quality factor, Q_i~1000, making further data acquisition very difficult.

We find that for most resonator widths, the quality factor of the resonator is substantially suppressed at magnetic fields well below the threshold associated with the center trace width. This is clearly visible for the $w = 3$ μm and $w = 6$ μm resonators. For the $w = 15$ μm resonator the internal quality factor at $B = 53$ mG is Q$_i = 2 \times 10^5$ which is a reduction by a factor of 25 from the value at $B = 0$ mG, Q$_i = 5 \times 10^6$.

5.3 Optimizing ground plane hole dimensions and location

In order to mitigate the loss from vortex motion typically holes are etched into the ground plane of superconducting circuits in order to trap the screening currents around the hole, similar to Ref. [109]. These artificial pinning centers allow the magnetic flux to penetrate the film through the holes with the screening currents surrounding the hole. Since there is no longer a normal metal core at the center of the radially decaying screening currents, there is no vortex and therefore no dissipation from there movement. However recent experiments suggest that there may be a trade-off; while the ground plane holes do limit the loss due to vortex motion [107] they may increase other loss mechanisms. One of these possible loss mechanisms may be due to the additional substrate surface exposed when creating the holes.

We test this hypothesis by measuring an array of ten CPW resonators ($w = 15$ μm,$g =$

10 μm) capacitively coupled to a common CPW transmission line, as seen in Fig. 5.3(c). Eight resonators are surrounded by an array of ground plane holes (GPH) with edge length $d = \{1\ \mu\text{m}, 2\ \mu\text{m}, 5\ \mu\text{m}, 10\ \mu\text{m}, 20\ \mu\text{m}, 40\ \mu\text{m}\}$ and separation between GPH of $s = 10\ \mu$m (see Fig. 5.3 inset). The two remaining resonators have no GPH surrounding them and are used as a control for the experiment. Based on the results of the previous experiment, the choice of $s = 10\ \mu$m should provide protection from vortex formation up to fields B~150 mG, much larger than the field where vortices begin to enter the center trace of the CPW.

First we measure the ability of the GPH to reduce the loss associated with the motion of vortices. In Fig. 5.3(a), we plot the internal loss, 1/Q$_i$, as a function of the applied magnetic field. For the resonators with no GPH we measure an immediate increase in the loss from the value at $B = 0$, consistent with the previous experiment. For the remaining resonators we find that the GPH effectively minimizes any loss associated with vortices, up until the vortices become trapped in the center trace. Above $B = 50$ mG, the resonator loss is dominated by the loss coming from vortices located in the center trace of the CPW, negating any benefit of the GPH.

Next we measure the power dependence of the resonator internal loss at $B = 0$ in an attempt to identify any additional loss associated with the GPH (Fig. 5.3(b)). Our control resonator which has no GPH surrounding it has a low power ($< n > \sim 1$) quality factor of Q$_i = 8 \times 10^5$ and high power ($< n > \sim 10^6$) quality factor of Q$_i = 5 \times 10^6$, both consistent with previous measurements of electron beam evaporated aluminum films [42]. For the remaining resonators with varying arrays of GPH surrounding them, we find nominally identical quality factors with

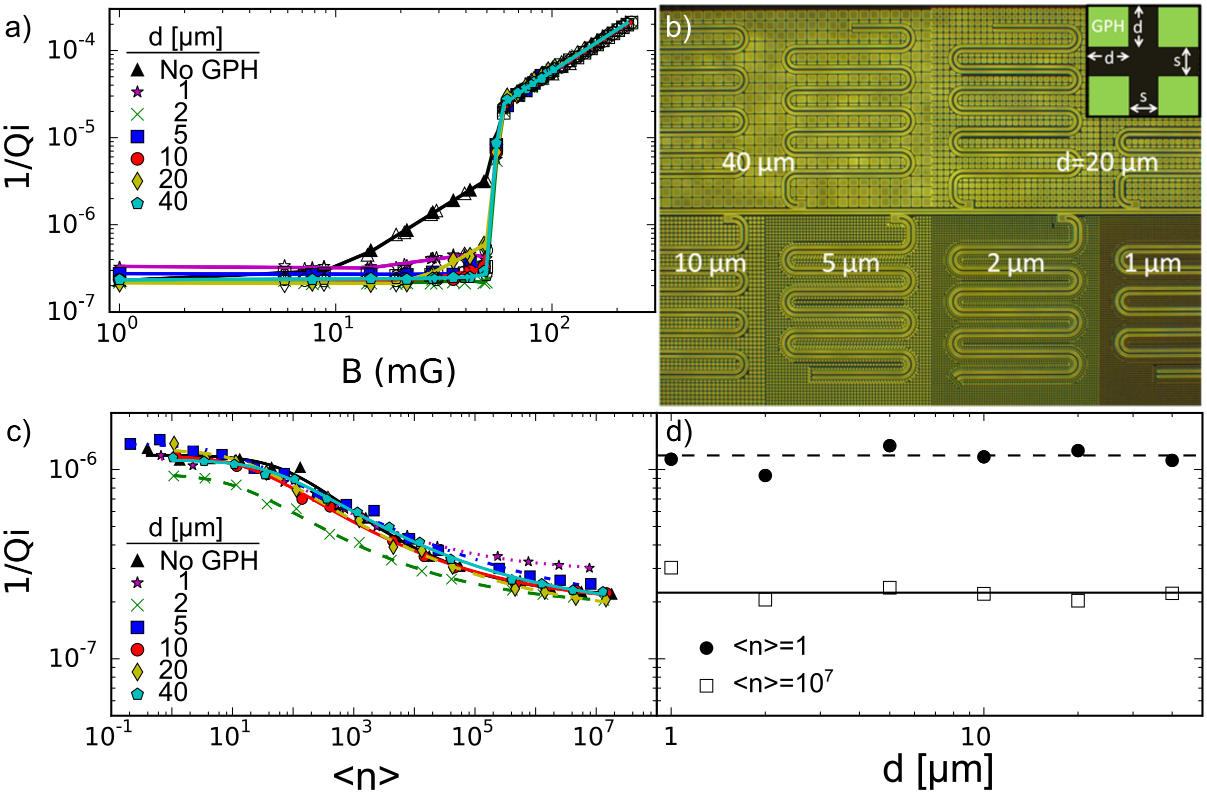

Figure 5.3: (a) Dissipation, 1/Q$_i$, as a function of magnetic field, B, for a vareity of ground plane hole (GPH) sizes. Data is measured at high drive powers corresponding to $< n >$ ~10^7 for large SNR. The data has been shifted to account for a 6.5 mG residual magnetic field. Open symbols are for negative applied magnetic fields. (b) Dark-field optical micrograph of resonators, dark regions correspond to Al while bright regions correspond to sapphire. (inset) Dimensions of GPH arrays: edge to edge separation of GPH, $s = 10\ \mu$m, and GPH edge lengths, $d = \{1\ \mu\text{m},\ 2\ \mu\text{m},\ 5\ \mu\text{m},\ 10\ \mu\text{m},\ 20\ \mu\text{m},\ 40\ \mu\text{m}\}$. (c) Power dependence of 1/Q$_i$ for same GPH. Data is taken with an applied B = 6.5 mG to zero the residual magnetic field. (d) Data from (c) replotted as a function of the GPH edge length, d, at high drive powers, $< n >= 10^7$, and low drive powers, $< n >= 1$. The solid (dashed) lines are the quality factors of the resonator with no GPH at high (low) drive powers. We find that these GPH arrays do not induce a measurable additional loss compared with no GPH.

no clear dependence on GPH size at both high and low drive powers. It is interesting to note that the $d = 40$ μm resonators have had ~$2/3$ of their ground plane removed leaving a square grid of wires, seemingly without effect on the quality factors.

These measurements show that GPH arrays are a very effective method of reducing loss from vortex motion, with no apparent additional loss at this level of quality factor and GPH geometry. A more sensitive experiment is required to test the hypothesis that the GPH can lead to additional surface loss.

In order to increase the sensitivity of the previous experiment, we place the GPH where the field strength is larger and fabricate resonators using MBE Al which have lower dielectric loss. The GPH are positioned in the middle of the center trace of the CPW as well as close to the ground plane edge. There are CPW resonators with center trace widths $w = 6$ μm and $w = 15$ μm. For both widths we have control resonators which do not have holes in the ground plane (NH). Resonators are also tested with a surrounding square grid of 2 μm $\times$ 2 μm holes in the ground plane (GPH), with a spacing of 12 μm on center same as the previously discussed experiment. The spacing between the edge of the CPW resonator and the first column of GPH is $w = 2\mu$m. There are also $w = 15$ μm CPW resonators with a single column of 2 μm $\times$ 2 μm holes in the middle of the center trace (CTH) and spaced 10 μm on center along the center trace's length. All the variations of resonators are located on the same die so that any possible inconsistencies from different runs or across the wafer are eliminated. In Fig. 5.4, we compare the power dependence of the internal quality factor of the various CPW resonators with different hole configurations.

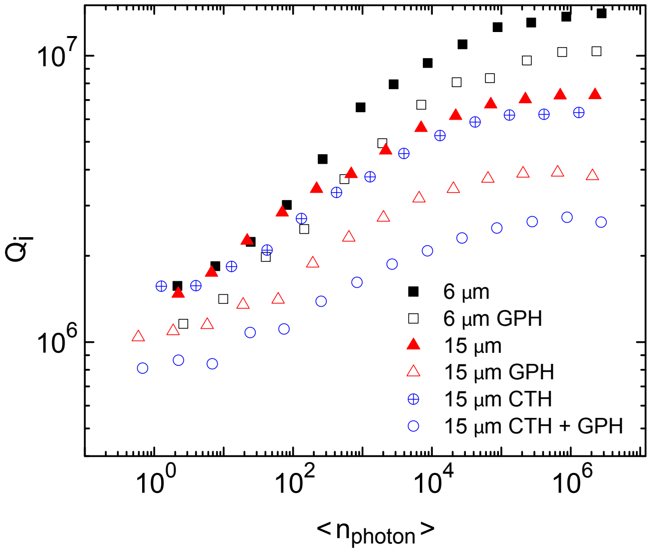

Figure 5.4: Power dependence of the internal quality factor Q_i versus average photon number in the resonator $\langle n_{photon} \rangle$, for resonators with $w = 6$ μm and 15 μm for different hole variations: (NH) no holes in the ground plane, (CTH) holes in the center trace of the resonator, and (GPH) holes in the ground plane. The holes sizes and configurations are provided in text. The typical low-power statistical error from a least-squares fit of Eq. (3.3) is ~3%, smaller than the symbol size.

The on-chip control resonators NH have quality factors which are comparable to the resonators of chapter 3. In Fig. 5.4, the GPH resonators show a substantial decrease in the quality factor consistent with the above hypothesis that the additional loss is due to the increase in the exposed substrate surface. The $w = 15\ \mu$m CTH resonator shows very similar behavior compared to NH, but it does have slightly lower Q_i especially at high power. It is interesting to note that the holes in the center trace have much less effect on Q_i than the ground plane holes. This is mostly likely due to there location in the middle of the center trace, which is where the fields are at a local minimum and smaller than near the ground plane edge of the CPW gap [110, 107]. This is in contrast to the array of holes that are present in the ground plane, where the nearest ground plane holes are close enough to the edge of the CPW gap to have significant electric fields.

5.4 Conclusion

We have measured significant additional loss in Al-on-sapphire CPW resonators due to magnetic vortex motion even at relatively small fields of order 10's of mG. Reducing this loss mechanism can be achieved through proper shielding and attention to removing any magnetic materials from inside the shield to minimize any stray fields. Through room temperature testing we measured that standard brass screws & washers, microwave connectors, and other components can have significant residual magnetic fields, so their use should be avoided in favor of magnetic free options.

Additional protection from vortex formation and subsequent loss is accomplished through an array of holes located in the ground plane of the superconducting circuit. However these holes can introduce an additional source of loss consistent with an increase in exposed surface of the substrate. This additional surface loss can be reduced if the holes located closest to the ground plane edge of the CPW resonator are kept at least $6\ \mu$ away. These results have been confirmed with more recent experiments [111].

The complete elimination of vortices may be undesirable since they are effective as quasiparticle traps [112]. Another potentially interesting experiment is to only partially etch the ground plane holes, which would not expose the surface of the substrate, but would still increase the pinning strength [62]. This would allow the vortex to still perform as a quasiparticle trap[112], while limiting the motion. Reference [113] had success with a similar experiment, where they etched a trench down the length of a CPW center stripline to pin the vortices. The etch removed 90 nm of a 150 nm thick film and was 200 nm wide on the order of the vortex core size. They found that the loss at $B = 0$ was nominally identical to a resonator without a trench, however their $Q \sim 1.5 \times 10^4$ is substantially smaller than those of Fig. 5.4 so they were probably unable to resolve the losses which dominate our resonators.

Chapter 6

Improving superconducting qubits

Complex superconducting circuits, in particular those with Xmon and gmon qubits, can require multiple layers of lithography, up to 20 various cleanroom tools, and 50 hours per wafer. Once the device is fabricated another several days are required to package and then cool the device to begin measurements. After this substantial investment of time and resources you may find that the device doesn't perform as expected. How does one determine what is the cause? Some problems and their solutions are sometimes obvious and determined using standard metrology tools, however these are usually discovered before the device is even measured. Improving qubit coherence beyond these conspicuous issues typically requires multiple device iterations requiring several months or more assuming you have some idea of what to look for. For example, after we decided to move from the phase qubit to the transmon qubit it took multiple researchers several years to obtain the high level of coherence of the Xmon and gmon qubits. This process required many improvements from our original design. Significant improvements in T_1 occurred over eight qubit design iterations that took place over a span of around two years. Resonators would directly guide the next alterations for some iterations. However resonators always served as ”witnesses” to the fabrication processes to ensure that we did not misidentify unexpected losses due to the new alterations.

6.1 Using resonators for improved qubit fabrication and circuit development

In chapters 3-5 I discussed reducing energy loss in planar resonators which should directly correlate to improvements in the Xmon capacitor, however these and other improvements still need to be realized in the qubit fabrication process. The following sections describe the further utility of resonators as dissipation testbeds.

6.1.1 Josephson Junctions

Previous work using phase qubits has shown that individual strongly coupled TLS can be located in the Josephson junction dielectric barrier [66, 102]. More recently, 3-D transmons have obtained much longer T_1 times [48]. A major difference between these devices is the physical size of the junctions, with phase qubits typically having several orders of magnitude greater junction area. The decreased dielectric volume of the transmon junctions allows individual devices to statistically avoid TLS over the measured frequency range.

The Xmon qubit junction process requires additional fabrication steps not needed for the simpler resonators of the previous chapters. At a minimum the Josesphson junctions require an argon ion mill for DC contact to previous metal layers, electron beam lithography (EBL) for small feature sizes, and lift-off metal deposition common for most EBL processing (see App. A for more detailed discussion on the fabrication). In order to better characterize these additional fabrication steps we employ a novel use of resonators recreate the Josephson junction fabrication to better characterize the quality factor of the resulting structures. A series of resonator experiments identified the argon ion milling of the substrate to reduce the quality factor of Al resonators on a sapphire substrate by a factor of two [114].

The factor of two reduction in quality factor from ion milling is in excellent agreement with the sputter deposited resonators discussed in Ch. 3 that used a similar ion mill clean before aluminum deposition. The use of resonators allowed for additional experiments not possible with the Xmons. For example we tested the quality factors of $\lambda/2$ resonators deposited by EBL lift-off deposition without the use of an ion mill, since no galvanic connection of the center trace is needed. We found these resonators also suffered a 50% reduction in quality factor due to residue left behind from the electron beam imaging resist. This residue can be effectively removed using a low energy oxygen clean such as room temperature ozone or downstream ashing at 150° C. Given the significant participation of the qubit's narrow junction electrodes, these reduced resonator quality factors indicate the T_1's of larger Xmon qubits (*e.g.* $w = 24\ \mu$m, $g = 24\ \mu$m) may be suppressed due to these fabrication steps. Further investigation into mitigating this additional loss channel is currently ongoing.

6.1.2 Airbridges

A qubit circuit requires more wiring complexity than the simpler resonator circuits discussed in previous chapters. The additional drive lines required for qubit manipulation can lead to dissipation through parasitic coupling to undesired slotline modes generated by improperly connected CPW grounds (discussed in chapter 2). Coupling to slotline modes can be heavily attenuated through the use of crossovers [59]. Typically, lossy dielectrics physically support the crossovers. However use of lossy dielectrics should be avoided near the qubit and certainly should not be part of the qubit structures, limiting their utility.

Wirebond crossovers which do not have a lossy dieletric for support are an alternative, however the large inductance due to their length leads to a significant impedance at microwave frequencies. Instead we have developed an airbridge crossover which has no lossy diectric support and whose short length has significantly less inductance than a wirebond, thus providing a good electrical short at microwave frequencies [59]. The airbridges are temporarily supported with photoresist that was reflowed to create a structurally strong arch shape.

We again measured resonators to ensure the additional fabrication steps required to produce the airbridges did not lead to additional losses. We also characterized any additional loss of the airbridge structures by including them over CPW resonators to connect the grounds. A third experiment used up to 10 airbridges connecting segmented pieces of the center trace of a CPW resonator to measure the added loss in more complex qubit structures such as the gmon discussed later in the chapter. The airbridge fabrication did not significantly reduce the resonator's quality factor, however the airbridge structures did introduce roughly equal loss whether used to connect the grounds or as the center conductor. For example 110 airbridges covering the resonator reduced the quality factor by a factor of two. The fabrication process of our airbridges is discussed in more detail in App. A and further details of the experiment in Ref. [59].

6.1.3 Witness and control resonators

In the proceeding sections and chapters I described how superconducting resonators can be used as dissipation testbeds, on-chip secondary thermometers, and even on-chip magnetic field detectors all without requiring any additional measurement wiring. They do require some real estate on the device, which may prohibit their use. Using these properties of resonators we identified the need for additional cold magnetic shielding in the dilution refrigerator to shield from stray fields associated with microwave circulators, SMA connectors, screws and washers. Additionally we used Al-on-silicon resonators to help identify potential thermalization issues associated with the new magnetic shielding in the DR.

Dedicated “control” resonator chips are included on the same wafer as the qubit dies and are etched during the same lithography step as the Xmon capacitor. They are nominally identical to the chip layouts used in the previous chapters. They contain one transmission readout line and 10-16 resonators typically having several CPW widths. However, since they are located on a physically separate chip they would use up additional microwave lines and space on the cryostat mix plate. For this reason these devices were rarely cooled down in the dilution refrigerator and mainly measured in an ADR.

“Witness” resonators are placed on the same devices as the qubits, when space allows. These are used in a space saving capacity and are typically capacitively coupled to the same readout line as the readout resonators. While these witness resonators can be extremely useful to help diagnose magnetic field problems, thermalization issues, fabrication failures etc., all but the narrowest resonators tend to suffer from dissipation from parasitic coupling to the qubit control lines, discussed in more detail below. The repeated use of witness and control resonators on nearly every wafer involving a new fabrication or circuit design led to the consistent and dramatic improvements in the qubit’s decay time, T_1, over a period of two years.

In the remainder of this chapter I will focus on four experiments which highlight the current state of the art in Xmon and gmon qubit technologies. The base metal for the first Xmon device is electron beam evaporated aluminum on a sapphire substrate, and has three uncoupled qubits per die. Next I discuss current progress toward transitioning qubit circuits from sapphire to silicon substrates. The final two devices are designed for specific algorithms. First of these is an MBE Al-on-sapphire Xmon device arranged in a nine qubit linear chain designed specifically for operating the repetition code (a primitive of the 2-D surface code discussed in chapter 1). Finally I will discuss a three qubit device based on the gmon adjustable coupling architecture. This device is used to simulate the inelastic collision of Na and He, which I will discuss in chapter 7.

6.2 Uncoupled Xmons: developmental devices

6.2.1 Al-on-Sapphire

Device

Several decoherence mechanisms (e.g. surface loss, radiative loss, magnetic vortices) depend on the width of CPW center trace and gap. With this in mind, the following circuit layout provides a comparison standard for troubleshooting new fabrication processes, such as moving to silicon substrates described later. As seen in Fig. 6.1, the device has 3 Xmon qubits as well as 5 CPW witness resonators both of varying widths. The center trace and gap of the Xmon capacitor are equal $w = g$, with $w = \{8\ \mu\text{m}, 16\ \mu\text{m}, 24\ \mu\text{m}\}$. For each qubit size there are two witness resonators, except there is only one $w = 24\ \mu$m witness resonator due to space limitations on the chip. The device design minimizes stray coupling of the qubits to other circuit elements, including the other two qubits. Witness resonators are located on the opposite side of the readout CPW transmission line as the qubits.

The device is fabricated using the process outlined in appendix A. The base wire metal is deposited using electron beam evaporation of Al in a Plassys double angle evaporator. A single sapphire wafer contains the two variations of the devices; half of the devices include airbridge crossovers and the other half of the devices do not have crossovers. In this way both of the variants are subject to the same fabrication processes. The airbridge crossovers are used to connect the ground plane on either side of the CPW control and readout lines. The purpose of the crossovers is to mitigate parasitic coupling to spurious slotline modes. One chip of each

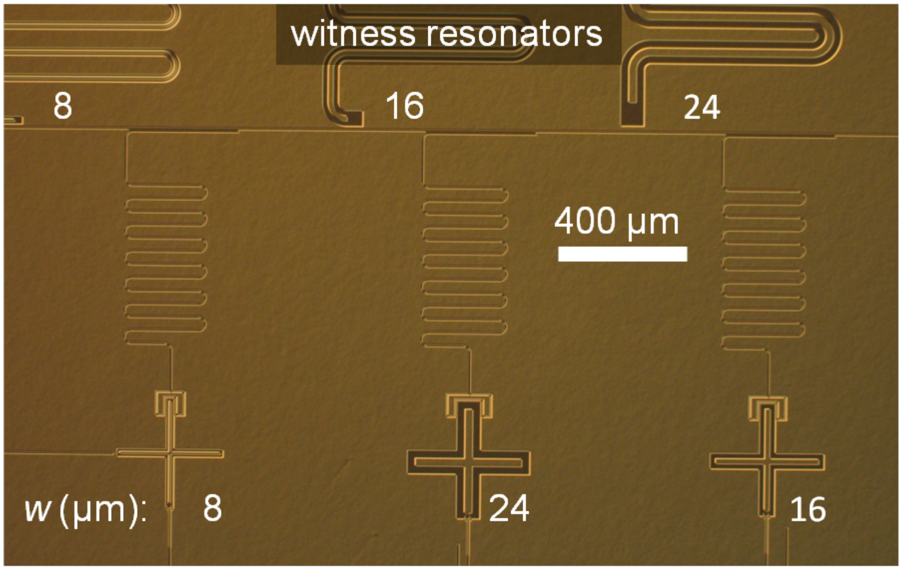

Figure 6.1: Differential interference contrast micrograph of the electron beam deposited Al-on-sapphire device. The device has three Xmon qubits with each capacitively coupled to an individual readout resonator. In addition there are five CPW witness resonators. Both the witness and readout resonators are measured using a common CPW feedline. The Xmon qubits and witness resonators have serval CPW widths to measure the scaling of T_1. The center trace and gap of the Xmon capacitors and witness resonators are equal $w = g$, with values $w = \{8\ \mu\text{m}, 16\ \mu\text{m}, 24\ \mu\text{m}\}$. This same mask set is used to fabricate devices on a Si substrate in the next section.

variant is then cooled in a dilution refrigerator and measured.

Xmon Results

Interestingly we find that the two chip variations, with and without airbridges, has no statically significant difference in T_1 spectra for any of the Xmon widths, as seen in Fig. 6.2. All of the spectra have large variability when comparing neighboring frequencies. However there is a general trend of higher T_1 at lower frequencies. The average T_1 for the $w = 8$ μm Xmons is 15 μs over the measured frequency range. We also plot the T_1 spectra as quality factors by multiplying the measured T_1 by the angular frequency, ω_0, at which the data was taken. When plotted in this manner we find no meaningful dependence on the frequency, with an average quality factor of 4×10^5 over the frequency range. For the Xmons with wider CPW capacitors we also measure a constant mean quality factor over the frequency range, with an average quality factor of 6×10^5 for both the $w = 16$ μm and $w = 24$ μm Xmons. These results are comparable to those previously measured using MBE Al-on-sapphire Xmons [53, 27, 28].

For the larger $w = 16$ μm and $w = 24$ μm Xmons we find an improved T_1 spectrum over the measured frequency range compared with the $w = 8$ μm qubit, qualitatively consistent with decreasing the surface loss associated with the CPW capacitor. However the $w = 16$ μm and $w = 24$ μm Xmons performed identically on both device variants. Furthermore their average T_1 improved by only 50% compared with the $w = 8$ μm qubit, much less than simulations and resonator measurements predict. This indicates that surface loss from the capacitor is not the only dominate loss mechanism in the these devices.

To gain further insight into what may be causing the increased dissipation in the qubits, we

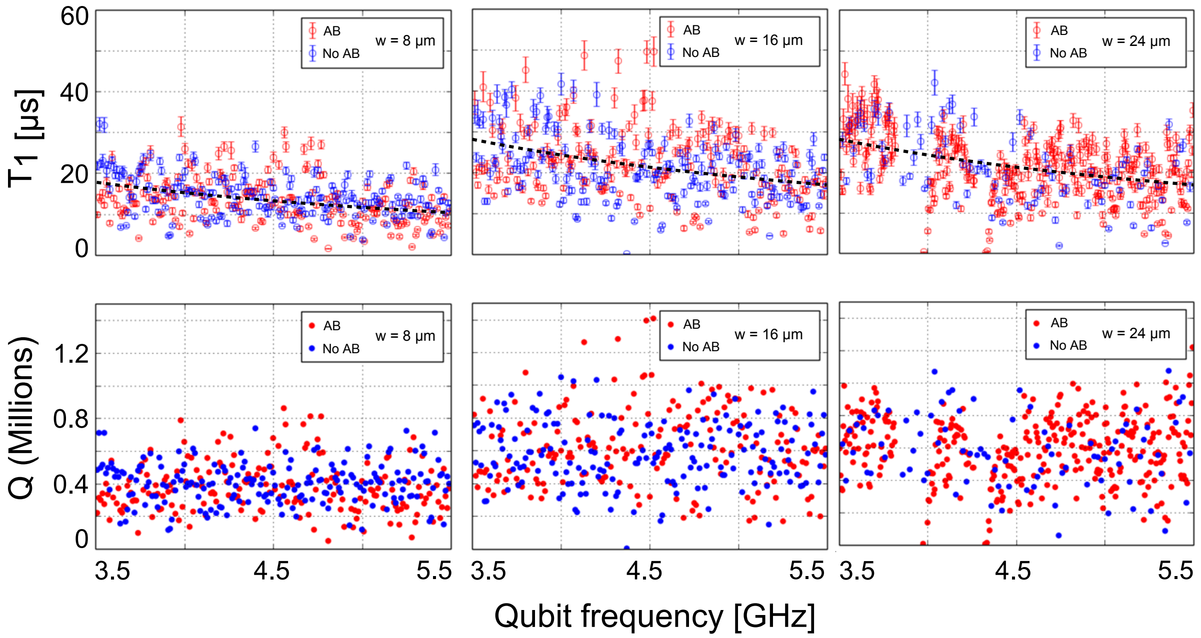

Figure 6.2: T_1 as a function of qubit frequency. Red symbols are for device which had airbridge crossovers on all readout, drive, and bias lines, blue symbols represent the device which has no crossovers present. Three Xmon CPW capacitor widths are plotted, with left, middle, and right plots having $w = g = \{8\ \mu\text{m}, 16\ \mu\text{m}, 24\ \mu\text{m}\}$, respectively. The lower row of plots is the same data converted to a quality factor, $Q = 2\pi f T_1$, where f is the qubit frequency. The black dashed line represents a quality factor of $Q_{Ave} = 4.0 \times 10^5$ for the $w = 8\ \mu$m and $Q_{Ave} = 6.0 \times 10^5$ for both $w = \{16\ \mu\text{m}, 24\ \mu\text{m}\}$ Xmons.

begin by assuming the control resonators fabricated with the qubit devices accurately predict the quality factor of the Xmon capacitor. The Xmon capacitors should then produce qubits which have quality factors of $Q_{8\mu m} = 5.5 \times 10^5$, $Q_{16\mu m} = 9.5 \times 10^5$, and $Q_{24\mu m} = 1.2 \times 10^6$ if they are only limited by the capacitor. The resonator data predicts the Xmons should have up to a factor of two larger quality factor than the average we measured.

We then subtract the expected Xmon capacitor loss from the average measured loss. Interestingly, this results in an additional parallel loss channel for each qubit with very similar quality factors of 1.5×10^6, 1.6×10^6, and 1.2×10^6 for $w = 8$ μm, 16 μm, 24 μm respectively, showing no clear dependence on CPW width. This may point to a common loss channel shared among all the qubits with a quality factor of ~1.4×10^6 over the measured frequency range and independent of these capacitor sizes. Of the dissipation sources listed in Chap. 2, only TLS located in the Josephson junctions themselves, the stray junction, or the junction electrodes fit the above criteria.

TLS located in the Josephson junction are unlikely candidates due to the very small volume of dielectric in the 200 nm × 300 nm × 1.5 nm typical of a Josephson junction used in transmon qubits. In this small dielectric volume regime it has been found that strong coupling to individual TLS is more common than weaker coupling to a bath more typical of resonators. This reasoning is also consistent with the many high coherence experimental results when using 3-D transmon qubits [48]. Due to the large electric fields, the TLS located in the junction are more likely to have strong coupling to the qubit, producing avoided level crossings in the qubit spectra and coherent swapping when on resonance. Neither of these appear in the current devices.

The current Xmon fabrication procedure produces a stray Josephson junction during the oxidation step in between the top junction electrode and the base wiring ground plane. In order to limit its effect we make the area of this junction orders of magnitude larger than the intended junction ($8\ \mu\mathrm{m}^2$ compared to $0.06\ \mu\mathrm{m}^2$). Therefore if the stray junction is indeed the source of dissipation we would expect to see a strong frequency dependence when biasing the SQUID of the Xmon, as this would decrease the electric fields inside the stray junction [115].

The junction electrodes contribute a small fraction of the total capacitance of typical qubits, for the Xmon architecture about 1% [114]. Additionally we measured the losses associated with the junction fabrication using resonators, described in section 6.1.1. The ion mill was found to cause a 2.5 times lower quality factor than resonators defined using a dry etch process. If 1% of the capacitance has a factor of 2.5 larger loss tangent than the remainder of the capacitance, this should not measurably alter the total quality factor. However the junction electrodes are in very close proximity to one another, about $1\ \mu$m separation, thus increasing their participation ratio by a factor of $(\frac{72\ \mu\mathrm{m}}{1\ \mu\mathrm{m}})^{0.86} \approx 40$ compared to the more spacious cross section of the $w = 24\ \mu$m CPW cross, assuming the participation of the more complicated junction geometry scales with effective width as is the case with the CPW cross-section. Accounting for the increased participation, we can quickly estimate the junction electrodes provide an additional parallel loss channel with quality factor of $\mathrm{Q}_{JE} \approx \frac{1}{10^{-2}}\frac{1}{40}\frac{1.2\times10^6}{2.5} = 1.2 \times 10^6$. While this is a simple estimate, it is interesting to note that this predicts the loss of the junction electrodes to be of the same order as the loss from the rest of the CPW cross. In addition, this estimate is in

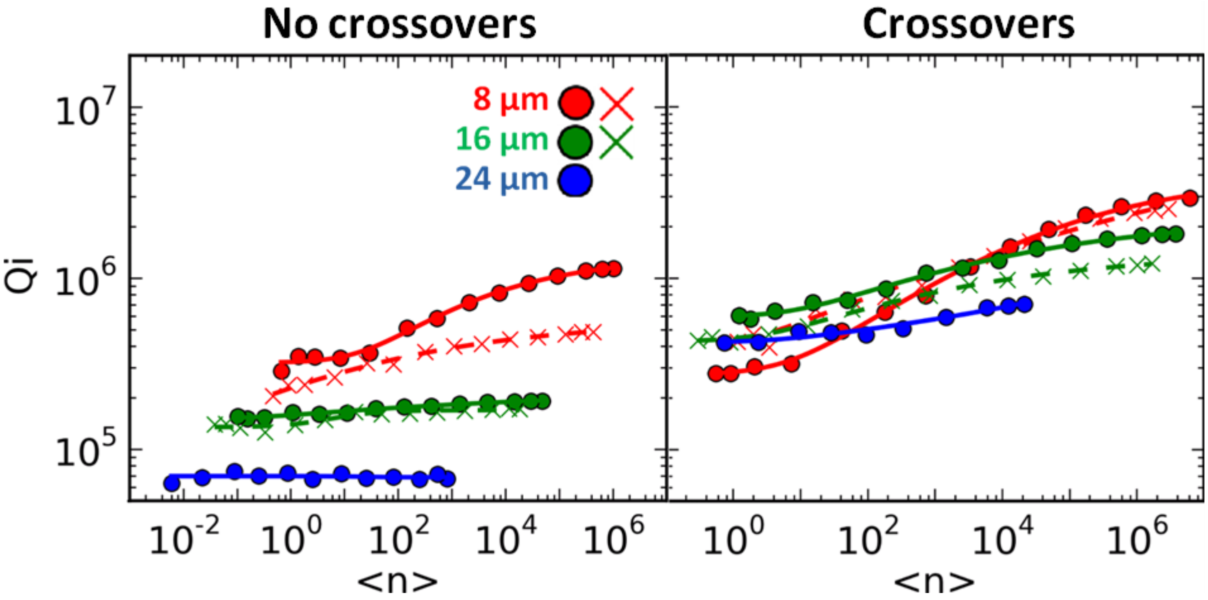

Figure 6.3: Power dependence of witness resonator quality factors on two nominally identical qubit devices containing either crossovers or no crossovers. The **Crossovers** device has airbridge crossovers connecting the CPW grounds for the readout and qubit control lines. The **No crossovers** device is identical to the **Crossovers** device except it does not include airbridge crossovers. Both devices were fabricated on the same wafer simultaneously. Note the **Crossovers** device has much larger quality factors than the **No crossovers** device, particularly for the larger resonators. This is consistent with the airbridge crossovers reducing the parasitic coupling between the resonators and the Xmon drive lines.

good qualitative agreement with the proposed parallel loss channel of 1.4×10^6 extracted from the various Xmon spectra. This leaves TLS located on the surfaces of the junction electrodes as a likely source of the suppressed T_1, and is the current focus of ongoing research [115, 116].

Witness resonator results

While the qubits have no measurable evidence that the presence of airbridge crossovers helped with dissipation, the on-chip witness resonators did show a very strong dissipation mechanism, coupling to the qubit drive lines. We confirmed the coupling to the drive lines as the cause of added loss by measuring, with a VNA, the transmission spectrum from one drive line to another, and indeed we were able to see the resonance peaks of both the large witness resonators and also the smaller qubit readout resonators. Due to this, we only use control resonator devices, which have no additional microwave lines, to measure the quality of a fabrication run. However, the witness resonators can provide insight into radiative losses, probably due to their much larger size compared with the qubit. Understanding and reducing this loss mechanism may become crucial as the number of control lines increases as we scale up to many individually controlled qubits per device.

In Fig. 6.3 we plot the power dependence of the quality factor for both the sample containing airbridge crossovers , labeled Crossovers, and the one without, No crossovers. It is immediately obvious that the sample with no airbridges shows substantially suppressed quality factors, especially for the larger CPW widths. By including airbridge crossovers in the circuit, high power quality factors are improved by an order of magnitude for the $w = 16\ \mu$m and $w = 24\ \mu$m witness resonators and by a factor of three for the $w = 8\ \mu$m witness resonator. However even with the major improvements due to the airbridges, the high power quality factors of these resonators are suppressed compared to the control resonator devices.

6.2.2 Al-on-Silicon

With the very promising and reproducible results of Al-on-Si CPW resonators discussed in chapter 4, we began this experiment eager to measure record T_1 values. Unfortunately much more development work was required than originally expected. Simply changing the substrate from sapphire to silicon led to many confusing and contradictory results as well as numerous changes in the fabrication required to minimize the interaction between Al and silicon. We are continuing to improve the process, which has been over a year in the making.

While troubleshooting this process, several different circuit designs have been used in these experiments. Each experiment is designed to enlighten us on general loss mechanisms or test a specific hypothesis. All of the designs have included three to five Xmon qubits and zero to five witness resonators, and typically included variations of the CPW capacitor width. In this section I will cover two of these devices, including an experiment using the identical mask set used for the previous experiment on a sapphire substrate. First I will discuss our initial design which has 5 Xmon qubits each with a $w = g = 24$ μm and no witness resonators.

Initial devices

The initial devices, fabricated using an electron beam deposited aluminum film on Si substrate, were initially only intended as practice for a subsequent MBE Al-on-Silcon wafer. This fabrication had numerous fabrication issues, described later in this section, but yielded functioning devices. Due to some critical cleanroom equipment being unavailable for an extended time, we decided to measure these devices while waiting for replacements. In Fig. 6.4, we see a 3-D micrograph of the device, consisting of five uncoupled Xmon qubits, all with $w = 24$ μm, $g = 24$ μm CPW capacitors. This qubit fabrication is the first to include airbridges. Many of the airbriges had collapsed during the junction fabrication steps and could be completely removed with very strong sonication in IPA. After redesigning the airbridge process, subsequent devices no longer had these bridge failures. At the time of this writing the best Xmons on

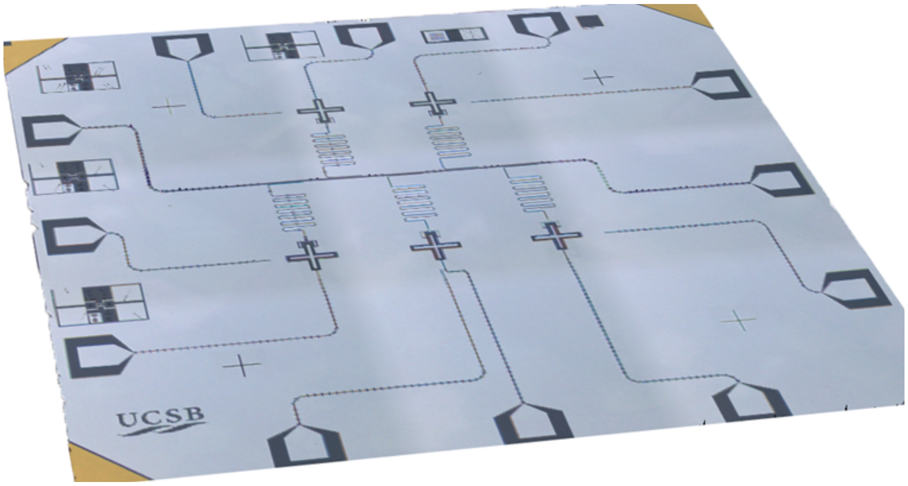

Figure 6.4: 3-D micrograph of initial Al-on-silicon Xmon device. It contains five Xmons ($w = 24\ \mu$m, $g = 24\ \mu$m) each with their own microwave drive and bias lines. The readout resonators are inductively coupled to a common CPW transmission line for readout.

silicon substrates came from this initial device, in spite of the numerous issues.

The first Al on Si Xmons show a lot of variability between qubits on the same device. In Fig. 6.5 we see the T_1 spectrum of a high coherence Xmon on the left, while the right plot shows a nominally identical qubit from the same device having lower T_1s over the full frequency range. Roughly half of the qubits on these devices show suppressed T_1s (below $10\ \mu$s) across most or all of the measured frequency spectrum, similar to Fig. 6.5(b).

The Xmons which have longer decay times, similar to Fig. 6.5(a), have a similar average T_1 to the Xmons on sapphire discussed in the previous section. For comparison, the black dashed line represents a quality factor of Q$= 6 \times 10^5$, the average quality factor of the sapphire devices in the previous section. It is interesting to note that while the average T_1 is similar, the Xmons on silicon show less variability in the spectrum. The majority of the spectrum on the left had $T_1 > 10\mu$s and only a few narrow frequency regions had $T_1 < 15\ \mu$s. This is typically not

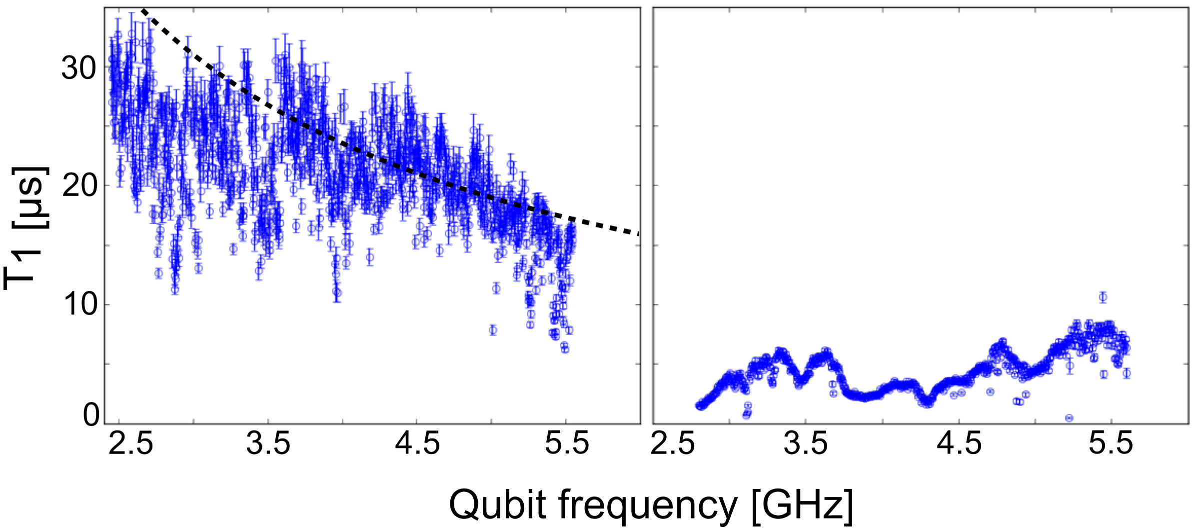

Figure 6.5: T_1 spectra for two Xmon qubits from the initial Al on Si device. The spectrum on the left shows good T_1s over the frequency range while in the right plot a nominally identical qubit on the same device has much lower T_1 over the same frequency range. The black dashed line represents a quality factor of Q= 6×10^5. The left qubit spectrum displays less variability compared to Xmons on sapphire. Error bars represent uncertainty in the fit to an exponential decay.

the case with Xmons on sapphire. Qubits having a cleaner spectrum are highly desirable when running a multi-qubit algorithm which requires strategically avoiding other qubits and defects when changing the qubit frequency.

Comparison with Xmons on sapphire

Subsequent experiments typically show only suppressed T_1's, while only a few qubits have T_1's comparable to the best of the initial device. Iteration time between measured devices was slow due to many fabrication complications that needed to be resolved. This included integrating a more robust airbridge fabrication process that could survive the Josephson junction fabrication as well as other fabrication issues discussed later in the section.

The most informative cooldown to date was with devices fabricated using the identical mask set as the Xmons on sapphire experiment described in the previous section. In Fig. 6.6a, the qubits show heavily suppressed $\mathrm{T}_1 < 15\ \mu$s over the measured frequency range with no clear CPW width dependence, indicating the Xmon capacitor is not the dominate source of loss. In Fig. 6.6(b) the Xmons on sapphire results from Fig. 6.2 are replotted on the same T_1 scale. The use of the same mask set eliminates design errors as possible explanations for the suppressed spectrum of Xmons on Si.

The control resonators provide further insight. In Fig. 6.6(c) and (d) we see similar quality factors of resonators on Si and sapphire substrates, with the resonators on Si performing slightly better. This provides further evidence that the Xmon capacitor is not the limiting factor for Si devices.

Research is still ongoing into the cause of the suppressed T_1 spectra on Si. We have already

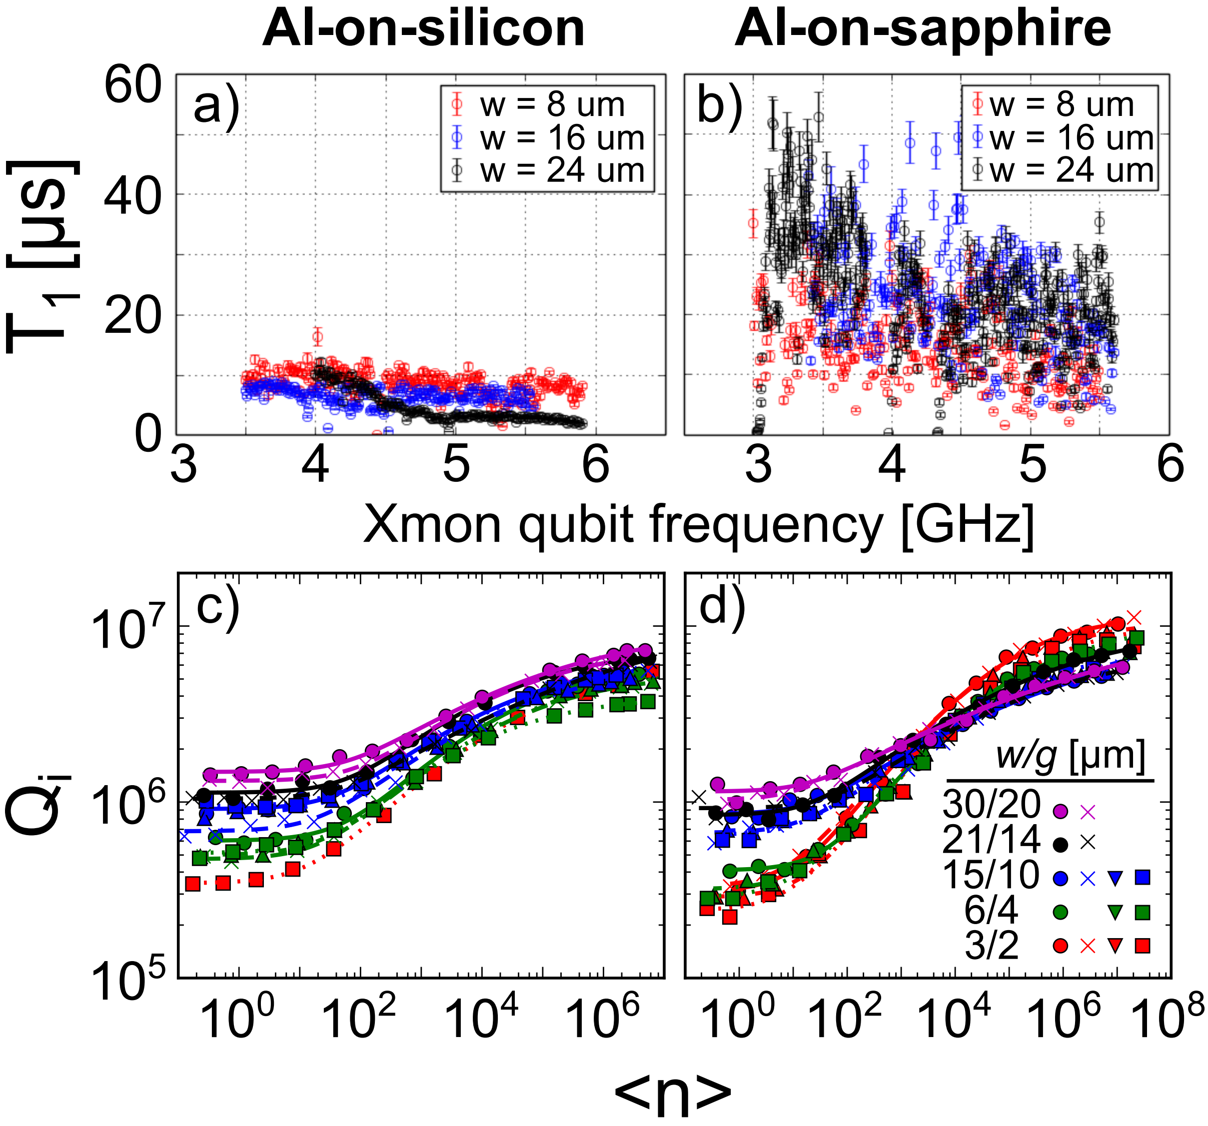

Figure 6.6: (a) and (b) T_1 spectra of Xmon qubits fabricated from Al on Silicon and Al on sapphire, respectively. Both devices are fabricated using the same mask set for direct comparison. The Al on Si Xmon have drastically reduced $T_1 < 12\ \mu s$ compared with the Al on sapphire $T_1 \sim 20\ \mu s$. (b) Same data shown in Fig. 6.2. Note the Al on Si Xmon T_1 spectra do not have a discernible width dependence. (c) and (d) Power dependence of internal quality factor, Q_i, for the control resonators fabricated during the same lithography step as the Xmon capacitors. The Al on Si resonators slightly outperform those on sapphire. This indicates that the Xmon capacitor is not the dominant source of loss in the devices on a Si substrate.

ruled out numerous possibilities through further check experiments, including poor thermalization, magnetic vortices, and the cryostat. It is interesting to speculate that if the Xmons on sapphire appear to be limited by some aspect of the Josephson junction fabrication, then this may be the limiting factor for the Xmon on Si devices as well. Currently we are in the process of testing the Josephson junction fabrication using resonators to determine if this is the cause. Preliminary results look very promising.

Fabrication complications

The simple fabrication of resonator devices, consisting of a single layer of photolithography, did not expose any unexpected processing complications for Al on silicon substrates. This led to the fantastic results obtained with resonators discussed in chapter 4. However more complex devices requiring multiple layers of lithography such as the Xmon qubit, exposed a large number of unfavorable interactions which are either smaller or not present when a sapphire substrate is used.

Al attack photoresist developer

It is well known that photoresist developer can etch aluminum. For example, developing for 60 s with Al on sapphire substrates and using AZ 300MIF, a TMAH based developer, etches 20 nm into the Al. However this etch rate can be greatly enhanced when the Al is deposited on a Si substrate. During the same develop step, a 100 nm film of Al can be completely etched away.

This uncontrollable etching of aluminum led us to explore another developer, AZ Dev 1:1,

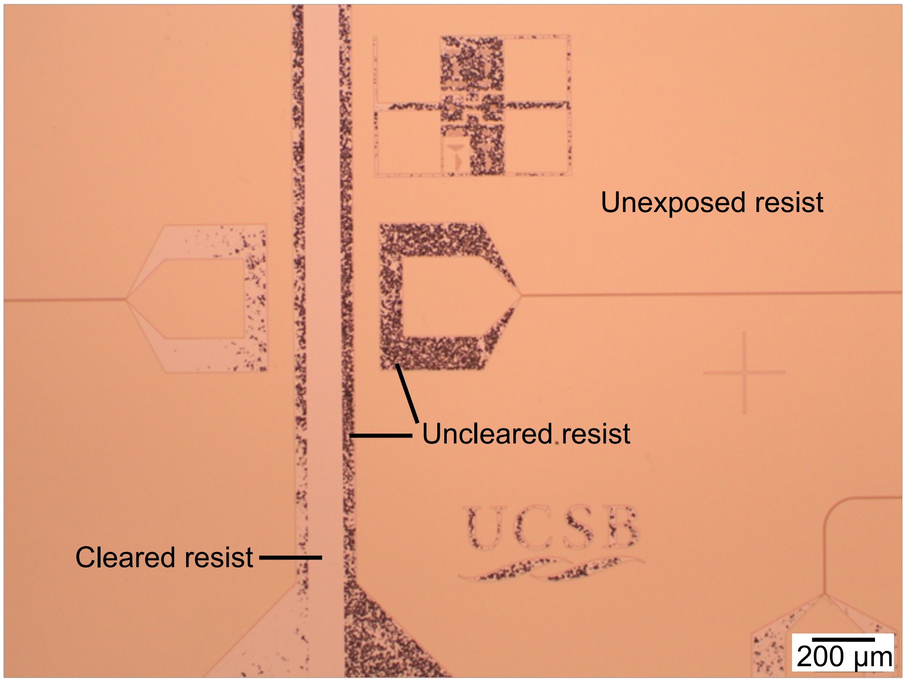

Figure 6.7: Micrograph of undeveloped resist immediately after developing for three minutes in AZ Dev 1:1, a Na based developer which has minimal attack of Al. The light regions are where the photo resist developed, the spotted black areas are where the photoresist did not develop completely, and the remaining region still has the unexposed photoresist. This issue does not arise when using a more common TMAH based developer. However the TMAH based developer can lead to a significant etch rate of Al, even etching completely through the 100 nm film in the 60 s development time.

a Na ion based developer, which has the least amount of aluminum attack according to the manufacturer. The Al attack decreased substantially when using this developer even though the development time increased to 3 min. It was still imperative to not expose bare Si during the development. The exposed Si would again cause the etch rate to dramatically increase. This may be due to an electro-chemical interaction between the Si wafer, Al film and the ionic developer.

Use of AZ Dev 1:1 developer would sporadically lead to incomplete development of the photoresist, as seen in Fig. 6.7. Since the Al is partially etched away where the resist has been removed, these wafers would be discarded and we would start over. We conducted many controlled tests to determine the parameters that led to incomplete development, but were not able to reproduce these results consistently. We varied bake temperature, bake time, develop time, and exposure time and found no consistent cause. We did notice that the incomplete development seemed to occur only when developing large areas across the wafer, such as the initial base wire etch, which develops the dicing marks, bond pads, and control CPW wiring.

Our current strategy is to develop the initial base wiring layer using AZ 300MIF to ensure the resist properly clears. Additionally, we use this effect to intentionally wet etch the Al in the area local to where the junctions will be fabricated. Since the developer does not etch Si, it is used as an etch stop providing a flat and smooth surface with no possibility of trenching or undercutting into the Si. This ensures the Josephson junction process makes good galvanic contact with both the cross of the Xmon and the ground plane. Additionally, it eliminates the possibility of polymer formation which can occur during the ICP dry etch, discussed next.

For the subsequent layers we use AZ Dev 1:1 and are much more careful about mask layout, ensuring that no Si is exposed (where Al has been etched away in a previous step) during the development process.

Polymer formation during fluorine quench

The dry etch of an Al film typically uses a BCl_3/Cl_2 chemistry in a reactive ion etcher (RIE). However, after the etch Cl byproducts remain on the exposed surfaces as well as the resist. After venting the RIE chamber the wafer is exposed to the atmospheric water which reacts with the byproducts forming HCl. The HCl will locally etch the Al uncontrollably.

To combat this attack, two approaches are used, a fluorine quench and a DI water soak. Before removing the wafer from the etcher, a fluorine based quench is performed to replace Cl byproducts with F. This works well because fluorine based chemistry do not etch Al [83]. Additionally any remaining concentration of HCl can be reduced by saturating the wafer in deionized (DI) water immediately after venting the chamber.

The standard protocol used at UCSB for sapphire substrates uses both of these strategies with success for years, the full process is described in appendix A. However when using the same protocol with Si substrates both strategies led to irreprodicable problems. The issues with submerging the wafer in DI water will be discussed in the next section.

We perform the *invacuo* quench using CF_4 for 10 s. As seen in Fig 6.8, this quench will sometimes cause a polymer to form at the sidewall edges. However this is not always the case, and may be dependent on other cleanroom user processes. One can deduce that the polymer must form after the BCl_3/Cl_2 etch is complete since the polymer does not appear to

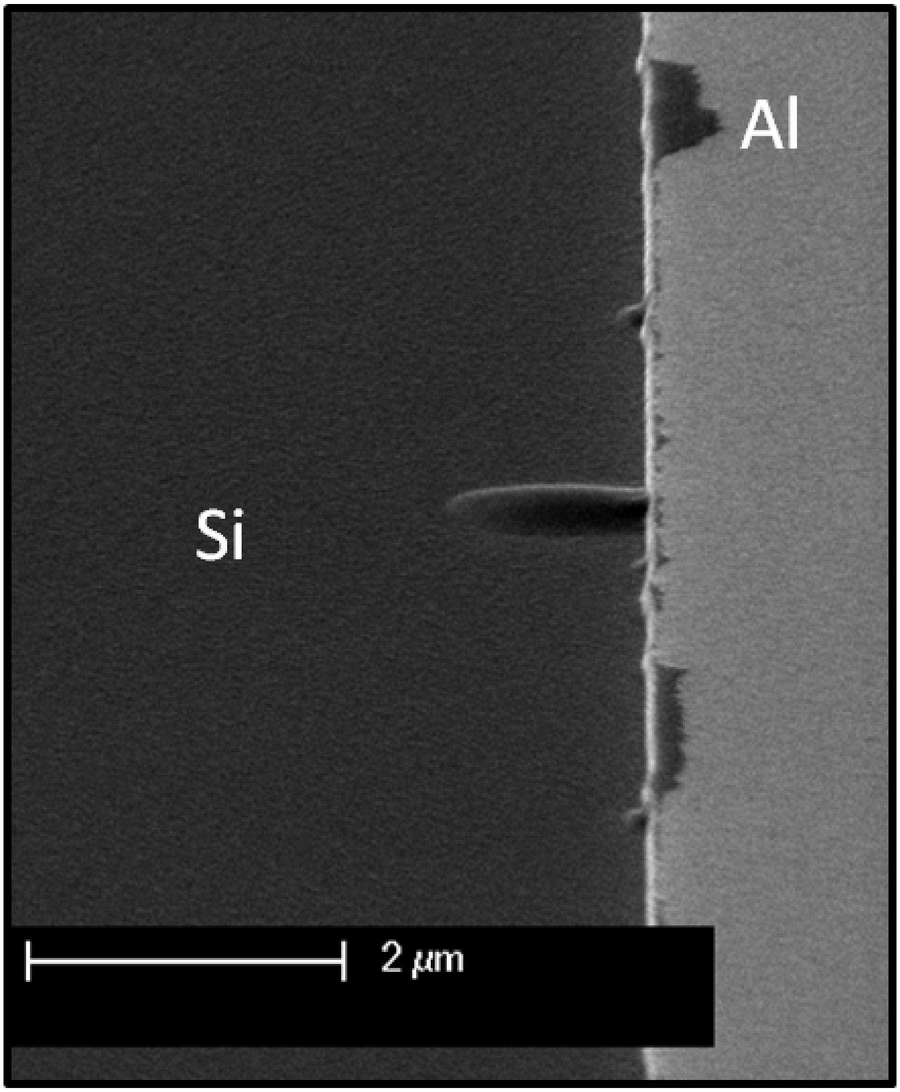

Figure 6.8: Scanning electron micrograph of polymer residue after performing a post etch CF_4 quench in an inductively coupled plasma reactive ion etcher. The formation of polymers is reduced below visible detection level when a SF_6 quench is used instead.

mask the area during the Cl etch. We have been unable to remove this polymer using heated NMP, acetone, ozone, downstream oxygen ashing, or nanostrip. Not surprisingly, this polymer appears to cause significant dissipation in resonators. Resonators measured with this polymer present had a 30% reduction in quality factor compared with resonators without any visible signs of this polymer.

The CF_4 chemistry has the potential to form polymers due to the higher carbon and lower fluorine availability compared with other fluorine chemistries such as SF_6. We have performed numerous tests which show a quench using SF_6 does not form the polymer, however there is a significant amount of etching (~100 nm)into the silicon substrate.

Spontaneous etching post dry etch

In the previous section I discussed the unintentional etching of Al caused by Cl byproducts, produced during a RIE etch, forming HCl once exposed to atmospheric water. The typical solution of submerging in DI water immediately after venting the system causes much more significant damage, as seen in 6.9. This process, which works well with sapphire, causes the aluminum to be etched up to 40 μm or more under the protective resist mask and can destroy the center trace of most CPW widths. The exact mechanism for this is unknown, but again it appears that the combination of Al and Si leads to an enhanced etching similar to the developer discussed earlier. Our current solution is to only use the SF_6 quench and do not perform any DI water treatment. Instead we immediately strip the photoresist in heated NMP solvent (see App. A. When following this procedure we do not see any etching.

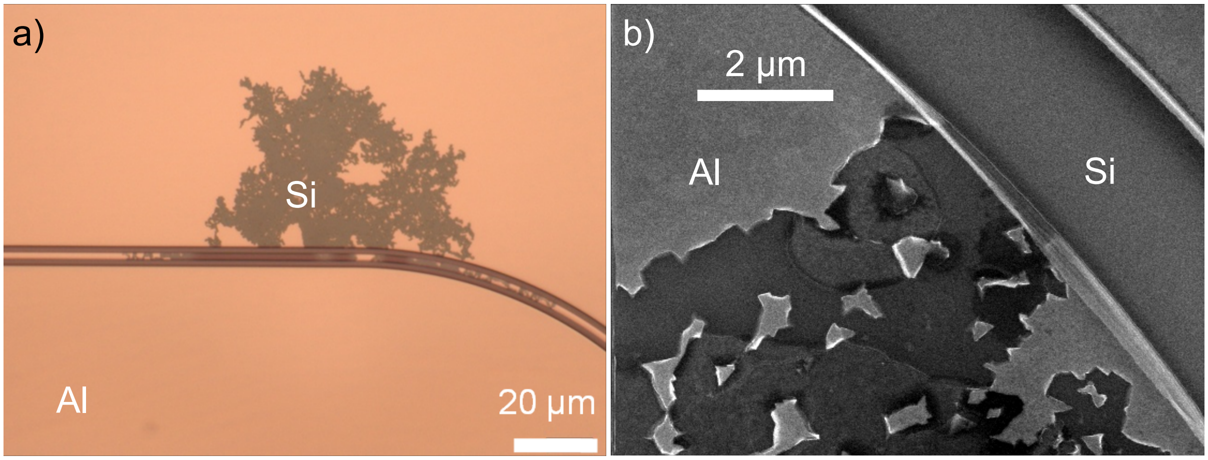

Figure 6.9: (a) Optical micrograph of Al-on-Si CPW where the Al has been unintentionally etch away. (b) Scanning electron micrograph of a similar unintentional etching of another CPW structure. The Al etching occurs while soaking in DI water. We soak the wafer in DI water immediately after removal from ICP etch chamber. The intention of this soak is to dilute the concentration of HCl which forms when the $AlCl_3$ etch products react with H_2O present in the atmosphere. On sapphire substrates the DI water reduces the etching of Al, while on Si the DI water soak seems to increase the etch rate. This issue is minimized by performing a fluorine quench before removing the wafer from the etch chamber. This is followed by immediately stripping the resist in solvents.

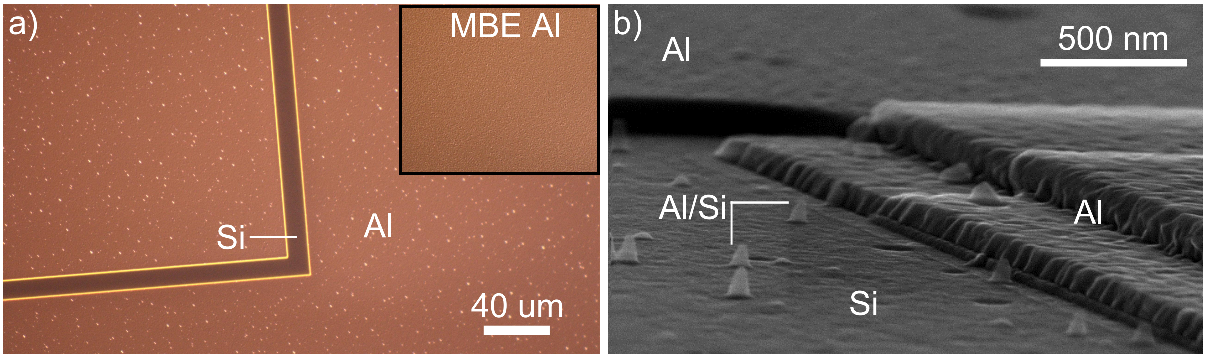

Figure 6.10: (a) Differential interference contrast micrographs and (b) scanning electron micrograph of a thin film of Al deposited on a hydrogen terminated Si(100) substrate. Several days to weeks after deposition hillocks would appear. Heating the wafer led to a significant increase in hillock density. (Inset) hillocks did not appear on wafers deposited in the MBE with an *insitu* pre-deposition UHV anneal (image is taken after 6 months of storage and heating during processing.)

Al-Si interaction: Hillocks

Interaction between Al and Si caused hillocks to form, as seen in Fig. 6.10. This happened on both hydrogen terminated Si and with the native oxide present. The number of hillocks increases dramatically if the wafer is heated above 200° C. Interestingly this did not appear on MBE heat treated wafers (Fig. 6.10(inset)). In the SEM image of Fig. 6.10(b) we see that hillock structures remain even after the Al is etched away.

6.3 Coupled Xmons: dedicated experiment devices

The proceeding devices contained several qubits on each die, where they where located far from each other to minimize coupling to other drive lines as well as qubits. The following devices are designed for performing experiments with specific algorithms. Both of these devices are at

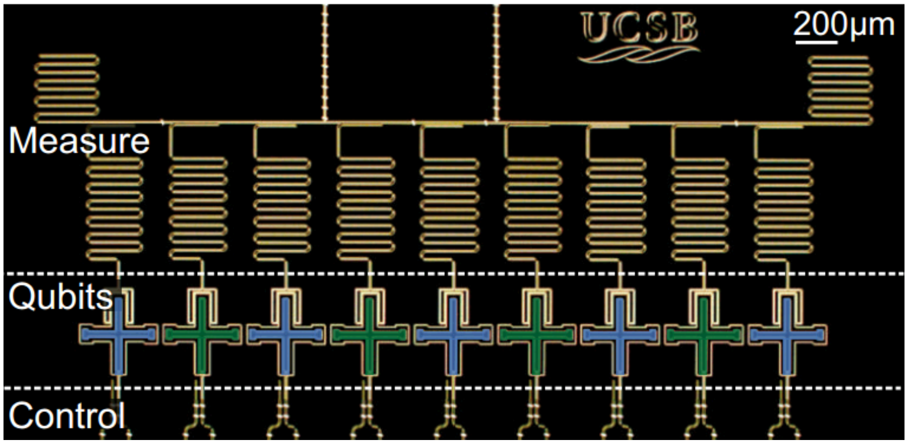

Figure 6.11: This figure reproduced from Ref. [28]. Darkfield optical micrograph of an Xmon on sapphire device. It contains nine qubits arranged in a linear chain with the dominate coupling between nearest neighbors. The capacitor for the qubits (shown in green and blue) is fabricated using high quality MBE deposited Al.

the frontier of superconducting qubit technology.

6.3.1 9-Xmon linear chain

Device

The device seen in Fig. 6.11 is the culmination of years of work for many researchers across several teams, whose individual successful experiments are combined to perform a first test of the fundamentals of quantum error correction based on the surface code [16]. This experiment uses fast high fidelity readout [117]; near quantum limited noise parametric amplifiers with large saturation power and bandwidth [118]; high fidelity gates [27]; which built upon high coherence Xmons [53] and resonators [42]. This device also hopes to answer if Xmons retain their high coherence when increasing the number of qubits and control lines.

The device consists of a 9-Xmon linear chain, as seen in Fig. 6.11 [28]. Each Xmon is capacitively coupled to a readout resonator which is inductively coupled to low quality factor $\lambda/2$ CPW resonator, used as a bandpass filter for fast high-fidelity readout. All of the Xmon CPW capacitors (shown in blue and green) are ($w = 24$ μm, $g = 24$ μm). Each Xmon has its own microwave drive and bias lines for individual control. The qubits Q_0 to Q_9 have increasing number from left to right.

The blue and green coloring of the Xmon capacitors denotes the intended use of each qubit in the repetition code. Blue coloring signifies it is a 'data' qubit and green for 'measure' qubits. The state of the measure qubits is readout during each cycle of the repetition code, while the data qubits are only readout at the end of the algorithm. For this reason the measure qubit readout resonators have stronger coupling to the bandpass filter and their maximum frequency is closer to the readout resonator.

Results

The T_1 spectra for all nine qubits are plotted in Fig. 6.12. For the data qubits (even numbers) we find the T_1s to be the best measured of any Xmons to date. T_1 ranges from several microseconds up to 100μs, with the average being ~30 μs. The spectra show a general improvement at lower frequencies, consistent with a constant quality factor, similar to the Al-on-sapphire Xmon described earlier in the chapter. The measure qubits show a steeper roll off at high frequency, likely from the Purcell limit of the strongly coupled low quality factor readout resonators that are optimized for fast readout. Some strongly coupled TLS are found in the spectra.

It is interesting to perform an energy loss budget with this sample, similar to the previous Xmon device where we identified the junction electrodes as a likely limitation. We can assume any additional loss associated with the junction electrodes is identical to the previous device since the fabrication procedure is identical.

Since we do not have measured resonator data at a similar size as the Xmon capacitor, we can estimate its value using the measured internal quality factor of a $w = 15$ μm, $g = 10$ μm MBE Al resonator and assume a similar scaling as the electron beam deposited Al control resonators. From Fig. 6.6(d) we see a 70% improvement going from $w + 2g = 35$ μm to $w + 2g = 70$ μm. The control resonators on previous MBE Al on sapphire devices had internal quality factors of 1.4×10^6. Now we can estimate that the Xmon capacitors in the current device have a quality factor of 2.4×10^6.

From the T_1 spectra of the nine Xmons we find an average measured quality factor of 8×10^5 (black dashed line in Fig. 6.12) across all nine qubits and over the majority of the frequency range. By subtracting the estimated loss of the capacitor from the average measured loss, we find a parallel loss channel with a quality factor of 1.2×10^6. This is in excellent qualitative agreement with the back of the envelope calculation, which estimated a quality factor of 1.2×10^6 for this loss channel; it is identical to the value calculated with the previous electron beam $w = 24$ μm Xmon. These results lend further evidence that the junction electrodes are an additional and dominate source of loss. If loss from this channel can be mitigated by both reducing its participation and improving the materials, we can expect to find T_1 consistently in the 100 μs range for MBE Al on sapphire Xmons.

Even with the most impressive Xmon T_1 spectra to date, frequency crowding was a concern with this device. The high fidelity two-qubit entangling gates used in the repetition code re-

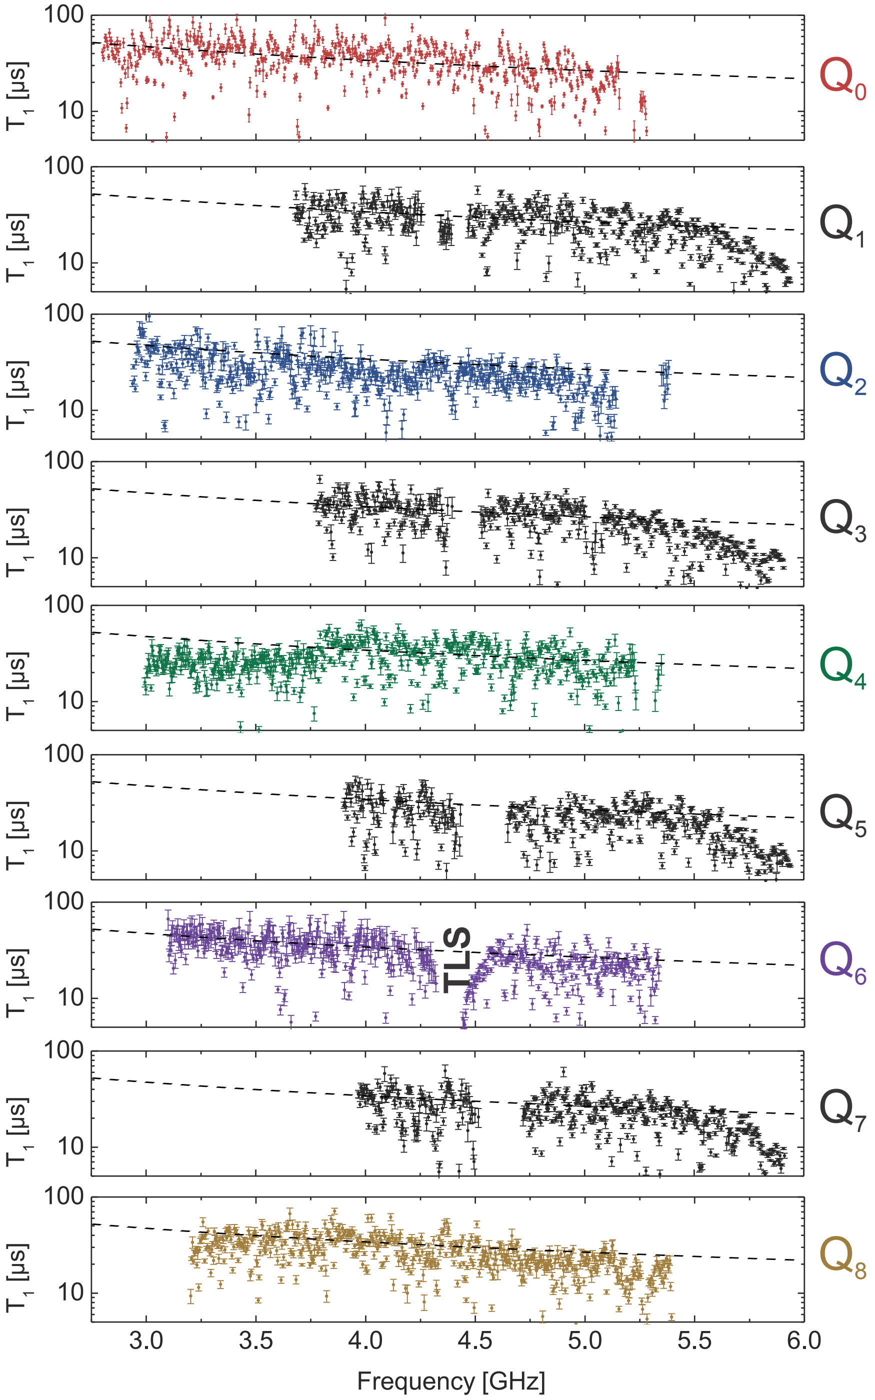

Figure 6.12: This figure reproduced from Ref. [28]. T$_1$ spectra for the nine Xmons shown in Fig. 6.11. The odd numbered qubits are coupled to readout resonators which have stronger coupling to the Purcell filter, which is the likely cause of the decreased T$_1$ at higher frequencies. The black dashed line represents a quality factor of $Q_{Ave} = 8 \times 10^5$ on all plots for reference. This is a 33% larger quality factor than the same CPW width capacitors as the electron beam deposited Xmon sample discussed earlier.

quire tuning the qubits through 800 MHz during the operation. Finding space in the frequency spectrum to perform the entangling gates with minimal interaction becomes problematic since there are nine qubits each with undesired couplings to next-nearest neighbor qubits and TLS. As the number of qubits continues to scale up, this problem will become unmanageable with this current quality of qubits. Improvements in materials are needed to reduce the density of TLS in the spectrum. A large improvement will allow a narrower CPW capacitor thereby minimizing next-nearest neighbor coupling.

6.3.2 3-gmon ring

Device

While the nine Xmon linear chain represents the state of the art for digital error corrected quantum computing, the three-gmon ring pushes the envelope further for analog computation. This device contains three gmon style qubits, with an adjustable coupler connecting each pair of qubits, as seen in Fig. 6.13 [119]. The added control of the adjustable coupling between neighboring qubits builds upon the versatility of the Xmon and allows for not only the qubit self-energy to be dynamically tuned but also the qubit-qubit coupling in very fast (nanosecond) timescales.

All metal layers of this device are deposited using electron beam evaporation. Unfortunately, the MBE was unavailable during the fabrication of the device. The gmon capacitors consist of a straight CPW segment with $w = 24$ μm and $g = 24$ μm. This is the first device with coupled qubits to include airbridges for crossovers, instead of SiO_2 or a:Si-H as has been done previously. Furthermore the complexity of this circuit required the use of airbridges to not only electrically short the many broken grounds in this device, but also be part of the qubit circuit as well as the readout CPW transmission line required to crossover the qubit and coupler control lines (Fig. 6.13(b) and (c)).

Airbridge crossovers allow for more efficient mutual coupling between the qubit and coupler linear inductors by sandwiching the coupler inductor (cyan) on both sides with the qubit inductor (green), as seen in Fig. 6.13(b) & c. Applying a flux in the coupler loop using the bias line seen in Fig. 6.13(c) changes the nonlinear inductance, L_c, of its junction. By varying L_c we can tune how much of each qubit excitation flows through the coupler, thus changing the coupling strength between qubits. For this device the coupling strength can be tuned between -16 MHz $< g/2\pi < +5$ MHz, where $2g$ is the splitting between energy levels.

Results

In Fig. 6.14 we see the T_1 spectra for the three qubits of this device. Even with the additional loss channel of the coupling structure, we find a respectable average T_1 time of $13.5\ \mu$s, with most in this frequency range having $5\ \mu\text{s} < T_1 < 20\ \mu\text{s}$. There is a trend of T_1 improving at lower frequencies, consistent with having a constant average quality factor $Q_{Ave} = 4.2 \times 10^5$ (represented by black dashed line in Fig. 6.14). These results indicate that the gmon architecture does cause additional loss when compared with the simpler Xmon structure.

The voltage divider between L_S and L_g determines how strong the electric fields located at the edges of the coupler, which allows us to determine the participation of lossy surfaces located in the coupler. In the original gmon design the additional coupling architecture participated

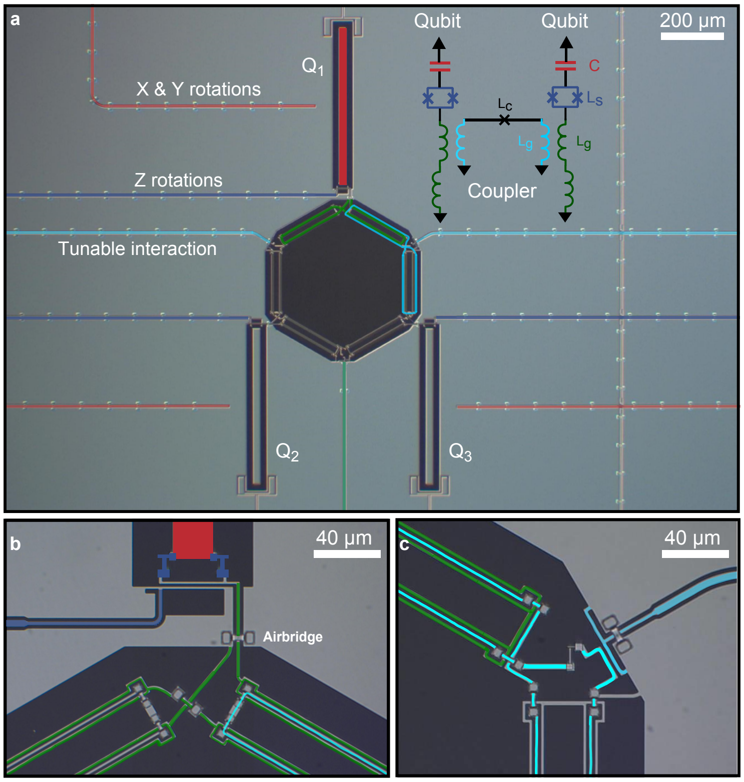

Figure 6.13: This figure reproduced from Ref. [119]. (a)Image of 3-gmon qubit chip. Individual microwave control of the qubits is achieved using capacitive coupling between the red lines and the large CPW gmon capacitor (Q_1 shown in red). Vertical silver CPW line is for readout of resonators which are capcitively coupled to the qubit capacitor. (b) Qubit SQUID loop and bias (blue). Linear inductor to ground for coupling network (green). (c) Adjustable coupler Josephson junction and bias (cyan). Note the ample use of airbridges required to improve coupling efficiency while maintaining low loss.

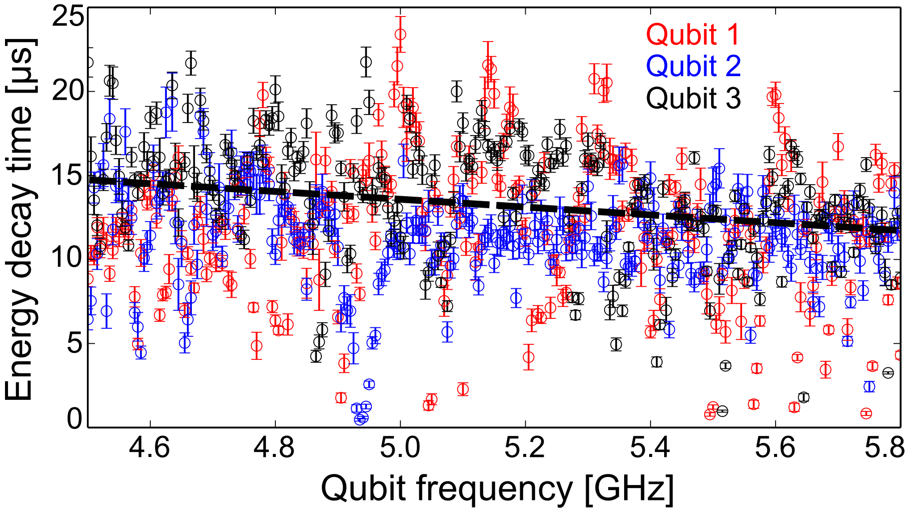

Figure 6.14: Energy decay time, T_1, of the three gmon qubits as a function of frequency. The black dashed line represents a quality factor of $Q_{Ave} = 4.2 \times 10^5$ and fits the data well over the frequency range. This quality factor is notably less than the electron beam deposited Xmons described at the beginning of the chapter.

$(L_S/L_g)^2 > 2000$, however for the current design $(L_S/2L_g)^2 > 130$ so it offers approximately a factor of twenty less protection. Additionally the participation of the narrow traces used in the coupling structure could be more than twelve times that of the wider CPW cross. After including the additional loss associated with the coupler junction electrodes, its stray junction as well as the airbridges, it is conceivable that the adjustable coupler structures contribute an additional loss channel comparable to the CPW cross. Understanding and mitigating this loss mechanism is a necessity if qubits will be coupled to more than two qubits in the future, as this will further decrease any protection offered by placing the coupler at a low voltage point of the circuit.

6.4 Conclusion

Planar Xmon qubits have achieved a high level of coherence. Xmons fabricated with MBE deposited Al perform roughly 33% better than standard e-beam evaporation. Using resonators, we have isolated several layers of Xmon fabrication and have identified their contributions to the total loss. The rapid and steady progress in qubits was spurred on through the use of witness and control resonators verifying that the fabrication induced loss was not the cause of the lower T_1 spectra of early devices. This allowed us to focus on the other dissipative elements of the designs. Further reduction of dissipation in these devices appears possible through improvement in the SQUID design and fabrication processes. This work has led to relatively complex circuits with high coherence that are capable of doing impressive demonstrations of quantum error detection and quantum simulation, the latter of which is discussed in chapter 7.

Chapter 7

Simulating quantum chemical dynamics with a quantum processor

7.1 Introduction

The possibility for atoms to form molecules is determined by the structure of the atomic orbitals. Molecular collision experiments are highly sensitive to atomic structure and provide a means for discerning between models of atomic orbitals. Predicting the results of these collisions requires computing the dynamics of strongly interacting electrons, a computationally difficult problem. In this chapter, using three fully-connected superconducting qubits, we generate the quantum dynamics resulting from a collision between He and Na atoms by encoding the problem into a time-dependent multi-qubit interaction. We measure the probability of exciting the Na atom at every instance in time for a wide range of collision parameters.

The scattering cross-section shows that there is a velocity where the excitation probability is maximized, in agreement with scattering experiments. The ability to generate arbitrary time-dependent Hamiltonians establishes our system as a platform to study dynamics underlying quantum chemistry.

7.2 Molecular chemistry is computationally challenging

Many physical and life sciences study the chemistry of molecular formation. Biologists study living organisms partly by investigating their fundamental building blocks - cells. Cells themselves are incredibly complex structures made up of molecules. Understanding the formation of these molecules as well as the interaction between molecules is very important. Is it possible to understand life from the molecular level up? The answer is clearly outside the scope of this thesis, however we can begin by studying simple molecule formation.

The Schrödinger equation provides the comprehensive framework for understanding the chemistry of molecules. Nevertheless, as Dirac noted, the application of it leads to complicated equations, which can only be tackled by numerical approaches. Since the inception of quantum mechanics, chemists have developed numerical methods for solving Schrödinger equation, which are frequently used in calculating binding energies and spatial structure of atomic and molecular orbitals. While implementing these methods are rather straightforward, reaching the desired accuracy are computationally demanding tasks.

These challenges led to a quest for finding alternative platforms that extend the abilities of commonly used supercomputers. Quantum computers are regarded as viable candidates for chemical computations, where various protocols have been proposed and debated [120]. Here, we present a programmable quantum chip and present the first experimental study of the dynamics of molecular interactions. Our implementation of this small quantum processor demonstrates the essence of the quantum simulator idea, which was originally envisioned by Richard Feynman [4].

Collision experiments provide a clear test of our knowledge of molecular orbitals. The simplest atomic collision to consider is between an inert noble gas with its closed valence orbital and an atom with a single valence electron, since there are fewer electrons and orbitals to consider in this system. Helium provides one ideal candidate as it has the largest energy gap between its ground and excited states of all the atoms, minimizing any interactions with these levels during the collision (Fig. 7.1). Many ions may contain a single valance electron, however only the neutral alkali metals do so naturally. Of these, sodium is an interesting choice to explore. Previously the interaction between sodium and helium was used to characterize the atmosphere of an extra-solar planet [121]. Additionally, Na and He have been used in previous collision experiments to determine the fundamentals of interactions [122].

Diabatic potentials and couplings, as described in Fig. 7.2, are used to describe the interaction of Na and He. In the center of mass basis, the interactions depend only on the separation between the atoms R. The shape of these potentials allows us to predict that there are no stable NaHe molecules in nature, since there is no local minimum therefore it is not energetically favorable to form a bound molecular state.

These potentials and couplings represent our best guess that describes the underlying physics.

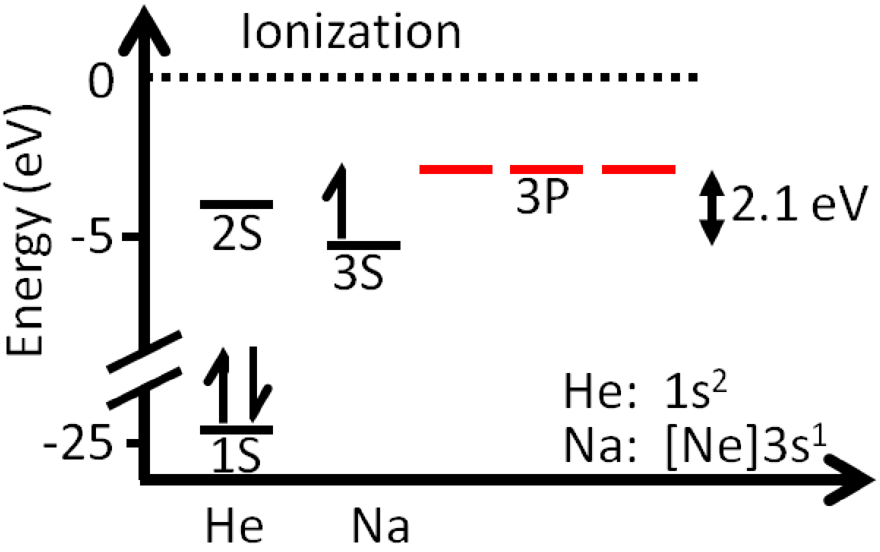

Figure 7.1: Depiction of the ground and first excited energy levels of Na and He relative to vacuum. The He atom has the largest energy gap of the noble gases and is an ideal candidate for our simulation since its electrons and higher orbitals can be neglected to high accuracy. The Na ground state has a single valence electron located in the Na($3s$) orbital. The first excited states are degenerate Na($3p$) orbitals which have an energy gap of 2.1 eV.

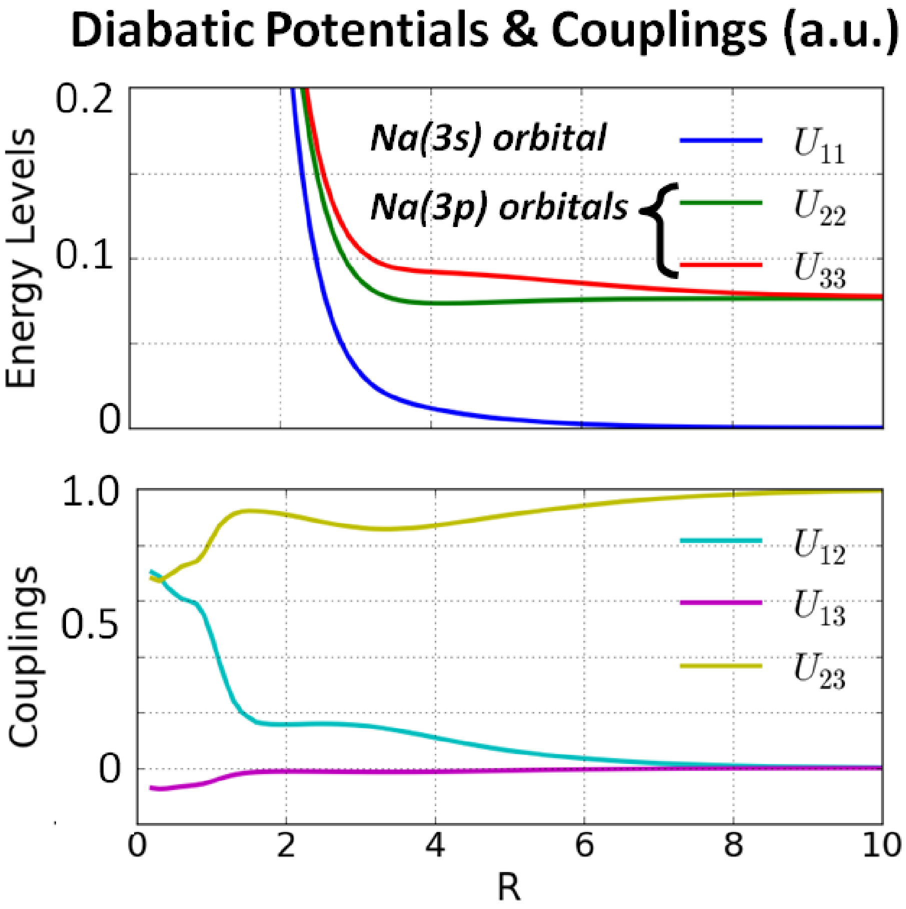

Figure 7.2: Previously calculated diabatic potentials and couplings as a function of the separation, R, between Na and He. At large separation the potentials associated with the Na$(3p)$ orbitals become degenerate and the coupling between the ground and excited states reduces to zero. Note the shapes of these potentials exclude the possibility of a NaHe molecule to form as a result of the collision.

Their accuracy is determined by how well they describe experimental observations. After these potentials have been generated, they need to be verified to produce the expected collision outcomes obtained from experiment. As the number of interacting electrons, orbitals, and atoms are increased, the computation quickly becomes classically intractable. Here we include only the interactions between the He ground state and the Na ground and first excited states. This three channel approximation is justified by previous experimental results described in the next section.

7.3 Olsen collision experiment

As mentioned in the previous section a previous scattering experiment was performed by Olsen *et al.*, who studied the collision between Na and He. They accelerated a beam of neutral Na atoms to collide with He target gas using the apparatus shown in Fig. 7.3 and described here for completeness. Ions were produced from NaCl which was evaporated in a separate oven. The Na ions were then accelerated between $0.6 - 60$ keV in the laboratory reference frame using either of two types of accelerators. The Na ion beam from the accelerator was partly neutralized by resonant charge transfer operated at ~5×10^{-3} Torr vapor pressure. Electrostatic deflector plates removed the remaining ionic component of the beam. A thermal beam detector measured the neutral beam power after it had passed through the collision cell, which contained the He target gas. A monochromator and photomultiplier tube measured the emitted photons of excited Na over the wavelengths $3000 - 8500$ Å, which encompassed the relevant transitions.

The probability of measuring exciting Na and measuring a photon depended on the energy

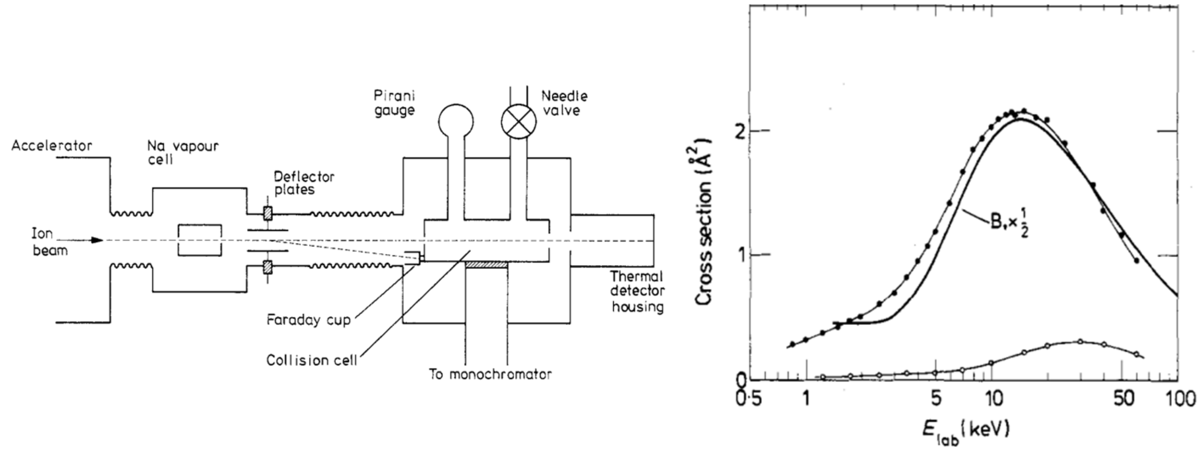

Figure 7.3: This figure is reproduced from Ref. [122]. (a) Schematic of original collider used in the scattering experiment. The apparatus creates an accelerated beam of neutral Na atoms which are then directed towards a collision cell which contains the target He gas. After a collision there is a probability to excite the Na valence electron. If the valence electron is excited from the collision it will emit a photon upon relaxation. The intensity of the emitted photons are measured and are proportional to the inelastic cross section. (b) Inelastic cross section as a function of Na beam energy in the lab frame for the Na transitions $3^2\mathrm{S}-3^2\mathrm{P}$ (closed symbols) and $3^2\mathrm{P}-3^2\mathrm{D}$ (open symbols). The wavelength of the measured photons distinguishes which transition has occurred. Note the peak in the cross section as well as the minimum contribution of the Na($3d$) orbitals.

of the Na beam. In Fig. 7.3(b) the inelastic cross section is plotted, with closed and open symbols corresponding to 3^2S-3^2P and 3^2P-3^2D respectively. We see that the 3^2S-3^2P transition dominates the cross section over the measured energy range. As the projectile energy is increased the measured 3^2S-3^2P cross section increases up to ~2 Å^2 when E_{lab} ~15 keV, above this energy the cross section decreased. The 3^2P-3^2D transition has an order of magnitude smaller peak in the cross section and occurs at higher energies.

The peak in the 3^2S-3^2P cross section is consistent with the Massey adiabatic criterion which predicts the cross section will be maximized when the time, τ, of collision is comparable to the frequency, ν, corresponding to the difference of the involved energy levels, ΔE, so when $\tau\nu \approx 1$ [123]. By substituting ν with $\Delta E/h$ and τ with a/v, where a is the "adiabatic parameter" a length of the order of the atomic dimensions involved, we produce the standard form of the adiabatic criterion

$$\frac{a|\Delta E|}{hv} \gtrsim 1. \tag{7.1}$$

In the remainder of this chapter I describe the simulation of this collision experiment over the same energy range as Olsen *et al.* using our 3 gmon ring quantum processor.

7.4 Mapping collision to qubits

As mentioned earlier, the most direct application of a quantum computer is to simulate other quantum systems, e.g., the collision of Na and He. Various algorithms can be used to perform these simulations, each with different hardware requirements. It is interesting to compare these different approaches to quantum simulation (e.g. digital vs analog), as this will change important platform requirements such as number of qubits, minimum gate fidelities, total simulation time, etc. For example, simulating this collision using a gate-based algorithm would require two qubits and more than 10^3 two-qubit gates. If we assume error rates of 10^{-3}, better than current state of the art, the fidelity of the simulation would be 0.36. So while a digital algorithm could be attempted using the current architecture, the chance of success is small. Significant improvements to gate fidelity or implementing quantum error correction (requiring many more qubits) would be required for this to be considered a viable simulation path.

An alternative simulation approach that is particularly interesting for simulating such problems was suggested in Ref. [19]. This approach uses analog control of the single excitation subspace (SES) of the qubits (e.g. $|100\rangle$, $|010\rangle$, and $|001\rangle$ for a 3 qubit device). This protocol has the capability to simulate any real, time-dependent Hamiltonian of an n-dimensional quantum system using an n-qubit quantum computer, in a simulation time that is independent of n. Simulating over the same parameter range of the Olsen experiment will require a minimum of 3-qubits and depends on the number of energy levels considered.

Perhaps the most difficult challenge in making an analog quantum processor is achieving accurate time-dependent control over the terms of the Hamiltonian of interest. A major challenge in this regard is having proper control lines at the user interface and establishing a one-to-one correspondence between them and the terms of the Hamiltonian at the processor. This requires independent calibrations to determine the system imperfections (see appendix B).

In various quantum platforms and experiments [124, 18], adiabatic dynamics were previously demonstrated; however, arbitrary time-dependent control far from the adiabatic limit has not been achieved. Another difficulty is preserving the quantum coherence of the system during the simulation. In fact, the two challenges are intertwined. The flexible design of superconducting circuits allows readily adding control lines to the device. However, these lines commonly have detrimental effects on the coherence, requiring careful design and fabrication. The success of this simulation experiment relies on overcoming both of these challenges.

We use the 3-gmon qubits device discussed in chapter 6 to perform the quantum simulation. The qubits are arranged such that each qubit is coupled to the two other qubits. The qubit-qubit coupling is mediated through an adjustable mutual inductance [54]. This design ensures minimal dissipation, since there is only a small potential difference of the linear inductors, comprised of narrow wires, relative to ground, as described in chapters 2 and 6.

The frequencies of the three qubits are independently tunable with a maximum frequency of 5.8 GHz down to near DC. The coupling strength $g/2\pi$ between qubits can be independently tuned between +5 MHz to -16 MHz going smoothly though zero, as shown in Fig. B.3. Both the qubit frequencies and coupling strengths can be varied on nanosecond time scales, allowing for very fast manipulations and arbitrary control shapes. These allows independent control of the 6 unique terms in a real 3x3 Hamiltonian

$$H(t) = \begin{bmatrix} \Delta_{11} & g_{12} & g_{13} \\ g_{21} & \Delta_{22} & g_{23} \\ g_{31} & g_{32} & \Delta_{33} \end{bmatrix} . \tag{7.2}$$

In order to generate the time-dependent collision Hamiltonian matrix elements, we need to project the full, many-body Hamiltonian into a similar form as Eq. 7.2. We use a 3-dimensional basis consisting of: Na($3s$) + He($1s^2$)[$1\,^2\Sigma^+$] and Na($3p$) + He($1s^2$)[$1\,^2\Pi^+$; $2\,^2\Sigma^+$]. This includes only the interaction between the He and Na ground states and He ground state with Na($3p$) orbitals (Na's first excited states). This approximation is justified by the experimental cross section data.

We use the diabatic potentials and nonadiabatic couplings which depend only on internuclear distance R, and assume a straight-line trajectory, $R(t) = \sqrt{b^2 + v^2t^2}$, where v is the velocity of the reduced mass and b is the impact parameter of the collision. The matrix elements for a collision with $b = 0.053$ Å and $v/c = 7.3 \times 10^{-4}$ are shown in Fig. 7.4(a). The majority of interaction happens within a few femto-seconds and with very large coupling strengths, neither of which are achievable with our quantum processor. Therefore we need to rescale the magnitude of the matrix elements.

7.4.1 Rescaling Hamiltonian for efficient simulation

The matrix elements seen in Fig. 7.4(a) need to be rescaled while keeping the unitary operator unchanged. The unitary propagator U is generated from the time dependent Hamiltonian by

$$U(t) = Te^{-i\int H(t')dt'}, \tag{7.3}$$

where T is the time-ordering operator. Without loss of information, we can include a scaling parameter λ in the integrand such that $H_{QC} = H(t)/\lambda$ and $dt_{QC} = dt \times \lambda$, which leaves the unitary propagator unchanged. For example if $\lambda = 10^9$, then the off-diagonal coupling matrix elements are $|g(t)/2\pi| < 5$ MHz and easily achievable with our quantum processor. However the total simulation time on the quantum computer now becomes 25 μs which is much longer than our decoherence times T_2~3 μs at the operating point. We could reduce the simulation time by truncating the simulation to include only small R.

A more complete approach, described in Ref. [19], allows λ to vary in time, as seen in the inset of Fig. 7.4(a). Essentially this maximizes the available resources by ensuring at least one of the matrix elements is always at a maximum operating value during the entire simulation. This protocol effectively speeds through the simulation when the atoms are far apart and slows the simulation during the collision when the couplings are largest. The resulting matrix elements are plotted in Fig. 7.4(b), where the matrix elements in Fig. 7.4(a) are rescaled using $\lambda(t)$ displayed in the inset.

We determine $\lambda(t)$ by choosing an appropriate maximum for the coupling strength g_{max} and the maximum detuning of the qubits from their initial frequency Δ_{max}. Both of these are selected by weighing the trade-offs between coherence and control.

The adjustable couplers have a much larger negative coupling strength compared to positive coupling. However reliably achieving these large coupling strengths for an arbitrary control pulse is difficult. This is due to imperfect calibrations and the large $\frac{dg}{d\Phi}$ at the larger coupling strengths. We set $|g_{\text{max}}| = 4.2$ MHz. This is the largest positive coupling strength for one of the couplers.

In order to determine Δ_{max}, we need to first choose the operating frequency of the qubits. This requires finding a region in qubit coherence spectrum which has little or no defects and is also closest to the flux insensitive point of the qubits to minimize dephasing during the simulation. Optimizing between these trade-offs, we chose to park the qubits at 5.65 GHz, approximately 70 MHz below the lowest flux insensitive point of the qubits. We then set

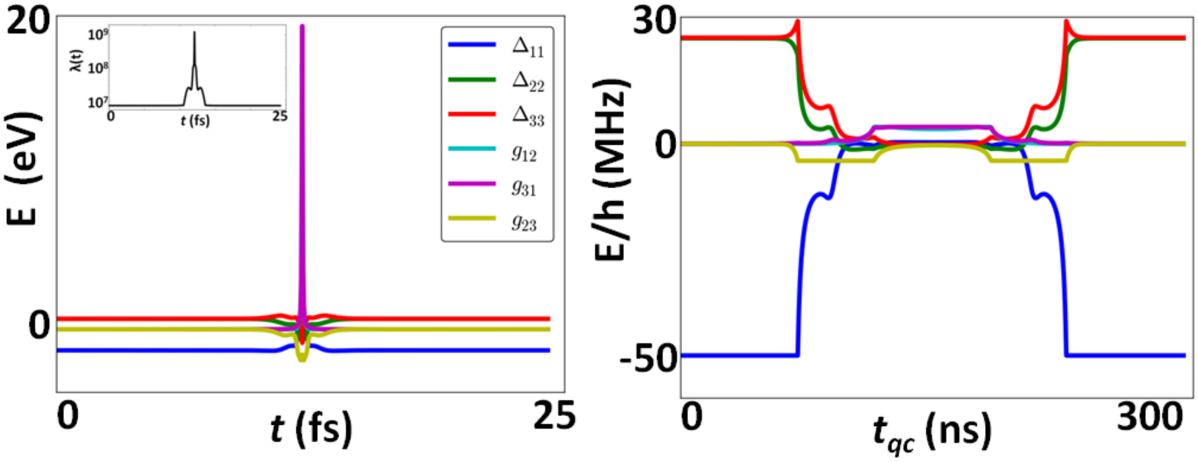

Figure 7.4: (a) The six independent Hamiltonian matrix elements ,as a function of time, for the collision of Na and He. The matrix elements are constructed in the three molecular channel basis and using the diatomic potentials and couplings plotted in Fig. 7.2. (Inset) Scaling parameter, λ, as a function of collision time. Note the highly nonlinear scaling allows for at least one matrix element to have a maximum $|g|$ or $|\Delta|$ the system is capable of. This allows for a much faster simulation time compared to constant scaling. (b) The same collision Hamiltonian matrix elements as (a) after rescaling using the $\lambda(t)$ plotted in the inset.

$|\Delta_{\text{max}}| = 50$ MHz to avoid any interaction with strongly coupled TLS.

Now that the collision Hamiltonian has been re-scaled for efficient simulation on our quantum processor, we need to characterize the imperfections in our experimental setup. We can then determine the signal required from the room temperature electronics in order to produce the desired Hamiltonian matrix elements at the device. This is done through a series of independent calibrations described in appendix B.

7.5 Results

7.5.1 Collision Dynamics

The voltages plotted in Fig. 7.5 are applied to the coupler and qubit bias lines for a simulation with parameters $v/c = 7.3 \times 10^{-4}$ and $b = 0.053$ Å. The legend on the left side of the figure depicts these six bias control lines. The voltages are determined by first generating the rescaled Hamiltonian matrix elements and then applying corrections based on the independent calibrations described in appendix B.

The pulse sequence of the simulation begins by populating q_1, the qubit representing the molecular channel with Na and He in their ground states. Next we apply the six time-dependent matrix elements. Finally we simultaneously measure the population of all three qubits. We run up to 5000 trials to measure the probability distribution of the excitation's location. Additionally, we can stop the evolution and measure the current probability distribution at any time during the collision. It is interesting to note that our platform allows the simulation to begin in the excited Na states (corresponding to q_2 and q_3), something which is much more difficult to achieve in a scattering experiment.

The results of these measurements are plotted in Fig. 7.6. The plotted probabilities are from averaging over 2500 trials and we ensure that the three probabilities sum to 1 by correcting for measurement visibility and T_1 decay.

For the first 60 ns of the simulation, the excitation remains in qubit q_1. After this time, the couplings increase, which initiates swapping of the excitation between the three qubits/chan-

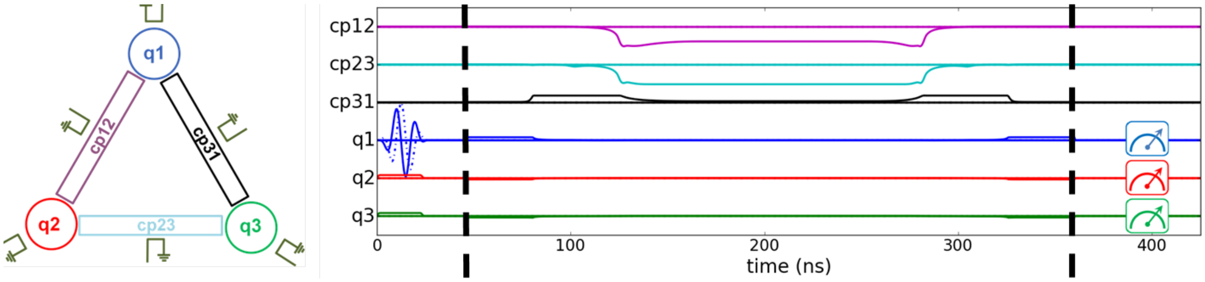

Figure 7.5: (a) Depiction of 3-gmon ring emphasizing the six independent controls, three for the individual qubit self-energies and three for the adjustable couplers. (b) The pulse sequence for a collision simulation with parameters $v/c = 7.3 \times 10^{-4}$ and $b = 0.053$ Å. We initialize the system by applying a pi pulse to q_1 which represents both atoms in their ground states. We then perform the simulation by dynamically applying the six control biases during the evolution. Here we have converted the matrix elements of Fig. 7.4(b) into applied voltages using the calibrations discussed in appendix B. Finally we measure the three qubits. We then repeat this process to obtain enough statistics to determine the transition probabilities of the collision.

nels. This complex pattern provides insight into the quantum dynamics occurring during the collision that are not observable in the scattering experiment. After 320 ns we see that the couplings have returned to zero indicating that the atoms are again well separated post-collision. The simulation of these collision parameters predicts a 60% probability of the excitation leaving the initial channel of the Na and He ground states and ending in the Na($3p$) orbital. This relates to the scattering experiment by noting that there is a 60% probability of a photon being emitted for possible detection.

Since the quantum processor contains only three qubits, we are able numerically compute, on a classical computer, the quantum evolution dynamics. This provides a benchmark to compare our quantum processor against. We construct the unitary propagator by means of Suzuki-Trotter expansion: dividing the total simulation time into 1000 bins where within each time bin the matrix elements are assumed constant [125]. We then evolve the initial state for-

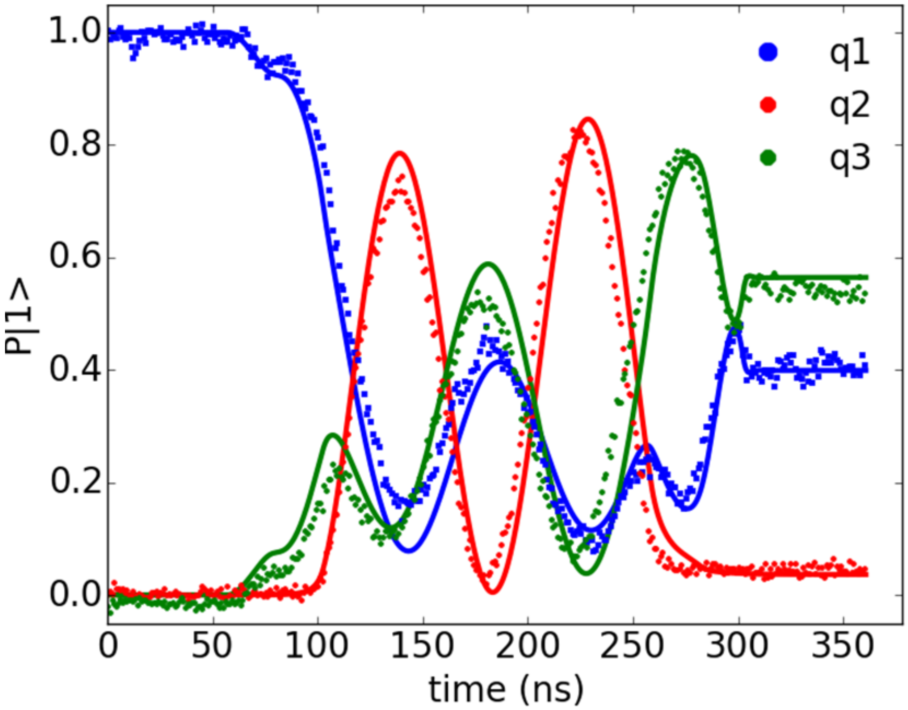

Figure 7.6: Probability of measuring the qubit in the excited state (P$|1\rangle$) as a function of simulation time for all three qubits. See Fig. 7.5 for corresponding pulse sequence. Each data point is the average of 2500 trials after stopping the quantum evolution at the specified simulation time. The data has been corrected for measurement visibility and T_1 decay. The solid lines are the result of numerically computing the evolution dynamics by trotterizing the time-dependent Hamiltonian. Therefore they represent an ideal version of our quantum processor. The deviations between the data and the solid lines are due to imperfect calibrations and corrections of pulse distortion and qubit dephasing during the simulation time.

ward through the collision using these discrete unitary propagators and plot the evolution as solid lines in Fig. 7.6. We find excellent agreement over the full simulation time between the ideal version of our quantum processor and the actual device performance. This shows that at the current level of qubit coherence, calibration routines, and control electronics we are able to accurately simulate the quantum dynamics of this collision.

7.5.2 Calculating the inelastic cross section

The scattering experiment cannot measure the quantum dynamics that we can explore with our quantum processor. Their measurement data is limited to the detection of photons, i.e., the final probabilities. For comparison we next simulate only the final excited state probabilities over the same energy range as the Olsen experiment. This spans over two orders of magnitude in projectile energy, corresponding to an order of magnitude variation of collision velocities, and a wide distribution of impact parameters. In total we simulate 10^4 distinct parameter combinations with 100 linearly spaced steps in both center of mass velocity, v, and impact parameter, b, the results of which are plotted in the left column of Fig. 7.7. The large number of distinct parameter combinations and therefore unique quantum evolutions preclude us from individual optimization, thus necessitating methodical calibrations. For comparison in the right column of Fig. 7.7, we plot the corresponding numerical computation done on a classical computer over the same parameter regime.

For large impact parameters, we expect there to be no interaction or population exchange between the states because the atoms effectively never "collide". For the range of velocities studied, we find that no population is transferred into the other qubits/molecular channels when $b > 3.0$ Å. For 1.0 Å $< b < 3.0$ Å we measure the greatest population transfer when v/c ~1.4×10^{-3}, with roughly equal probability of finding the excitation in qubit q_2 or q_3. For other velocities in this range of impact parameters, we find a decrease in the probability of exciting Na. When $b < 1.0$ Å the collision results in a substantial (75%) probability of exciting Na over the full range of measured velocities. When $v/c < 0.8 \times 10^{-3}$ the excitation is mainly

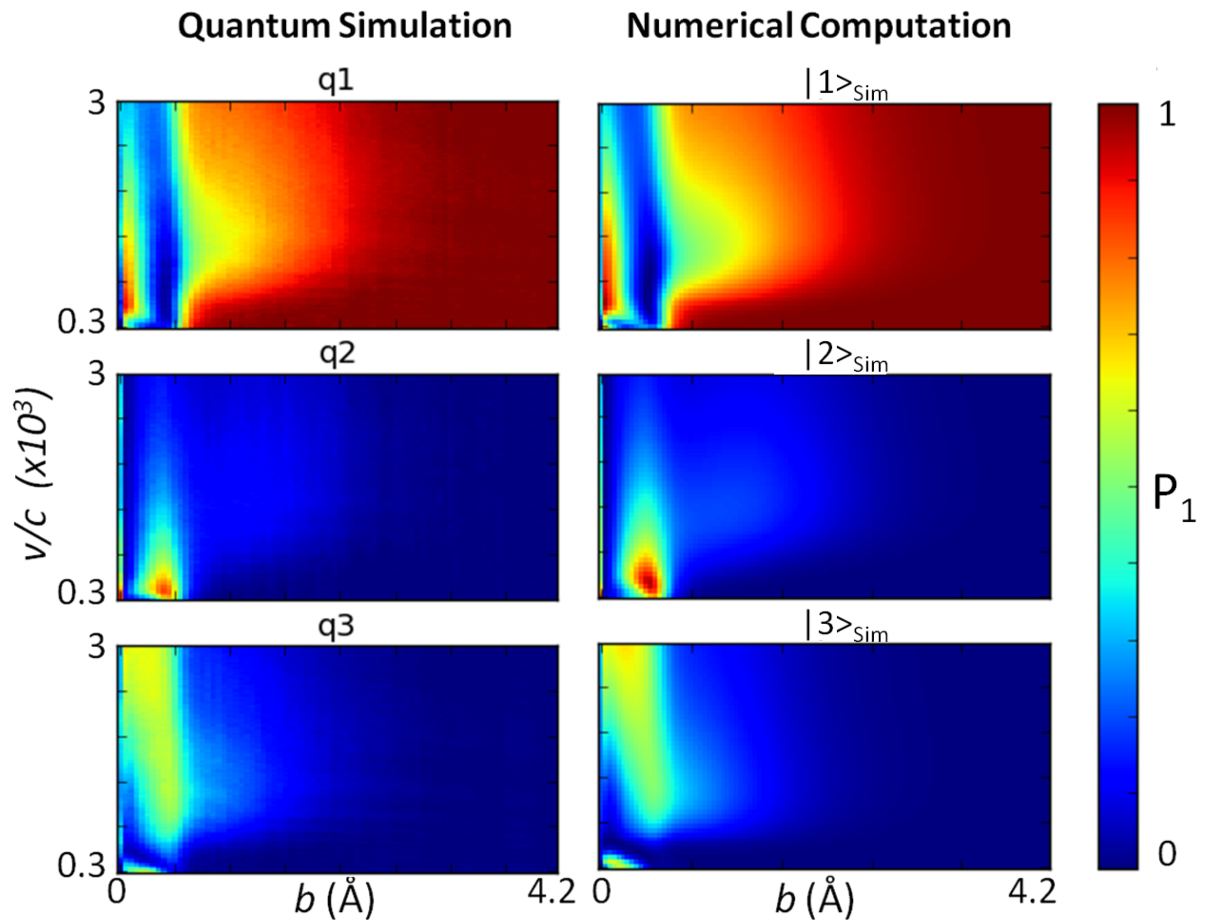

Figure 7.7: P$|1\rangle$ (false color) as a function of collision impact parameter, b, and velocity, v. The left column are data from simulations performed on our quantum processor after correcting for readout fidelity and T_1 decay. The right column are numerical computations of the ideal version of our quantum processor. We find excellent agreement over the full simulation range.

found in q_2. Above this velocity we find the excitation is in q_3.

We now have the required simulation results to calculate an inelastic cross section. First we note that the scattering experiment did not have the ability to measure an individual collision's impact parameter. Inherently they are sampling over all impact parameters by using a beam of many Na atoms colliding with the He gas. For comparison with the experimentally determined cross section, we integrate the final probability of exciting the Na atom weighted by the impact parameter

$$\text{Cross section} = 2\pi \int_0^\infty b \times (1 - P_{|1\rangle}) db. \tag{7.4}$$

In Fig. 7.8 we plot the inelastic cross section as a function of the center of mass kinetic energy. The original scattering experimental data has been replotted here in the center of mass frame (red squares). The numerical computation performed on a classical computer representing the ideal version of our quantum processor is shown as the solid black line. Qualitatively, the numerical computation agrees very well with the experimental data, including the position of the peak and roll-off at higher and lower energies. Quantitatively the numerical computation predicts approximately twice the cross section measured by the Olsen experiment over the full measurement range. This discrepancy may be due to incorrect diabatic potentials or systematic errors in the original experimental data.

We find excellent agreement between the inelastic cross section measured using our quantum processor and the numerical computation over the full range of the original scattering experiment. The mean absolute percent error is 8.7%, while the maximum percent error is 15.8%. We hypothesis that the discrepancy between the numerical computation and the quantum processor is from dephasing during the simulation time. The scattering experimental data measures the peak in the cross section at 2.1 keV, in agreement with the adiabatic criterion discussed earlier. For the numerical computation and quantum processor datasets, we find the peak positions to be E$_{peak}$ ~2.6 keV.

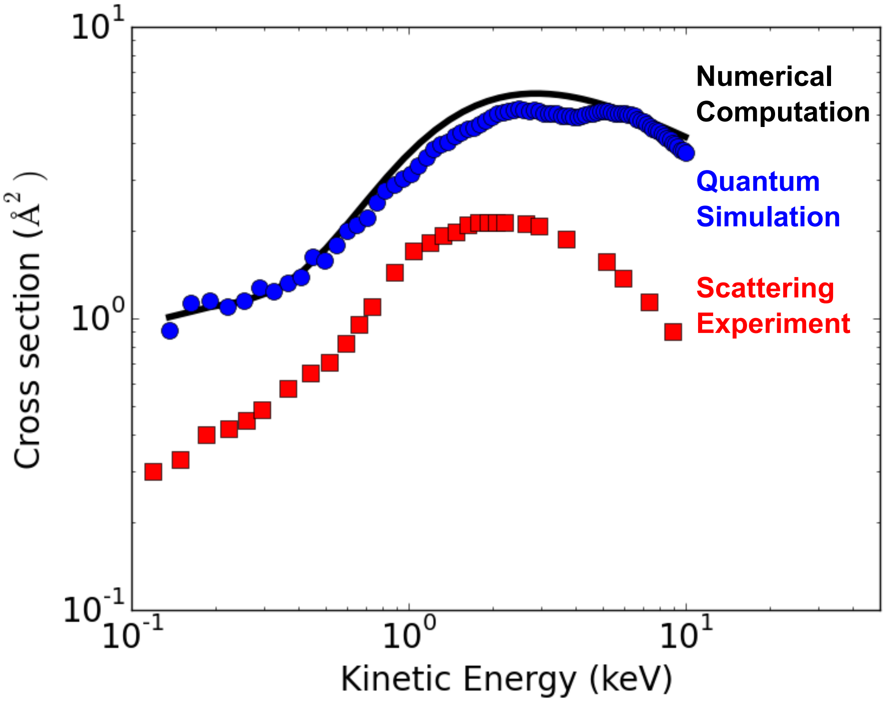

Figure 7.8: Inelastic cross section of Na and He collision as a function of center of mass (COM) kinetic energy. The red squares are the original Olsen collision experiment replotted in the COM frame. The solid black line is the cross section from the numerically computed data of Fig. 7.7 after integration according to Eq. 7.4. We perform the same integration of the data obtained using our quantum processor (blue circles). The deviation between the numerical computation and the original Olsen collision experiment may be due to imperfect diabatic potentials or possible systematic error in the collision experiment. The deviation between the simulation performed on our quantum processor and the numerical computation is due to imperfect calibration of pulse errors or more likely dephasing of the qubits during the simulation.

7.6 Conclusion

We have simulated the inelastic collision of Na and He using a superconducting quantum processor containing three fully connected gmon qubits. We simulated the collision over the same range of kinetic energies as a previous scattering experiment. We find that both the original experiment and our quantum simulation produce a peak in the cross section in agreement with the adiabatic criterion. Furthermore the quantum simulation predicts the experimentally measured cross section to within a factor of two over the full range of energies, and within 8.7% of the numerical computation. The accuracy of this quantum dynamics simulation shows the maturity of superconducting qubit control technology and paves the way for potentially solving more complicated problems using superconducting qubits.

Chapter 8

Conclusion and Outlook

8.1 Conclusion

A scattering cross section of the inelastic collision between Na and He atoms has been simulated using a superconducting quantum processor. The generated quantum chemical dynamics as well as the final scattering cross section agree very well with numerical computations performed on a classical computer and a previous scattering experiment. Generating the cross section required fast arbitrary time-dependent modulations of both the individual qubit frequencies and the qubit-qubit coupling strengths. This allows complete control over the single excitation subspace of the three qubits. The vast parameter space associated with generating the cross section essentially eliminated the possibility of fine-tuning the individual pulse sequences. Instead we developed a series of independent calibrations to minimize the effects from pulse distortions and other non-idealities.

The success of these simulations required high coherence quantum elements while maintaining controllability. Similar requirements are necessary for testing the fundamentals of quantum error correction. We developed the high level of qubit coherence over several years using both simpler qubit circuits and superconducting resonators. We have measured hundreds of resonators over this time. The results have clearly shown the importance of improving both the metal-substrate interface as well as the exposed interfaces by cleaning with a non-destructive method, for example without physical ion bombardment. These improved interfaces have led to better quality factors on both single crystal sapphire and silicon substrates. Besides dielectric loss, resonator experiments have improved our understanding of other energy loss mechanisms coming from quasiparticles, magnetic vortices, and radiation. We applied this knowledge to improve the qubit design, fabrication, and measurement, leading to a thirty-fold improvement in T_1 times over a two year period.

In this work we have achieved record high quality factors on silicon substrates ($Q_i^{LP} > 5 \times 10^6$). These resonators exhibit a previously undiscovered temperature dependence of the quality factor which is not predicted by the standard tunneling model of two level systems (TLS). Instead these results are consistent with a model of interacting TLS. Further research is required to better understand and utilize this behavior.

8.2 Outlook

The work in this thesis has led to a level of coherence and control where we can perform nontrivial analog simulations and test the fundamentals of quantum error correction. More complicated circuits and algorithms may be developed using current technology. However this will limit the number of Xmon qubits to 30 - 50. Algorithms performed on several tens of qubits are unlikely to perform better than classical computers. The prevailing view is many more qubits will be required in order to compete with the immense classical computing power available. With more qubits the probability of an error occurring during an algorithm increases at least linearly with the number of qubits, assuming uncorrelated errors. In this context it is unclear if physical qubits will ever have enough coherence and control to compete with classical computers directly, so error correction may be a necessity. In either case, increasing the number of qubits and improving their coherence are the top priorities for future work.

8.2.1 Further improvements in qubit coherence

The discrepancy between the results of Xmons on sapphire compared to those on silicon highlights the need for an improved understanding of the microscopic nature of dielectric loss. In lieu of this knowledge, we have used resonator experiments as testbeds to reduce losses. While this has led to many improvements, these experiments still take significantly more time than many other material characterization techniques. Furthermore a better microscopic model for dielectric loss may lead to more easily identifying good material systems to test, narrowing this huge parameter space.

In chapters 3 and 4 I described the significant effort to reduce energy loss through high quality materials. Poor interfaces can compromise the use of these ideal materials and dominate the loss of qubits and resonators. While it is clear that producing thin clean interfaces is vital to minimizing their contribution to loss, the microscopic source of the loss is still not well understood. Further characterization of these interfaces beyond the techniques used in this thesis are required, the following are some techniques which may provide some additional insight.

X-ray photoelectron spectroscopy (XPS) may be used to measure the substrate surface after various cleaning protocols to learn what contaminates may be present before deposition of the superconductor. Additionally, measuring the exposed surfaces of the device just prior to cooling down for measurement will inform what contaminates these surfaces. Interpretation of these measurements will have some uncertainty as it will not account for changes to the surfaces due to the deposition of the superconductor as well as adsorbates which may freeze onto the exposed surface of the device during cooldown. Another researcher began a study using this technique to investigate contaminates remaining after removing the native oxide of Si using various procedures.

Thermal desorption spectroscopy (TDS) may also provide information of the contaminates on the surfaces. TDS involves heating the sample while measuring the desorbed molecules using a quadrupole mass spectrometer or a time-of-flight mass spectrometer. This method may reveal both the nature of contaminants as well as the temperature where they are effectively removed. Systematically fabricating resonators whose substrates are heated to precise temperatures that removes only certain contaminates may prove very enlightening.

In order to investigate the superconductor-substrate interface, alternative characterization methods are required that can directly measure this buried interface. Some promising techniques include electron energy loss spectroscopy (EELS), energy-dispersive X-ray spectroscopy

(EDS), or dynamic secondary ion mass spectrometry (DSIMS). The first two techniques involve creating a lamella using a focused ion beam (FIB) and subsequent analysis using a scanning transmission electron microscopy (STEM). The preparation and handling of the sample may influence the results. Additionally these methods can not be used to determine hydrogen contamination, a possible candidate for TLS leading to dielectric loss[126].

Detecting hydrogen experimentally is very difficult, especially small concentrations located in a thin interfacial volume. One possibility DSIMS uses a continuous beam of ions to sputter material away while some fraction of this material is ionized and measured using a mass spectrometer. However since hydrogen is typically the major constituent of the residual gas in vacuum systems (either in the form of H_2O for HV or H_2 for UHV) the SNR may be to small to detect small amounts of hydrogen contamination located in thin interface.

The predominant focus of this thesis has been improving energy dissipation in qubits and resonators. This progress has not directly increased the dephasing times of Xmons significantly. Most of the previous work to understand dephasing has been performed with SQUIDS that have relatively simple geometries in comparison to the SQUID of the Xmon qubit. Understanding the differences between these devices will inform of possible near term improvements. However even flux noise in SQUIDs is still an open problem. Flux noise may also benefit from a similar systematic approach in improving the materials and interfaces used to fabricate these devices.

8.2.2 Scaling up to larger qubit systems

A fault-tolerant universal quantum computer will likely require more than one million physical qubits. This is a daunting task and clearly beyond the scope of this thesis. However, given the current technology, one can envision scaling up to several hundred coherent qubits, a relatively large step from the nine Xmons already achieved. Even this will require innovations in device packaging, control wiring and electronics, reducing cross talk, and fabrication stability.

Our current packaging uses bulky SMA connectors around the perimeter to control the device. This is clearly not scalable to even 50 individually controlled qubits. Moving to smaller microwave connectors which integrate through the face of the package will allow the number of connectors to scale with the area rather than the parameter of the package. Work in this direction is currently underway. The wire bond connections that bring the control signal on to the chip have the same scaling issue. Moving to larger chips to increase the number of control wires will likely lead to lower frequency box modes, degrading the performance. Additionally, routing the control wires to the qubits becomes significantly more difficult when moving to a 2-D array of qubits.

Currently airbridges are the most sophisticated on-chip, low-loss wiring technology used with Xmon and gmon qubits. However their utility is still rather limited. As more control lines are routed past and cross over other circuit elements, crosstalk becomes an important issue. More sophisticated 3-D integrated wiring schemes, such as flip-chip bonding and through substrate vias (TSV), are needed for a 2-D grid of coupled qubits to each retain their own individual control and measurement.

Manually calibrating three gmons and adjustable couplers or nine Xmons is easily accomplished, although somewhat time consuming. Hundreds of qubits will require diligent work to minimize crosstalk, simplifying the task as much as possible. The 3-D wiring integration methods may help with this issue. More sophisticated fabrication will allow additional ground planes to shield circuit elements from one another. We need to understand how this will effect qubit coherence. This additional processing will likely lead to less predictable qubit parameters, such as frequency and coupling. Understanding these trade-offs is vital for successfully scaling to hundreds of coherent qubits.

Appendix A

Nanofabrication

A.1 Fabrication Processes

A.1.1 Qubit fabrication overview

Note the substrate dependent steps.

1. Control wiring
    (a) Clean substrate and deposit MBE Al as in section A.1.2.
    (b) Optical lithography as in section A.1.7.
    (c) ICP etch as in section A.1.3.
2. Au alignment marks for e-beam write
    (a) Optical lithography as in section A.1.7.

   (b) Deposit (Ti 10 nm)/(Au 150 nm) using liftoff process.

3. Crossovers
   (a) SiO_2 crossovers as in section A.1.5 (Sapphire).
   (b) Airbridge crossovers as in section A.1.5 (Silicon).
4. Xmon capacitor etch
   (a) Optical lithography as in section A.1.7.
   (b) ICP etch as in table A.1.3.
5. Quarter wafer (optional)
   (a) Dice wafer into quarters (Sapphire/Si(111))
   (b) Cleave wafer into quarters (Si(100))
6. Josephson junctions
   (a) Spin e-beam resist as in section A.1.4 Resist stackup.
   (b) Expose with parameters from section A.1.4 E-beam writer.
   (c) Deposit junctions as in section A.1.4 Deposition process.

A.1.2 High quality aluminum deposition

See chapter 3 (4) for a through description of the MBE cleaning and deposition processes for sapphire (silicon) substrates.

Table A.1: BCl_3Cl_2 etch recipe for Al films.

	Step 1	Step 2	Step 3	Step 4	Step 5
BCl_3	20 SCCM	20 SCCM	20 SCCM	0 SCCM	0 SCCM
Cl_2	40 SCCM	40 SCCM	40 SCCM	0 SCCM	0 SCCM
CF_4	0 SCCM	0 SCCM	0 SCCM	50 SCCM	50 SCCM
Pre.	3.0 Pa	0.7 Pa	0.7 Pa	2.0 Pa	2.0 Pa
SRC FWD	300 W	300 W	300 W	700 W	700 W
BIAS FWD	0 W	0 W	70 W	0 W	20 W
Step Time	5 s	5 s	33 s	5 s	5 s

Table A.2: BCl_3Cl_2 etch recipe for trenching into sapphire substrate.

	Step 1	Step 2	Step 3	Step 4	Step 5
BCl_3	20 SCCM	20 SCCM	20 SCCM	0 SCCM	0 SCCM
Cl_2	20 SCCM	20 SCCM	20 SCCM	0 SCCM	0 SCCM
CF_4	0 SCCM	0 SCCM	0 SCCM	50 SCCM	50 SCCM
Pre.	3.0 Pa	0.3 Pa	0.3 Pa	2.0 Pa	2.0 Pa
SRC FWD	900 W	900 W	900 W	700 W	700 W
BIAS FWD	0 W	0 W	100 W	0 W	20 W
Step Time	5 s	5 s	90 s	5 s	5 s

A.1.3 Etching

Sapphire substrates

We etch the Al film using a Panasonic E6261 inductively coupled plasma (ICP) dry etcher, with a BCl_3Cl_2 chemistry, with parameters shown in table A.1. Immediately after etching, the wafer is immersed in DI water for 10 minutes to dilute any residual Cl that is embedded in the resist before it can react with water in the atmosphere to form HCl and cause nibbling. Photoresist is cleaned via sonication in Acetone, IPA and then the wafer is spun dry.

We have also developed a stronger etch that sequentially etches Al and then the sapphire substrate, as discussed in chapter 3. The recipe is shown in table A.2.

Table A.3: BCl_3Cl_2 etch recipe for Al films. SF_6 trench and undercut Si substrate.

	Step 1	Step 2	Step 3	Step 4	Step 5
BCl_3	20 SCCM	20 SCCM	20 SCCM	0 SCCM	0 SCCM
Cl_2	40 SCCM	40 SCCM	40 SCCM	0 SCCM	0 SCCM
SF_6	0 SCCM	0 SCCM	0 SCCM	50 SCCM	50 SCCM
Pre.	3.0 Pa	0.7 Pa	0.7 Pa	2.0 Pa	2.0 Pa
SRC FWD	300 W	300 W	300 W	700 W	700 W
BIAS FWD	0 W	0 W	70 W	0 W	20 W
Step Time	5 s	5 s	33 s	5 s	30 s

Silicon substrates

For silicon substrates we use the same BCl_3Cl_2 etch for Al, but we quench using SF_6 instead of CF_4. The CF_4 would sometimes leave polymer deposits after the etch, see chapter 6 for more details. The SF_6 quench etches several hundred nanometers into the Si and undercuts the Al (discussed in chapter 4). The typical recipe used is shown in table A.3.

A.1.4 Josephson junctions

We developed a Josephson junction (JJ) process based off of the Dolan bridge technique [127, 128]. The bridges fabricated using this process have enough structural integrity to span up to 4 μm long after an agressive Ar-ion mill used to remove the native AlO_x of the base metal. The ion mill is necessary to make galvanic contact to the existing base Al layer. Typically we fabricate the JJ last to make the process as reproducible as possible.

Resist stackup

We use a bilayer of resist to create the Dolan bridge (500nm MAA EL9 and 300nm PMMA 950K A4). We verify the thickness of the resist using a Woolam M2000DI variable angle spectroscopic ellipsometer. The PMMA thickness of 300nm is used to ensure structural integrity during the ion mill. The recipe is as follows, note the steps which are substrate dependent:

1. Spin MAA at 1500 rpm for 45 seconds.
2. Bake at 160° C for 10 minutes.
3. Spin PMMA at 2000 rpm for 40 seconds.
4. Bake at 160° C for 10 minutes.
5. (Sapphire only) deposit 10nm of Au at 0.5 Å/s in NRC 3117 three source thermal evaporator. The Au prevents charging of the sapphire substrate during the e-beam write.
6. Expose in e-beam writer, see section A.1.4 E-beam writer.
7. (Sapphire only) Strip Au in Gold etchant type TFA (KI 18-42 ppm solution) for 10 seconds, DI rinse.
8. Develop in 3:1 IPA,MIBK mixture for 45 seconds with gentle agitation.
9. Immerse in IPA for 10 seconds with gentle agitation.
10. Blow dry with N_2 at 10 PSI until dry.

We have tested the Au anti-charging layer on Si and found that it did not noticeably change the process. Therefore we omit this step when using Si substrates.

E-beam writer

We use a 100kV JEOL JBX-6300FS system for e-beam writing. By varying the dose we can expose either the MAA layer only, or expose both the PMMA and MAA. By doing this, we can directly define the undercut for the Dolan bridge and other features. We use the following doses:

- MAA dose - 350 μC/cm^2
- MAA (deep undercut) dose - 600 μC/cm^2
- PMMA+MAA dose - 1500 μC/cm^2
- Clearing dose - 2000 μC/cm^2

The strong clearing dose is used to minimize the resist residue on the base metal contact pads. This ensures that we make good contact using the ion mill discussed in section A.1.6. The MAA (deep undercut) dose is used for undercuts greater than 400 nm.

Deposition process

We use a Plassys e-beam evaporation system for the double-angle junction deposition.

1. Pump down overnight to let resist and chamber outgas. Final pressure is typically $< 10^{-7}$ mbar.
2. Ar-ion mill for 3 minutes 30 seconds to remove native oxide, see section A.1.6.
3. ebeam evaporate Ti for 3 minutes at 0.2nm/s for gettering.

4. First angle deposition at 62.2° from normal to the substrate. Deposit 65 nm of Al at 1 nm/s. (deposits ~30 nm film on substrate).
5. Oxidize at 5 mbar for 20-50 minutes (oxidation parameters can vary).
6. Second angle deposition normal to the substrate. Deposit 100 nm of Al at 1 nm/s.
7. Remove resist using NMP strip, as in section A.1.8.

Process reliability

We find this process to be extremely reliable, in the sense that since 2011 over 100 transmon qubits on both Si and sapphire substrates have given 100% yield. Other members of the group have used this process to yield devices with thousands of junctions [129]. We also find that our junctions do not age over the time span of days, as other groups have seen. This is likely due to the ion mill removing organic residue before deposition [130], and possibly the long liftoff in 80° solvent could anneal the junctions to some degree.

The process requires running test junctions before each qubit device, since oxidation parameters seem to drift substantially over time. Xmon devices are relatively simple, requiring only slight differences in junction critical currents between neighboring devices as is the case with the 9 Xmon sample. It is fairly straightforward to yield good devices when a test run is done beforehand. However more complicated devices such as the gmon qubits with the adjustable coupler requires a substantially larger critical current. Achieving both the ideal qubit junction resistance and coupler junction resistance on the same device can be a daunting task. For this reason we sometimes use a full 3 inch wafer during the e-beam write instead of running quarter wafers to allow for more variations in junction size across the wafer. Reliably fabricating the intended parameters for our junctions is still on going research.

A.1.5 Crossovers

We use crossovers for our qubit devices to reconnect the ground plane, which has been broken up due to the numerous qubit control lines.

Airbridge crossovers

In addition to using airbridges as ground plane connectors, we also use them as integral components of the qubits and adjustable couplers, see chapter 6 for further discussion.

We have experimented with fabricating airbridges at various steps in the device fabrication. When airbridge survival is critical, if they are integral to the qubit structure, we fabricate them after the junctions. Otherwise the junctions are fabricated last for better reproducibility. Processing after the junctions are fabricated typically leads to an increase in resistance, up to 20%, depending on the processing steps and temperatures.

The fabrication process has changed slightly from what was originally published [59]. We now use SPR 955 0.9 μm resist to fabricate the airbridges, producing a ~1.3 μm tall airbridge. A thinner resist is used so that the top of the bridges are completely covered in resist during the junction fabrication. The taller bridges were destroyed during the ion mill and lift off processes.

1. Begin with wafer that has CPW control lines already etched.

2. Expose and develop contact windows SPR 955 0.9 μm resist, section A.1.7.

3. Reflow resist by baking at 160 C for 3 min on hot plate.
4. Ar-ion mill in Plassys system to remove native AlO_x, as in section A.1.6.
5. Deposit 350 nm Al at 1 nm/s in Plassys system.
6. Define bridges, expose and develop SPR 955 0.9 μm resist, section A.1.7.
7. Downstream oxygen ash using Gasonics at 150 ° C for 1 min.
8. Etch using Transene Al Etchant Type A at 30° C until Al has cleared for 5 seconds.
9. Downstream oxygen ash using Gasonics at 150 ° C for 2 min (repeat once).
10. Liftoff using NMP, as in section A.1.8.

SiO_2 crossovers

The SiO_2 crossovers have been used only on some of the sapphire substrate based devices. First, the SiO_2 is extremely lossy and thus limits its use nearby or in the qubit structures. Second, the SiO_2 is ebeam deposited so it can have trapped charge which could induce a 2-DEG at the interface with the silicon substrate leading to loss. We thus only use airbridge crossovers on Si. The main benefit of the SiO_2 crossovers is the dielectric support makes them robust against further fabrication steps, such as sonication, which can sometimes remove or collapse airbridges if care is not taken. This is now a deprecated process, reproduced here for completeness. We first deposit the insulating SiO_2 using liftoff. We then complete the crossover also using liftoff of Al.

SiO_2:

1. Begin with wafer that has CPW control lines already etched.
2. Expose and develop SPR 955 0.9 μm resist, section A.1.7.
3. Deposit 200 nm SiO_2 at 0.3 nm/s in custom Temescal e-beam evaporation system.
4. Liftoff using NMP, as in section A.1.8.

Al:

1. Expose and develop SPR 955 0.9 μm resist, section A.1.7.
2. Ar-ion mill in Plassys system to remove native AlO_x, as in section A.1.6.
3. Deposit 200 nm Al in Plassys system.
4. Liftoff using NMP, as in section A.1.8.

The added capacitance of the SiO_2 crossovers requires altering the CPW structure to maintain a 50 Ω impedance.

A.1.6 Argon ion mill

In order to make contact with a previously deposited Al layer, we use an Ar-ion mill to remove its native oxide. We use a Kaufman source with beam energy of 400 eV, 21 mA current, and width of 3.2" at the substrate for a dose of $\sim$ 2.6 x $10^{15}\text{cm}^{-2}\text{s}^{-1}$. We mill for 3.5 min to ensure we break through the oxide [128]. We use the same milling parameters for the airbridge process, where we measured their critical current and found good contact was made, see Ref. [59].

A.1.7 Optical lithography

All developers are basic and thus attack Al, some more than others. The problem is much worse if dissimilar metals and semiconductors, such as Si, are placed in contact with both the developer and Al. Common TMAH developers such as AZ 300 MIF can completely etch through Al during the standard 60s develop time. AZ DEV 1:1 is Na ion based and the attack of Al is much less, but it is does not have as wide a processing window as TMAH developers. When using either developer care must be taken to not simultaneously expose nearby Si and Al, or the etch rate is drastically increased, presumably due to electrochemistry.

1. Cover wafer with HMDS for 30 s while on spinner.
2. Spin at 3000 RPM for 45 s.
3. Change spinners as HMDS vapor can contaminate resist.
4. Spin SPR 955 0.9 μm resist at 3000 RPM for 45 s.
5. Pre-exposure bake for 90° C for 90 s.
    (a) Expose in GCA 200 I-Line Wafer Stepper for 0.4 s if using AZ 300 MIF developer.
    (b) Expose in GCA 200 I-Line Wafer Stepper for 0.8 s if using AZ Dev 1:1 developer.
6. Post-exposure bake at 110° C for 90 s.
    (a) Develop in AZ 300 MIF developer with gentle agitation for 60 s.
    (b) If etching Al is a concern develop with AZ Dev 1:1 3 min.

7. Immerse in DI water with gentle agitation for 60 s.
8. N_2 blow dry at 20 psi until dry.

SPR 955 0.9 μm photoresist should be filtered with 0.45 μm filter for devices with many small features, such as narrow CPW.

A.1.8 Resist stripping procedures

We use several resist stripping procedures. We have found the NMP to remove resist residue better than the acetone based strip [114]. Note, that one must keep device side of wafer 'wet' when transferring between beakers.

NMP

1. Soak in Shipley 1165 (N-methyl-2-pyrrolidine) at 80° for 10 min (up to 3 hours for lift off).
2. Sonicate in second beaker of 1165 at 80° for 5 min.
3. Sonicate in IPA for 3 min
4. Spin dry

Acetone

1. Sonicate in Acetone for 3 min.
2. Sonicate in second beaker of Acetone for 3 min.

3. Sonicate in IPA for 3 min.

4. Spin dry

A.2 Fabrication equipment

The following is a list of equipment used during the fabrication and characterization of superconducting devices.

A.2.1 Deposition

Multiple deposition systems are used during the fabrication of Xmon and gmon superconducting qubits.

Veeco 930 MBE system

We use this system to clean the substrate and deposit the base wiring metalization layer. It consists of three chambers: load-lock, staging chamber, and main deposition chamber. This system supports up to three inch wafers.

The load-lock contains quartz heat lamps for initial wafer and chamber degassing up to $200\,^{\circ}\mathrm{C}$. It can reach a base pressure of ($P_{\text{base}} = 1 \times 10^{-8}$ Torr) using a dedicated turbo pump.

The staging chamber is located between the load lock and main chamber. It can reach a base pressure of ($P_{\text{base}} < 5 \times 10^{-11}$ Torr) using a dedicated ion pump and also includes a titanium sublimation pump (TSP) used during the system bakeout procedure. This chamber also contains a substrate heater capable of reaching $800\,^{\circ}\mathrm{C}$.

The main deposition chamber is used for in-situ substrate cleaning and thin film deposition. It can reach a base pressure of ($P_{\text{base}} < 5 \times 10^{-11}$ Torr) using a variety of UHV pumps consisting of: Varian 2000HT turbo pump, Oxford cryo pump, and an ion pump with TSP and surruonding cryoshield. Two additional cryo-shields can flow liquid nitrogen to isolate sources and reduce outgassing during depositions.

For substrate cleaning, the main chamber has a substrate heater capable of $1200\,^{\circ}\mathrm{C}$ in an oxygen rich environment ($P_{\text{O2}} \leq 1 \times 10^{-6}$ Torr). The chamber also has a Veeco UNI-Bulb RF plasma source for activating oxygen or nitrogen source gases in order to make them more reactive.

For deposition, the main chamber is equipped with a custom Thermionics three pocket electron beam evaporator with a 6 kW power supply for thin film deposition. It has an Infincon quartz crystal microbalance for determining the deposited film thickness.

For characterization, the main chamber is equipped with a SRS RGA100 quadrupole mass spectrometer residual gas analyzer and a 15 kV STAIB reflection high-energy electron diffraction system (RHEED) for *in-situ* surface characterization.

Plassys E-beam evaporator

The double angle junction deposition is performed in this tool as well as both types of crossovers. This two chamber system, consisting of a load lock and deposition chamber, is a general purpose E-beam evaporator with up to 6 sources. This system can reach a base pressure of ($P_{\text{base}} < 4 \times 10^{-8}$ Torr). A Kaufmann ion mill is used for cleaning and oxide removal. The sample stage has multi-axis rotation, including continuous rotation about the substrate normal direction.

Custom E-beam evaporator - dielectrics

This is a custom e-beam evaporator for general use in the UCSB cleanroom. It is used solely for the deposition of dielectrics, including the SiO_2 for crossovers. It has a Temescal 10 kV power supply with four source material pockets. It can reach a base pressure of ($P_{\text{base}} < 5 \times 10^{-7}$ Torr). Substrates up to four inches can be used with an optional heater.

Temescal E-beam evaporator - metals

The Temescal e-beam evaporator is a general use tool in the UCSB cleanroom. Its main purpose is for metal deposition and Ge. We mainly use this system to deposit the Ti/Au alignment marks for junction e-beam write. It has a custom built load-lock which allows for fast turnarounds and keeps the source metals under high vacuum, and can reach a base pressure of ($P_{\text{base}} < 5 \times 10^{-7}$ Torr). The power supply is 10 kV and has 8 source material pockets. Samples up to four inches can be used.

A.2.2 Lithography

We define all of our structures using the GCA 200 stepper except the Josephson junctions, which uses the JEOL E-beam writer.

GCA 200 wafer stepper

The GCA wafer stepper is an i-line (365 nm) step and repeat exposure tool. The system accepts piece parts (down to our 6 mm × 6 mm dies) up to 6 diameter wafers using manual wafer loading. It can expose up to a maximum square die size of 14.8 mm × 14.8 mm. The system has an Olympus 2145 (N.A. = 0.45) lens that reduces the mask image by 5x. It is capable of resolving features as small as 0.5 μm, however we avoid features smaller than 1.5 μm in design. The system performs a 3-point wafer leveling, and has the capability of a manual and optional local alignment. We design for 0.5 μm alignment error and typically achieve better than this (~0.5 μm) when using 3 inch wafers. The i-line intensity is 420 mW/cm^2 using a 1000 W Hg arc lamp.

JEOL JBX-6300FS E-beam writer

The JEOL e-beam writer uses a thermal field emission electron gun with a ZnO/W emitter. It has a minimum spot-size at the substrate of 2 nm, and operates at 100 kV only. It has a unique two lens/deflector scanning system, with either a 8 nm minimum linewidth for a 62.5 μm × 62.5 μm scan field or a 25 nm minimum linewidth with a 500 μm × 500 μm scan field. We use the latter as our feature sizes are greater than 100 nm. The maxiumum deflector scan speed is 25 MHz. It has a 150 mm × 150 mm writable area with stage postition control of 0.6 nm accuracy and a 10 mm/s maximum stage speed.

A.2.3 Etching

Panasonic E626I Inductively coupled plasma reactive ion etcher

The ICP etcher is a general use tool located in the UCSB cleanroom. We use this tool for all of the dry etching described in this thesis. It has 1000 W ICP power and 500 W RF substrate power, and room temperature to $80\,^{\circ}\mathrm{C}$ operation. The wafers are cooled using a flow of He accross the back-side and clamped using an electrostatic chuck to maintain stable temperatures during etching. The system includes many source gases: Cl_2, BCl_3, CF_4, CHF_3, SF_6, Ar, N_2, and O_2, although some are routed through a switch panel and cannot be used together. The system accepts 6 inch wafers, so smaller pieces need to be mounted onto a carrier wafer using diffusion pump oil. Multiple wafers can be run automatically using a cassette loading system.

Xetch XeF_2 etcher

The XeF_2 etcher uses a purely chemical etch process used for isotropic etching of Si, Ge and some refractory metal such as Mo and Ta. The etch process consists of multiple cycles of pumping the chamber to ~0.3 Torr, then filling the chamber with XeF_2 to the desired pressure, with the option of also including N_2 gas. Two XeF_2 outgassing chambers are included to decrease the time between fill cycles. The system includes a transparent lid with external inspection microscope for realtime monitoring of the etch process.

SPTS Primaxx uEtch vapor HF etcher

This tool is used to etch SiO_2 while having a large selectivity with Al and Si. It uses a cyclic pump and purge protocol similar to atomic-layer deposition (ALD). To begin the entire chamber and process gases are preheated to $45\,^{\circ}\mathrm{C}$ under vacuum. Then process gases anhydrous HF vapor, ethanol vapor (EtOH), and N_2 are delivered into the chamber and after some time pumped out of the chamber before repeating the cycle. The ethanol vapor (EtOH) is a catalyst used to ionize the HF and activate the etching. Process gas flows, temperature, and pressure must be controlled so that liquid water does not condense onto the etching surface, which can lead uncontrolled etch rates and greatly reduced selectivity to Al.

Gasonics Aura 2000 downstream oxygen asher

Oxygen ashers are used to remove resist and other organic contaminants. This system forms a plasma by exciting the oxygen gas with a RF microwave source. The heated sample is located in the etching chamber downstream from the plasma chamber. This separation eliminates possible damage from ion bombardment. This system is capable of handling 8 inch wafers and heating them from $100\,^{\circ}\mathrm{C}$ to $350\,^{\circ}\mathrm{C}$ using heat lamps.

A.2.4 Characterization

FEI Sirion scanning electron microscope

This is the main SEM in the UCSB cleanroom. It has a thermal field emission source and uses an accelerating voltage between 200 V to 30 kV with resolution down to 1.5 nm. The ultra high resolution mode allows magnification between 2500x to 1,200,000x. The sample stage can tilt between tilt = $-15\,^{\circ}$ to $+75\,^{\circ}$. The system is also equipped with an EDS for elemental analysis of materials.

Asylum MFP-3D atomic force microscope

The Asylum AFM is located in a separate microscopy lab belonging to the Materials Research Laboratory (MRL) at UCSB. It is housed in a AEK 2002 acoustic isolation enclosure. The AFM is mounted on top of a Herzan TS-150 vibration isolation table, which actively minimizes vibrations between 0.7 Hz to 1 kHz up to 40 dB. The AFM is capable of 90 μm of travel in (X,Y) and 15 μm in Z. It is controlled using IGOR pro software and is capable of many operational modes, although we typically use AC mode, which controls using feedback from the amplitude of deflection.

Panalytical MRD PRO Materials Research Diffractometer

We use this tool to measure the crystallinity of our thin films using X-ray diffraction (XRD). The X-ray source is a 2.2 kW Philips ceramic sealed tube. The Cu Kα radiation (1.5405 Å) is filtered using a high flux hybrid monochromator or 4-crystal Ge(220)/(440). The signal is detected using a sealed proportional counter. The system has a 5 arcsec resolution. It is controlled using X'PERT data collector.

Appendix B

Calibrations

B.1 Qubit Calibrations

There are many calibrations required to run Xmon and gmon qubit experiments. The standard Xmon qubit calibrations have been previously described in Ref. [131, 128]. The following are additional calibrations required when using an adjustable coupling scheme. In particular, it is desirable to separate the control parameters of coupling strength and qubit frequency.

Given the nature of the gmon coupling architecture, it is inherent that tuning the coupling strength will also change the inductance of the qubit and therefore its frequency. Since the gmon qubits are frequency tunable, we can apply a compensation control signal. Models describing this effect are not accurate to the sub-MHz precision that is required for the analog simulations described in this thesis. More accurate models are needed as the number of qubits and couplers increase to minimize the time required for calibrations.

We begin the calibrations by first characterizing the cross talk between each flux bias line to each quantum element. Once this is corrected for, we can determine the energy splitting as well as the qubit frequencies as a function of coupler flux bias. After applying these calibrations, we then minimize the detuning between qubits as a function of coupler flux bias. Now that we have orthogonalized the coupling strength and the qubit frequencies, we carefully calibrate the coupling strength as a function of coupler flux bias.

B.1.1 Z crosstalk

Crosstalk to the qubits is measured to be several percent or smaller. However crosstalk to couplers can be much larger, up to 30%. This is due to the large inductive loops of the couplers, as seen in Fig. 6.13. The crosstalk matrix M_Φ is shown below, defined as $\Phi_{\text{actual}} = (1 + dM_\Phi)\Phi_{\text{ideal}}$, with Φ the flux threaded through each qubit's superconducting quantum interference device (SQUID) loop or coupler loop.

$$dM_\Phi \approx 10^{-3} \left(\begin{array}{cccccc|c} C_{12} & C_{23} & C_{31} & Q_1 & Q_2 & Q_3 & \\ 0 & 94 & 92 & -34 & -26 & 149 & C_{12} \\ 123 & 0 & 32 & 40 & -113 & -140 & C_{23} \\ 107 & 111 & 0 & -336 & 111 & -6 & C_{31} \\ 47 & 4 & -48 & 0 & 53 & -45 & Q_1 \\ -19 & 24 & 18 & 10 & 0 & 35 & Q_2 \\ 18 & 21 & -17 & -8 & 40 & 0 & Q_3 \end{array}\right)$$

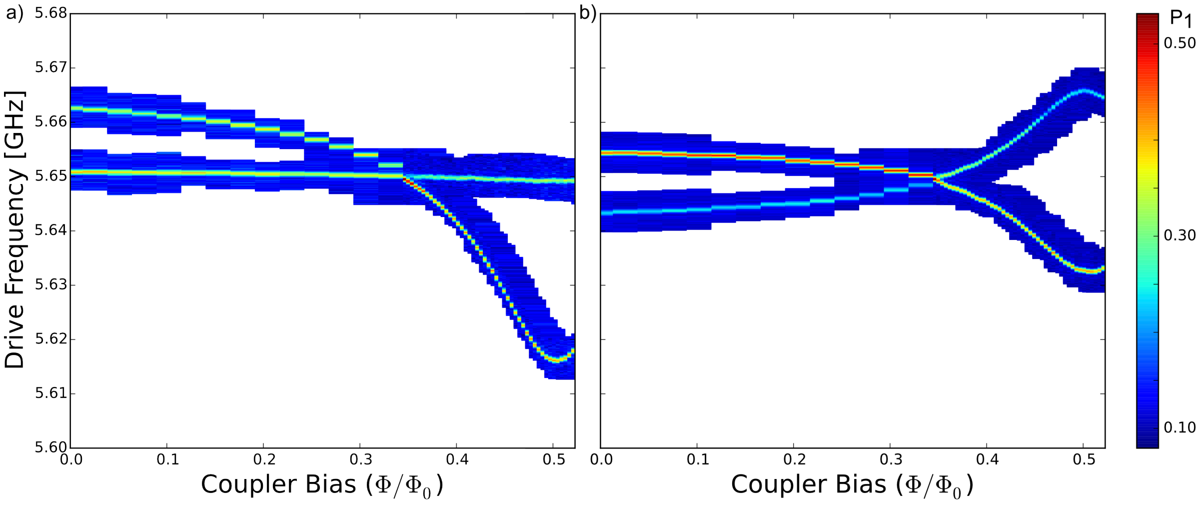

Figure B.1: (a) Spectroscopy as a function of coupler flux bias. Note how the qubit frequencies are shifting as the coupler flux bias is varied. b) Spectroscopy of the same qubits and coupler as (a). Here we apply a compensation pulse to each qubit, to keep qubit frequencies constant as a function of coupler bias. The separation between the spectroscopy peaks at a given coupler flux bias is equal to $2g/2\pi$

B.1.2 Two qubit spectroscopy versus coupling strength

The frequency of the qubits shift when the coupling strength is varied, as seen in Fig. B.1(a), where we plot two qubit spectroscopy as a function of coupler flux bias. From this data we learn the coupling strength g as well as the shift in the resonance frequency of the qubits. In Fig. B.1(b), we apply a compensation pulse to both qubits such that the mean frequency of the spectroscopy peaks remains constant. This ensures that varying the coupling strength between one pair of qubits does not effect the detuning between another pair of qubits. We use an adaptive algorithm to minimize the measurement time. This scan is performed for each pair of qubits while the third qubit is far detuned.

B.1.3 Minimize qubit detuning vs coupler bias

The qubits need to remain on resonance when varying the coupling strength. We begin with the coupling between a pair of qubits turned off and the two qubits on resonance, while the third qubit is far detuned. We excite one qubit, Q_A, and turn on the coupling for the time, $t_{swap} = 2\pi/4g$, to completely swap the excitation to the other qubit, Q_B. Finally we measure Q_A. We repeat at the same coupling strength while varying the frequency of Q_B. The minimum probability remaining in Q_A corresponds to zero detuning between the qubits. The inset of Fig. B.2 shows the pulse sequence. The frequency span and t_{swap} are decided adaptively based on the coupling strength. We can then apply a correction pulse to the qubits to zero the detuning at all coupling strengths.

B.1.4 Coupling strength vs bias

The simulation of the collision cross section requires high accuracy in the coupling strength. Unfortunately, the two qubit spectrocopy method was not accurate or precise enough. Here we improve upon the previous method by measuring the swap of a single excitation between a pair of qubits as a function of time, as seen in Fig. B.3(a). We measure up to 15 periods of swapping to determine the coupling strength, $g/2\pi$, as a function of coupler flux bias as seen in Fig. B.3(b). This method generates a very smooth calibration curve, showing the high precision of such a measurement.

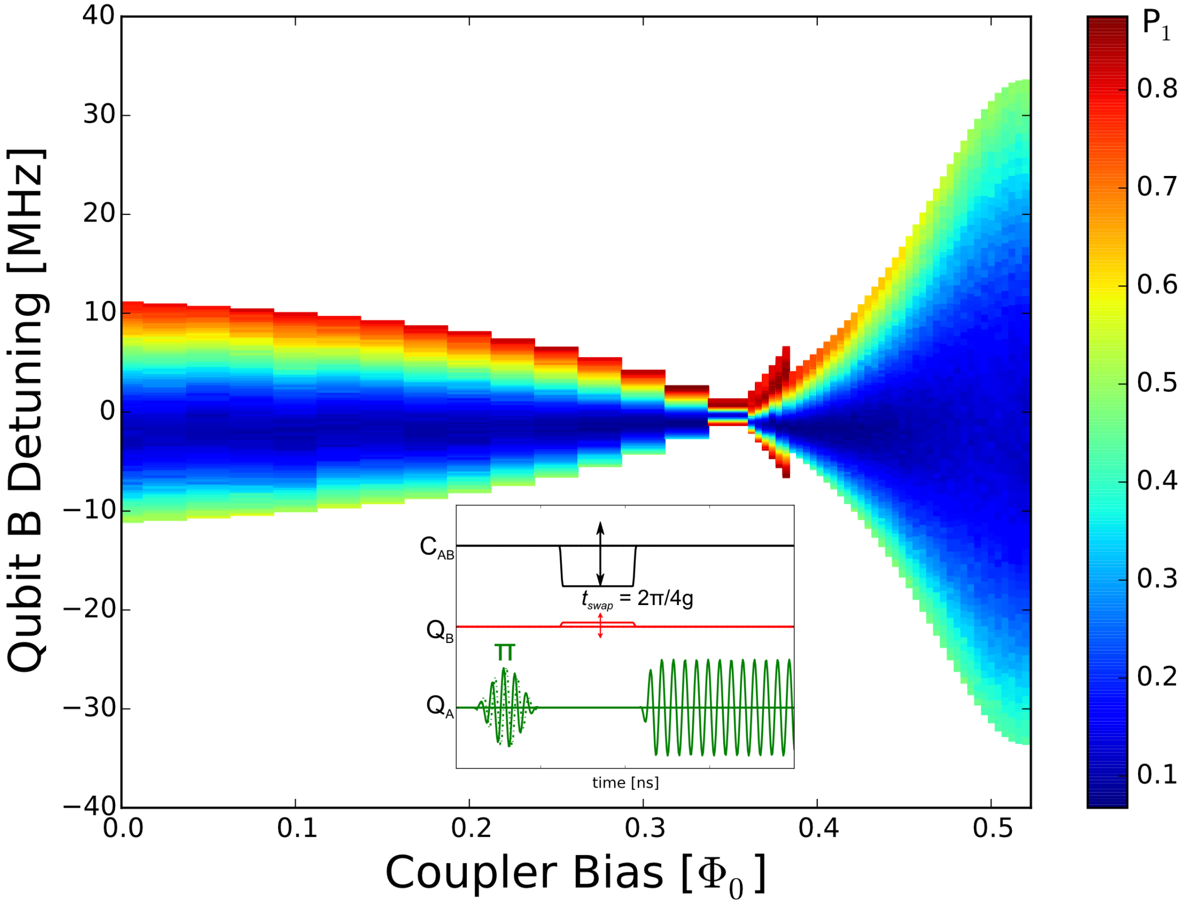

Figure B.2: (Inset) Pulse sequence. We excite qubit Q_A and turn on the coupling strength g for the time, $t_{swap} = 2\pi/4g$, to completely swap the excitation to the other qubit Q_B. Finally we measure Q_A. We minimize the detuning between qubits, by minimizing the population in Q_A. We do this for all coupling strengths. The third qubit is far detuned. We use this information to apply compensation pulses to both qubits.

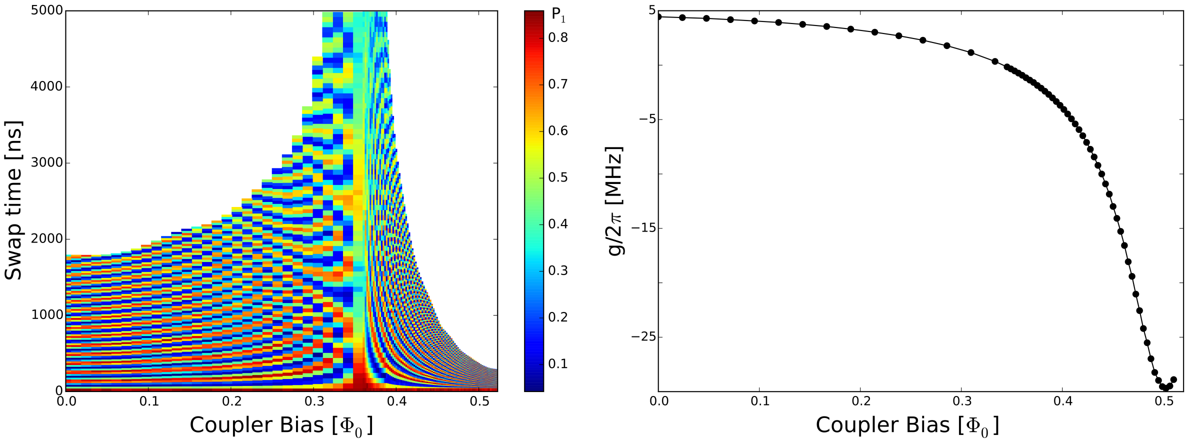

Figure B.3: (a) The population of the initially excited qubit as a function of swap time and coupler flux bias. (b) Coupling strength $g/2\pi$ as a function of coupler flux bias, determined using data of (a).

Appendix C

Measurement setup

The shielding and wiring configuration is similar between the wet dilution refrigerators (DR) used for qubit measurements and the adiabatic demagnetization refrigerators (ADR) used with resonators. They both consist of radiation and magnetic shielding as well as filtering of the microwave lines.

C.1 Adiabatic demagnetization refrigerator

The resonator measurements described in this thesis have all been performed in an High Precision Devices, Inc ADR model 103. The ADR provides very fast turn around times. It takes roughly 24 hours to cool to base temperature and about 12 hours to warm back up. We typically measure two samples boxes, as seen in Fig. C.1(a), up to four are possible, each with up to 16 resonators. The measurement time per resonator is dominated by the measurements at low drive power, near a single photon circulating in the resonator, where the SNR is small. For resonators with $Q_i \sim Q_c$ at low drive power, we can measure a full power sweep of a resonator in about 15 min, with a threshold of no more than 5% error in the internal quality factor at the lowest powers. We typically scan about 15 drive powers spaced 5 dB, which allows us to measure each sample in roughly 8 hours. The ADR slowly warms up to 100 mK in about 40 hours, easily allowing both samples to be measured. The full measurement cycle time is typically 2-3 days.

The sample box is made from machined Al alloy 6061, as seen in Fig. C.1(a). The center conductor of the nonmagnetic panel mount SMA connector is soldered to a Rogers board (RO43450B) which is 37 mil wide by 13 mil thick with 1.4 mil thick of Cu cladding. We wirebond from the Cu trace to the bond pads on the device. This is the same box used with qubits. The box shown has 12 SMA ports but only two are used for the resonator measurements. After the sample is wired and mounted to the 50 mK cold plate, we enclose the experiment in the first layer of shielding, a gold plated annealed copper light-tight shield. We coat the inside bottom of the can with "IR-Black" coating to absorb stray IR radiation [64], as seen in Fig. C.1(b). In Fig. C.1(c) We see the Cryoperm magnetic shield installed in the modified ADR setup. The wiring diagram in Fig. C.1(d) shows how all of the shielding and wiring are assembled. Additionally the ADR has the typical radiation shielding at the 4 K and 50 K stages.

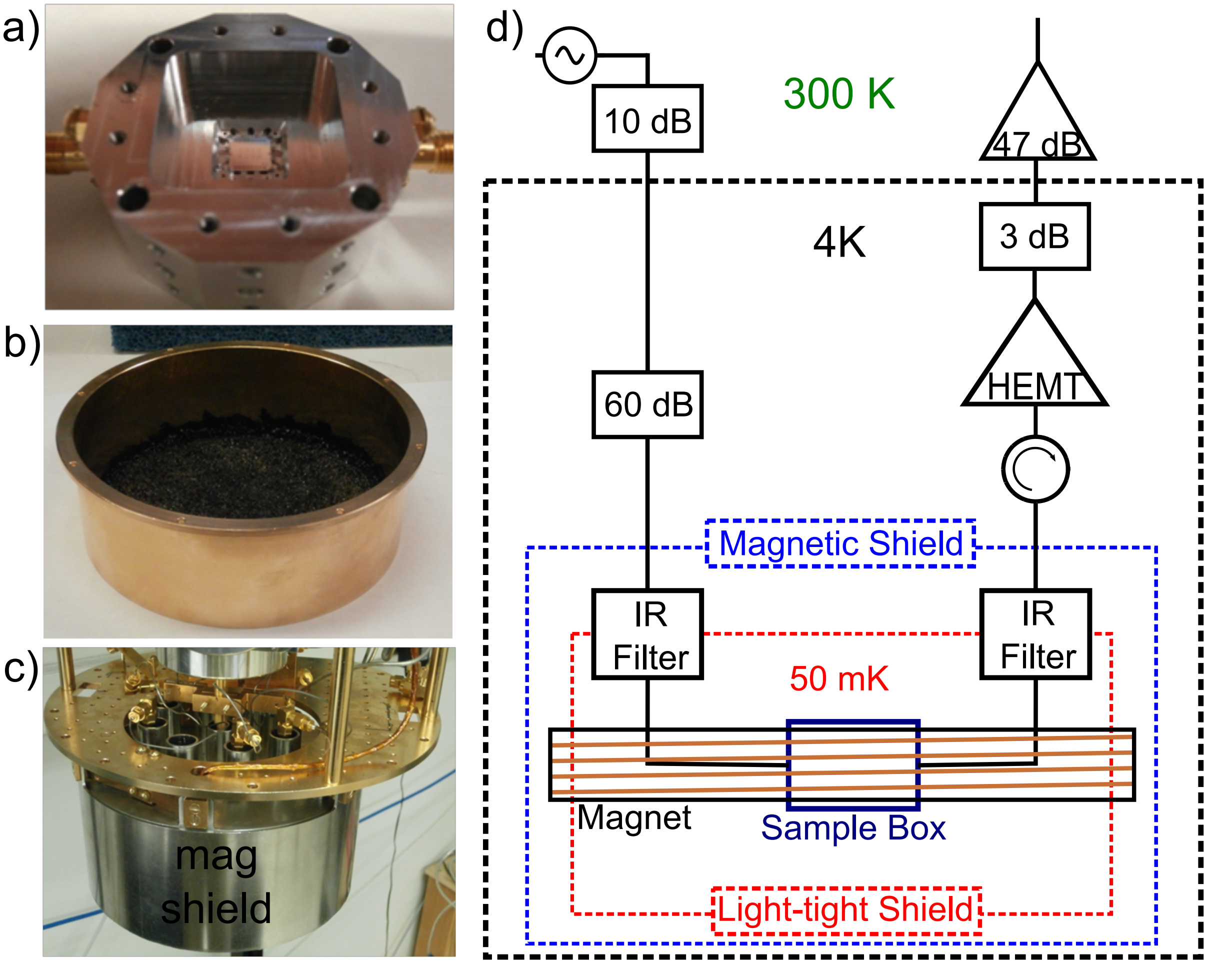

Figure C.1: a) Machined aluminum sample mount with nonmagnetic SMA connectors, screws and washers. b) Gold plated annealed copper can, component of light-tight shield. We coat the inside of the can with “IR-Black” coating to absorb stray IR radiation. [64]. c) Cryoperm magnetic shield installed in ADR setup. d) Typical wiring diagram for resonator measurement. We include a magnet for field cooling measurements described in chapter 4. The IR filters are custom made [64]. The HEMT is a Low Noise Factory LNC4_8A and the circulator is a QuinStar CTH1392KS.

C.2 Dilution refrigerator

We use a custom-built wet dilution refrigerator for qubit experiments. The base temperature of the mix plate is ~20 mK measured using a calibrated ruthenium oxide thermometer. The cooldown process takes 2-3 days to reach base temperature and warming up typically takes 36 hours. The mix plate has enough space for several multi-qubit devices. This is mainly limited by the magnetic shielding surrounding each individual qubit sample box, which is identical to that shown in Fig. C.1(a). A large copper can similar to Fig. C.1(b) encloses the entire mix plate. There is an additional radiation shield at the still, followed by the 4 K, and finally 77 K. Additional filtering and custom electronics control and readout the qubits, described in Fig. C.2.

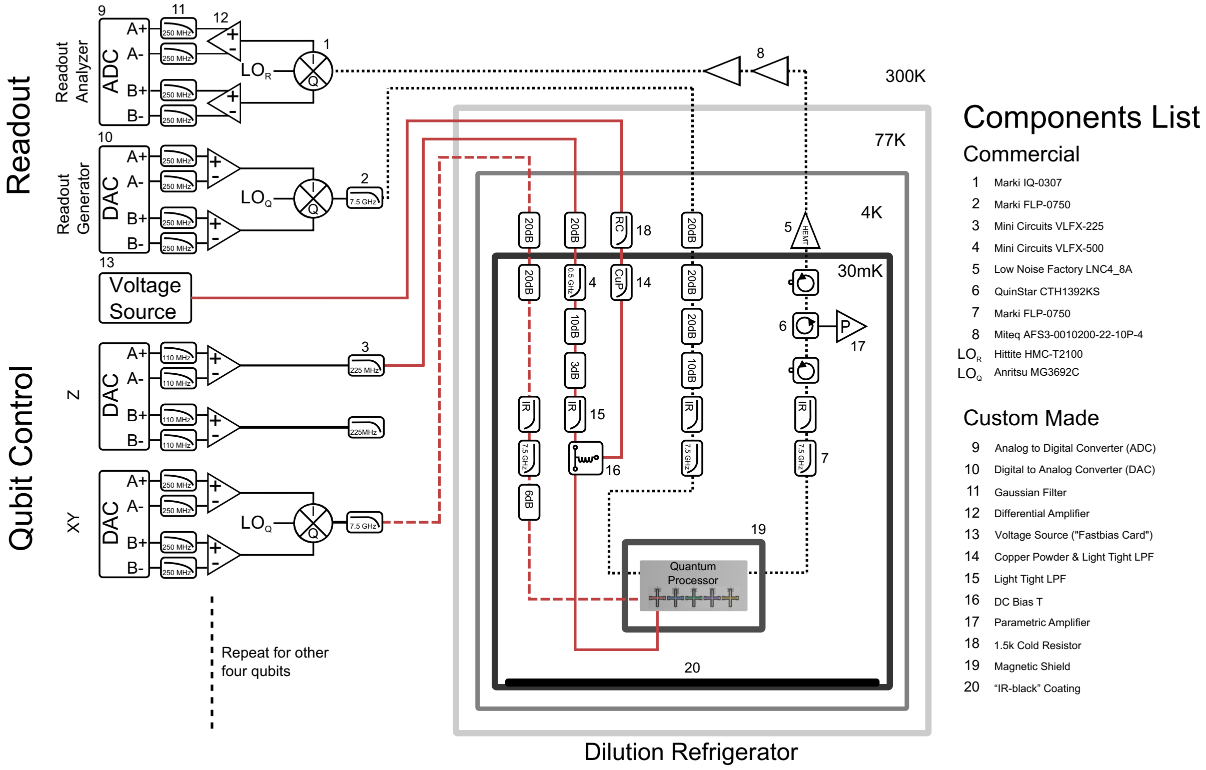

Figure C.2: This figure reproduced from Ref. [27]. Diagram which details the filtering, control wiring, and control electronics used in a typical qubit experiment. One digital to analog converter channel (DAC) controls rotations around X, Y, and Z axis for each qubit. Additionally, we can apply a static frequency offset to a qubit's Z control line through a DC bias tee. All of the control and readout wires go through multiple stages of filtering and attenuation to minimize stray signals and noise. We multiplex the frequencies of the readout resonators so that all can be readout simultaneously using a single DAC and the same measurement line. The transmitted signal is amplified by a wide bandwidth impedance matched parametric amplifier [118], followed by a high electron mobility transistor (HEMT) amplifier, and finally a room temperature low noise microwave amplifier before demodulation and state discrimination by the analog to digital converter (ADC).

Bibliography

[1] G. E. Moore. Cramming more components onto integrated circuits. *Electronics*, 38(8):114–117, 1965.

[2] Top500. Top500 supercomputer list, 2015. [Online; accessed 3-December-2015].

[3] Richard Cleve, Artur Ekert, Chiara Macchiavello, and Michele Mosca. Quantum algorithms revisited. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 454, pages 339–354. The Royal Society, 1998.

[4] Richard P Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6):467–488, 1982.

[5] Susumu Yamada, Toshiyuki Imamura, and Masahiko Machida. 16.14 tflops eigenvalue solver on the earth simulator: Exact diagonalization for ultra largescale hamiltonian matrix. In *High-Performance Computing*, volume 4759 of *Lecture Notes in Computer Science*, pages 402–413. Springer Berlin Heidelberg, 2008.

[6] Andreas M Läuchli. Numerical simulations of frustrated systems. In *Introduction to Frustrated Magnetism*, pages 481–511. Springer, 2011.

[7] Sergei Isakov, William Sawyer, Gilles Fourestey, Adrian Tineo, Neil Stringfellow, and Mattias Troyer. High-performance exact diagonalization techniques. *CUG2012 Final Proceedings*, 2012.

[8] Gilbert Laporte. The vehicle routing problem: An overview of exact and approximate algorithms. *European Journal of Operational Research*, 59(3):345–358, 1992.

[9] Eugene L Lawler. The traveling salesman problem: a guided tour of combinatorial optimization. 1986.

[10] Gilbert Laporte. The traveling salesman problem: An overview of exact and approximate algorithms. *European Journal of Operational Research*, 59(2):231–247, 1992.

[11] Rajesh Matai, Murari Lal Mittal, and Surya Singh. *Traveling salesman problem: An overview of applications, formulations, and solution approaches*. INTECH Open Access Publisher, 2010.

[12] Maria Schuld, Ilya Sinayskiy, and Francesco Petruccione. An introduction to quantum machine learning. *Contemporary Physics*, 56(2):172–185, 2015.

[13] Lov K Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219. ACM, 1996.

[14] Peter W Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *35th Annual Symposium on Foundations of Computer Science, 1994 Proceedings.*, pages 124–134. IEEE, 1994.

[15] David P. DiVincenzo. The physical implementation of quantum computation. *Fortschritte der Physik*, 48(9-11):771–783, 2000.

[16] Austin G Fowler, Matteo Mariantoni, John M Martinis, and Andrew N Cleland. Surface codes: Towards practical large-scale quantum computation. *Physical Review A*, 86(3):032324, 2012.

[17] Paul I Bunyk, Emile M Hoskinson, Mark W Johnson, Elena Tolkacheva, Fabio Altomare, Andrew J Berkley, Roy Harris, Jeremy P Hilton, Trevor Lanting, Anthony J Przybysz, and J. Whittaker. Architectural considerations in the design of a superconducting quantum annealing processor. *IEEE Transactions on Applied Superconductivity*, 24(4):1–10, 2014.

[18] M.W. Johnson, M.H.S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A.J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. S. Truncik, S. Uchaikin, J. Wang, B. Wilson, and G. Rose. Quantum annealing with manufactured spins. *Nature*, 473(7346):194–198, 2011.

[19] Michael R. Geller, John M. Martinis, Andrew T. Sornborger, Phillip C. Stancil, Emily J. Pritchett, Hao You, and Andrei Galiautdinov. Universal quantum simulation with prethreshold superconducting qubits: Single-excitation subspace method. *Phys. Rev. A*, 91:062309, Jun 2015.

[20] Jan Benhelm, Gerhard Kirchmair, Christian F Roos, and Rainer Blatt. Towards fault-tolerant quantum computing with trapped ions. *Nature Physics*, 4(6):463–466, 2008.

[21] TP Harty, DTC Allcock, CJ Ballance, L Guidoni, HA Janacek, NM Linke, DN Stacey, and DM Lucas. High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit. *Physical Review Letters*, 113(22):220501, 2014.

[22] M Veldhorst, JCC Hwang, CH Yang, AW Leenstra, B De Ronde, JP Dehollain, JT Muhonen, FE Hudson, Kohei M Itoh, A Morello, and AS Dzurak. An addressable quantum dot qubit with fault-tolerant control-fidelity. *Nature nanotechnology*, 9(12):981–985, 2014.

[23] Mete Atatüre, Jan Dreiser, Antonio Badolato, Alexander Högele, Khaled Karrai, and Atac Imamoglu. Quantum-dot spin-state preparation with near-unity fidelity. *Science*, 312(5773):551–553, 2006.

[24] Neil A Gershenfeld and Isaac L Chuang. Bulk spin-resonance quantum computation. *Science*, 275(5298):350–356, 1997.

[25] CA Ryan, Martin Laforest, and Raymond Laflamme. Randomized benchmarking of single-and multi-qubit control in liquid-state NMR quantum information processing. *New Journal of Physics*, 11(1):013034, 2009.

[26] L DiCarlo, JM Chow, JM Gambetta, Lev S Bishop, BR Johnson, DI Schuster, J Majer, A Blais, L Frunzio, SM Girvin, and RJ Schoelkopf. Demonstration of two-qubit algorithms with a superconducting quantum processor. *Nature*, 460(7252):240–244, 2009.

[27] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O/'Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and John M. Martinis. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497):500–503, 2014.

[28] J. Kelly, R. Barends, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Yu Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O/'Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and John M. Martinis. State preservation by repetitive error detection in a superconducting quantum circuit. *Nature*, 519(7541):66–69, 2015.

[29] Jeremy L O'Brien, Geoffrey J Pryde, Andrew G White, Timothy C Ralph, and David Branning. Demonstration of an all-optical quantum controlled-not gate. *Nature*, 426(6964):264–267, 2003.

[30] Markus Greiner, Olaf Mandel, Tilman Esslinger, Theodor W Hänsch, and Immanuel Bloch. Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. *Nature*, 415(6867):39–44, 2002.

[31] Charles P Slichter. *Principles of Magnetic Resonance*, volume 1. Springer Science & Business Media, 2013.

[32] Markus Ansmann. *Benchmarking the Superconducting Josephson Phase Qubit - The Violation of Bell's Inequality*. PhD thesis, University of California Santa Barbara, 2009.

[33] Daniel Sank, R. Barends, Radoslaw C. Bialczak, Yu Chen, J. Kelly, M. Lenander, E. Lucero, Matteo Mariantoni, A. Megrant, M. Neeley, P. J. J. O'Malley, A. Vainsencher, H. Wang, J. Wenner, T. C. White, T. Yamamoto, Yi Yin, A. N. Cleland, and John M. Martinis. Flux noise probed with real time qubit tomography in a Josephson phase qubit. *Physical Review Letters*, 109(6):067001, 2012.

[34] John Clarke and Frank K Wilhelm. Superconducting quantum bits. *Nature*, 453(7198):1031–1042, 2008.

[35] Hui Wang, Chuntai Shi, Jun Hu, Sungho Han, C Yu Clare, and RQ Wu. Candidate source of flux noise in squids: adsorbed oxygen molecules. *Physical Review Letters*, 115(7):077002, 2015.

[36] P. J. J. O'Malley, J. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, I.-C. Hoi, E. Jeffrey, A. Megrant, J. Mutus, C. Neill, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, A. N. Korotkov, A. N. Cleland, and John M. Martinis. Qubit metrology of ultralow phase noise using randomized benchmarking. *Physical Review Applied*, 3(4):044009, 2015.

[37] Peter K Day, Henry G LeDuc, Benjamin A Mazin, Anastasios Vayonakis, and Jonas Zmuidzinas. A broadband superconducting detector suitable for use in large arrays. *Nature*, 425(6960):817–821, 2003.

[38] Max Hofheinz, H Wang, Markus Ansmann, Radoslaw C Bialczak, Erik Lucero, Matthew Neeley, AD O'Connell, Daniel Sank, J Wenner, John M Martinis, and AN Cleland. Synthesizing arbitrary quantum states in a superconducting resonator. *Nature*, 459(7246):546–549, 2009.

[39] H. Wang, M. Hofheinz, J. Wenner, M. Ansmann, R. C. Bialczak, M. Lenander, Erik Lucero, M. Neeley, A. D. OConnell, D. Sank, M. Weides, A. N. Cleland, and John M. Martinis. Improving the coherence time of superconducting coplanar resonators. *Applied Physics Letters*, 95(23):233508, 2009.

[40] Matteo Mariantoni, H. Wang, T. Yamamoto, M. Neeley, Radoslaw C. Bialczak, Y. Chen, M. Lenander, Erik Lucero, A. D. O'Connell, D. Sank, M. Weides, J. Wenner, Y. Yin, J. Zhao, A. N. Korotkov, A. N. Cleland, and John M. Martinis. Implementing the quantum Von Neumann architecture with superconducting circuits. *Science*, 334(6052):61–65, 2011.

[41] Aaron D. OConnell, M. Ansmann, R. C. Bialczak, M. Hofheinz, N. Katz, Erik 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(11):112903, 2008.

[42] A. Megrant, C. Neill, R. Barends, B. Chiaro, Yu Chen, L. Feigl, J. Kelly, Erik Lucero, Matteo Mariantoni, P. J. J. OMalley, D. Sank, A. Vainsencher, J. Wenner, T. C. White,

Y. Yin, J. Zhao, C. J. Palmstrm, John M. Martinis, and A. N. Cleland. Planar superconducting resonators with internal quality factors above one million. *Applied Physics Letters*, 100(11):113510, 2012.

[43] JP Turneaure and Ira Weissman. Microwave surface resistance of superconducting niobium. *Journal of Applied Physics*, 39(9):4417–4427, 1968.

[44] Matthew Reagor, Hanhee Paik, Gianluigi Catelani, Luyan Sun, Christopher Axline, Eric Holland, Ioan M. Pop, Nicholas A. Masluk, Teresa Brecht, Luigi Frunzio, Michel H. Devoret, Leonid Glazman, and Robert J. Schoelkopf. Reaching 10 ms single photon lifetimes for superconducting aluminum cavities. *Applied Physics Letters*, 102(19):192604, 2013.

[45] Daniel L. Creedon, Yarema Reshitnyk, Warrick Farr, John M. Martinis, Timothy L. Duty, and Michael E. Tobar. High Q-factor sapphire whispering gallery mode microwave resonator at single photon energies and millikelvin temperatures. *Applied Physics Letters*, 98(22), 2011.

[46] Benjamin A Mazin. *Microwave kinetic inductance detectors*. PhD thesis, California Institute of Technology, 2004.

[47] Jens Koch, M Yu Terri, Jay Gambetta, Andrew A Houck, DI Schuster, J Majer, Alexandre Blais, Michel H Devoret, Steven M Girvin, and Robert J Schoelkopf. Charge-insensitive qubit design derived from the cooper pair box. *Physical Review A*, 76(4):042319, 2007.

[48] Hanhee Paik, D. I. Schuster, Lev S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. R. Johnson, M. J. Reagor, L. Frunzio, L. I. Glazman, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf. Observation of high coherence in Josephson junction qubits measured in a three-dimensional circuit qed architecture. *Physical Review Letters*, 107(24):240501, 2011.

[49] Brian David Josephson. Possible new effects in superconductive tunnelling. *Physics Letters*, 1(7):251–253, 1962.

[50] Philip W Anderson and John M Rowell. Probable observation of the Josephson superconducting tunneling effect. *Physical Review Letters*, 10(6):230, 1963.

[51] John M Martinis and Kevin Osborne. Superconducting qubits and the physics of Josephson junctions. *arXiv preprint cond-mat/0402415*, 2004.

[52] Vladimir Manucharyan. *Superinductance*. PhD thesis, Yale University, 2012.

[53] R. Barends, J. Kelly, A. Megrant, D. Sank, E. Jeffrey, Y. Chen, Y. Yin, B. Chiaro, J. Mutus, C. Neill, P. O'Malley, P. Roushan, J. Wenner, T. C. White, A. N. Cleland, and John M. Martinis. Coherent Josephson qubit suitable for scalable quantum integrated circuits. *Physical Review Letters*, 111(8):080502, 2013.

[54] Yu Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, A. Megrant, J. Y. Mutus, P. J. J. O'Malley, C. M. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, Michael R. Geller, A. N. Cleland, and John M. Martinis. Qubit architecture with high coherence and fast tunable coupling. *Phys. Rev. Lett.*, 113:220502, Nov 2014.

[55] A Vayonakis and J Zmuidzinas. Radiative losses from 2-d apertures, 2001.

[56] Robert W Jackson. Considerations in the use of coplanar waveguide for millimeter-wave integrated circuits. *IEEE Transactions on Microwave Theory and Techniques*, 34(12):1450–1456, 1986.

[57] J Wenner, M Neeley, Radoslaw C Bialczak, M Lenander, Erik Lucero, A D OConnell, D Sank, H Wang, M Weides, A N Cleland, and John M Martinis. Wirebond crosstalk and cavity modes in large chip mounts for superconducting qubits. *Superconductor Science and Technology*, 24(6):065001, 2011.

[58] John M Martinis, Rami Barends, and Alexander N Korotkov. Calculation of coupling capacitance in planar electrodes. *arXiv preprint arXiv:1410.3458*, 2014.

[59] Zijun Chen, A. Megrant, J. Kelly, R. Barends, J. Bochmann, Yu Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, J. Y. Mutus, P. J. J. O'Malley, C. Neill, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, A. N. Cleland, and John M. Martinis. Fabrication and characterization of aluminum airbridges for superconducting microwave circuits. *Applied Physics Letters*, 104(5):052602, 2014.

[60] Michael Tinkham. *Introduction to superconductivity*. Courier Corporation, 1996.

[61] Mark C Sweeney and Martin P Gelfand. Simple vortex states in films of type-I Ginzburg-Landau superconductor. *Physical Review B*, 82(21):214508, 2010.

[62] BLT Plourde, DJ Van Harlingen, N Saha, R Besseling, MBS Hesselberth, and PH Kes. Vortex distributions near surface steps observed by scanning squid microscopy. *Physical Review B*, 66(5):054529, 2002.

[63] DC Mattis and John Bardeen. Theory of the anomalous skin effect in normal and superconducting metals. *Physical Review*, 111(2):412, 1958.

[64] R. Barends, J. Wenner, M. Lenander, Y. Chen, R. C. Bialczak, J. Kelly, E. Lucero, P. OMalley, M. Mariantoni, D. Sank, H. Wang, T. C. White, Y. Yin, J. Zhao, A. N. Cleland, John M. Martinis, and J. J. A. Baselmans. Minimizing quasiparticle generation from stray infrared light in superconducting quantum circuits. *Applied Physics Letters*, 99(11):113507, 2011.

[65] Antonio D Córcoles, Jerry M Chow, Jay M Gambetta, Chad Rigetti, JR Rozen, George A Keefe, Mary Beth Rothwell, Mark B Ketchen, and M Steffen. Protecting superconducting qubits from radiation. *Applied Physics Letters*, 99(18):181906, 2011.

[66] John M. Martinis, K. B. Cooper, R. McDermott, Matthias Steffen, Markus Ansmann, K. D. Osborn, K. Cicak, Seongshik Oh, D. P. Pappas, R. W. Simmonds, and Clare C. Yu. Decoherence in Josephson qubits from dielectric loss. *Physical Review Letters*, 95(21):210503, 2005.

[67] J. Wenner, R. Barends, R. C. Bialczak, Yu Chen, J. Kelly, Erik Lucero, Matteo Mariantoni, A. Megrant, P. J. J. OMalley, D. Sank, A. Vainsencher, H. Wang, T. C. White, Y. Yin, J. Zhao, A. N. Cleland, and John M. Martinis. Surface loss simulations of superconducting coplanar waveguide resonators. *Applied Physics Letters*, 99(11):113513, 2011.

[68] David S Wisbey, Jiansong Gao, Michael R Vissers, Fabio CS da Silva, Jeffrey S Kline, Leila Vale, and David P Pappas. Effect of metal/substrate interfaces on radio-frequency loss in superconducting coplanar waveguides. *Journal of Applied Physics*, 108(9):093918, 2010.

[69] Moe S Khalil, FC Wellstood, and Kevin D Osborn. Loss dependence on geometry and applied power in superconducting coplanar resonators. *IEEE Transactions on Applied Superconductivity*, 21(3):879–882, 2011.

[70] GJ Grabovskij, LJ Swenson, O Buisson, C Hoffmann, A Monfardini, and J-C Villégier. In situ measurement of the permittivity of helium using microwave NbN resonators. *Applied Physics Letters*, 93(13):134102, 2008.

[71] Jeremy M Sage, Vladimir Bolkhovsky, William D Oliver, Benjamin Turek, and Paul B Welander. Study of loss in superconducting coplanar waveguide resonators. *Journal of Applied Physics*, 109(6):063915, 2011.

[72] C. Neill, A. Megrant, R. Barends, Yu Chen, B. Chiaro, J. Kelly, J. Y. Mutus, P. J. J. O'Malley, D. Sank, J. Wenner, T. C. White, Yi Yin, A. N. Cleland, and John M. Martinis. Fluctuations from edge defects in superconducting resonators. *Applied Physics Letters*, 103(7):072601, 2013.

[73] ME Day, M Delfino, W Tsai, A Bivas, and KN Ritz. Study of silicon surfaces bombarded with noble gas ions in an electron cyclotron resonance plasma. *Journal of Applied Physics*, 74(8):5217–5224, 1993.

[74] WA Phillips. Two-level states in glasses. *Reports on Progress in Physics*, 50(12):1657, 1987.

[75] J Ahn and JW Rabalais. Composition and structure of the Al_2O_3 $\{0001\}$-(1×1) surface. *Surface Science*, 388(1):121–131, 1997.

[76] C Niu, K Shepherd, D Martini, J Tong, JA Kelber, DR Jennison, and A Bogicevic. Cu interactions with α-Al_2O_3 (0001): effects of surface hydroxyl groups versus dehydroxylation by Ar-ion sputtering. *Surface Science*, 465(1):163–176, 2000.

[77] R. Barends, N. Vercruyssen, A. Endo, P. J. de Visser, T. Zijlstra, T. M. Klapwijk, P. Diener, S. J. C. Yates, and J. J. A. Baselmans. Minimal resonator loss for circuit quantum electrodynamics. *Applied Physics Letters*, 97(2), 2010.

[78] Michael R Vissers, Jeffrey S Kline, Jiansong Gao, David S Wisbey, and David P Pappas. Reduced microwave loss in trenched superconducting coplanar waveguides. *Applied Physics Letters*, 100(8):082602, 2012.

[79] Chang Hyun Jeong, Dong Woo Kim, Kyong Nam Kim, and Geun Young Yeom. A study of sapphire etching characteristics using BCl3-based inductively coupled plasmas. *Japanese Journal of Applied Physics*, 41(10R):6206, 2002.

[80] YP Hsu, Shoou-Jinn Chang, Yan-Kuin Su, Jinn-Kong Sheu, CH Kuo, CS Chang, and Shih-Chang Shei. ICP etching of sapphire substrates. *Optical Materials*, 27(6):1171–1174, 2005.

[81] H Furukawa, T Fukui, T Tanaka, A Noma, and D Ueda. A novel fabrication process of surface via-holes for GaAs power FETS. In *Gallium Arsenide Integrated Circuit (GaAs IC) Symposium, 1998. Technical Digest 1998., 20th Annual*, pages 251–254. IEEE, 1998.

[82] J.H. Wu, J.A. Del Alamo, and K.A. Jenkins. A high aspect-ratio silicon substrate-via technology and applications: through-wafer interconnects for power and ground and Faraday cages for SOC isolation. In *Electron Devices Meeting, 2000. IEDM '00. Technical Digest. International*, pages 477–480, Dec 2000.

[83] Kirt R Williams, Kishan Gupta, and Matthew Wasilik. Etch rates for micromachining processing-part II. *Microelectromechanical Systems, Journal of*, 12(6):761–778, 2003.

[84] Gregory TA Kovacs, Nadim I Maluf, and Kurt E Petersen. Bulk micromachining of silicon. *Proceedings of the IEEE*, 86(8):1536–1551, 1998.

[85] DB Fenner, DK Biegelsen, and RD Bringans. Silicon surface passivation by hydrogen termination: A comparative study of preparation methods. *Journal of Applied Physics*, 66(1):419–424, 1989.

[86] Takeshi Hattori, Stefan Heusler, JP Webb, and Takeshi Hattori. *Ultraclean surface processing of silicon wafers: secrets of VLSI manufacturing*. Springer Science & Business Media, 2013.

[87] S Ohya, B Chiaro, A Megrant, C Neill, R Barends, Y Chen, J Kelly, D Low, J Mutus, P J J OMalley, P Roushan, D Sank, A Vainsencher, J Wenner, T C White, Y Yin, B D Schultz, C J Palmstrm, B A Mazin, A N Cleland, and John M Martinis. Room temperature deposition of sputtered TiN films for superconducting coplanar waveguide resonators. *Superconductor Science and Technology*, 27(1):015009, 2014.

[88] M Naruse, Y Sekimoto, T Noguchi, A Miyachi, T Nitta, and Y Uzawa. Development of crystal Al MKIDs by molecular beam epitaxy. *Journal of Low Temperature Physics*, 167(3-4):373–378, 2012.

[89] P Dumas, YJ Chabal, and P Jakob. Morphology of hydrogen-terminated Si (111) and Si (100) surfaces upon etching in HF and buffered-HF solutions. *Surface Science*, 269:867–878, 1992.

[90] Taka-aki Miura, Michio Niwano, Daisei Shoji, and Nobuo Miyamoto. Kinetics of oxidation on hydrogen-terminated Si (100) and (111) surfaces stored in air. *Journal of Applied Physics*, 79(8):4373–4380, 1996.

[91] Le Thanh Vinh, M. Eddrief, C. A. Sbenne, P. Dumas, A. TalebIbrahimi, R. Gunther, Y. J. Chabal, and J. Derrien. Low temperature formation of Si (111) 7×7 surfaces from chemically prepared H/Si (111)-(1×1) surfaces. *Applied Physics Letters*, 64(24):3308–3310, 1994.

[92] James P Becker, Robert G Long, and John E Mahan. Reflection high-energy electron diffraction patterns of carbide-contaminated silicon surfaces. *Journal of Vacuum Science & Technology A*, 12(1):174–178, 1994.

[93] NF Mott. Conduction in non-crystalline materials: III. localized states in a pseudogap and near extremities of conduction and valence bands. *Philosophical Magazine*, 19(160):835–852, 1969.

[94] Rami Barends. *Photon-detecting superconducting resonators*. TU Delft, Delft University of Technology, 2009.

[95] Jiansong Gao, Miguel Daal, Anastasios Vayonakis, Shwetank Kumar, Jonas Zmuidzinas, Bernard Sadoulet, Benjamin A Mazin, Peter K Day, and Henry G Leduc. Experimental evidence for a surface distribution of two-level systems in superconducting lithographed microwave resonators. *Applied Physics Letters*, 92:152505, 2008.

[96] LJ Brillson, ML Slade, AD Katnani, M Kelly, and G Margaritondo. Reduction of silicon-aluminum interdiffusion by improved semiconductor surface ordering. *Applied Physics Letters*, 44(1):110–112, 1984.

[97] Gerald Chanin, EA Lynton, and B Serin. Impurity effects on the superconductive critical temperature of indium and aluminum. *Physical Review*, 114(3):719, 1959.

[98] B Abeles, Roger W Cohen, and GW Cullen. Enhancement of superconductivity in metal films. *Physical Review Letters*, 17(12):632, 1966.

[99] F Ladieu, J Le Cochec, P Pari, P Trouslard, and P Ailloud. Dielectric constant of glasses: evidence for dipole-dipole interactions. *Physical Review Letters*, 90(20):205501, 2003.

[100] J. Burnett, L. Faoro, I. Wisby, V. L. Gurtovoi, A. V. Chernykh, G. M. Mikhailov, V. A. Tulin, R. Shaikhaidarov, V. Antonov, P. J. Meeson, A. Ya. Tzalenchuk, and T. Lindstrm. Evidence for interacting two-level systems from the 1/f noise of a superconducting resonator. *Nature Communications*, 5, 2014.

[101] F Hoehne, Yu A Pashkin, O Astafiev, L Faoro, LB Ioffe, Y Nakamura, and JS Tsai. Damping in high-frequency metallic nanomechanical resonators. *Physical Review B*, 81(18):184112, 2010.

[102] J Lisenfeld, C Müller, JH Cole, P Bushev, A Lukashenko, A Shnirman, and AV Ustinov. Measuring the temperature dependence of individual two-level systems by direct coherent control. *Physical Review Letters*, 105(23):230504, 2010.

[103] Clemens Müller, Jürgen Lisenfeld, Alexander Shnirman, and Stefano Poletto. Interacting two-level defects as sources of fluctuating high-frequency noise in superconducting circuits. *Physical Review B*, 92(3):035442, 2015.

[104] Lara Faoro and Lev B Ioffe. Interacting tunneling model for two-level systems in amorphous materials and its predictions for their dephasing and noise in superconducting microresonators. *Physical Review B*, 91(1):014201, 2015.

[105] PM Tedrow, G Faraci, and R Meservey. Measurement of the surface inductance and penetration depth of superconducting aluminum. *Physical Review B*, 4(1):74, 1971.

[106] Gheorghe Stan, Stuart B Field, and John M Martinis. Critical field for complete vortex expulsion from narrow superconducting strips. *Physical Review Letters*, 92(9):097003, 2004.

[107] C Song, TW Heitmann, MP DeFeo, K Yu, R McDermott, M Neeley, John M Martinis, and BLT Plourde. Microwave response of vortices in superconducting thin films of Re and Al. *Physical Review B*, 79(17):174512, 2009.

[108] Eric Bronson, Martin P Gelfand, and Stuart B Field. Equilibrium configurations of pearl vortices in narrow strips. *Physical Review B*, 73(14):144501, 2006.

[109] Stuart B Field, SS James, J Barentine, V Metlushko, G Crabtree, H Shtrikman, B Ilic, and SRJ Brueck. Vortex configurations, matching, and domain structure in large arrays of artificial pinning centers. *Physical Review Letters*, 88(6):067003, 2002.

[110] Guo-Chun Liang, Yao-Wu Liu, and Kenneth K Mei. Full-wave analysis of coplanar waveguide and slotline using the time-domain finite-difference method. *IEEE Transactions on Microwave Theory and Techniques*, 37(12):1949–1957, 1989.

[111] B Chiaro, A Megrant, A Dunsworth, Z Chen, R Barends, B Campbell, Y Chen, A Fowler, I C Hoi, E Jeffrey, J Mutus, C Neill, P J J O'Malley, C Quintana, P Roushan, D Sank, A Vainsencher, J Wenner, T C White, and J Martinis. Dielectric surface loss in

superconducting resonators with flux-trapping holes. Manuscript submitted for publication, 2016.

[112] I Nsanzineza and BLT Plourde. Trapping a single vortex and reducing quasiparticles in a superconducting resonator. *Physical Review Letters*, 113(11):117002, 2014.

[113] C Song, MP DeFeo, K Yu, and BLT Plourde. Reducing microwave loss in superconducting resonators due to trapped vortices. *Applied Physics Letters*, 95(23):232501, 2009.

[114] C. M. Quintana, A. Megrant, Z. Chen, A. Dunsworth, B. Chiaro, R. Barends, B. Campbell, Yu Chen, I.-C. Hoi, E. Jeffrey, J. Kelly, J. Y. Mutus, P. J. J. O'Malley, C. Neill, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, A. N. Cleland, and John M. Martinis. Characterization and reduction of microfabrication-induced decoherence in superconducting quantum circuits. *Applied Physics Letters*, 105(6):062601, 2014.

[115] John M Martinis and A Megrant. Ucsb final report for the csq program: Review of decoherence and materials physics for superconducting qubits. *arXiv preprint arXiv:1410.5793*, 2014.

[116] Chen Wang, Christopher Axline, Yvonne Y Gao, Teresa Brecht, Y Chu, L Frunzio, MH Devoret, and RJ Schoelkopf. Surface participation and dielectric loss in superconducting qubits. *Applied Physics Letters*, 107(16):162601, 2015.

[117] Evan Jeffrey, Daniel Sank, J. Y. Mutus, T. C. White, J. Kelly, R. Barends, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. Megrant, P. J. J. O'Malley, C. Neill, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and John M. Martinis. Fast accurate state measurement with superconducting qubits. *Physical Review Letters*, 112(19):190504, 2014.

[118] J. Y. Mutus, T. C. White, R. Barends, Yu Chen, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, J. Kelly, A. Megrant, C. Neill, P. J. J. O'Malley, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, K. M. Sundqvist, A. N. Cleland, and John M. Martinis. Strong environmental coupling in a Josephson parametric amplifier. *Applied Physics Letters*, 104(26):263513, 2014.

[119] C Neill, P Roushan, M Fang, Y Chen, M Kolodrubetz, Z Chen, A Megrant, R Barends, B Campbell, B Chiaro, A Dunsworth, E Jeffrey, J Kelly, J Mutus, PJJ O'Malley, C Quintana, D Sank, A Vainsencher, J Wenner, TC White, A Polkovnikov, and JM Martinis. Ergodic dynamics and thermalization in an isolated quantum system. *arXiv preprint arXiv:1601.00600*, 2016.

[120] Ivan Kassal, Stephen P Jordan, Peter J Love, Masoud Mohseni, and Alán Aspuru-Guzik. Polynomial-time quantum algorithm for the simulation of chemical dynamics. *Proceedings of the National Academy of Sciences*, 105(48):18681–18686, 2008.

[121] CY Lin, PC Stancil, H-P Liebermann, P Funke, and RJ Buenker. Inelastic processes in collisions of Na (3 s, 3 p) with He at thermal energies. *Physical Review A*, 78(5):052706, 2008.

[122] J Ostgaard Olsen, T Andersen, and N Andersen. Excitation of Li and Na in collisions with He and Ne: measurements of Li I and Na I resonance-line emission and polarization. *Journal of Physics B: Atomic and Molecular Physics*, 10(9):1723, 1977.

[123] HSW Massey. Collisions between atoms and molecules at ordinary temperatures. *Reports on Progress in Physics*, 12(1):248, 1949.

[124] J Brooke, D Bitko, TF Rosenbaum, and G Aeppli. Quantum annealing of a disordered magnet. *Science*, 284(5415):779–781, 1999.

[125] Masuo Suzuki. Generalized Trotter's formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems. *Communications in Mathematical Physics*, 51(2):183–190, 1976.

[126] Luke Gordon, Hazem Abu-Farsakh, Anderson Janotti, and Chris G Van de Walle. Hydrogen bonds in Al_2O_3 as dissipative two-level systems in superconducting qubits. *Scientific Reports*, 4, 2014.

[127] GJ Dolan. Offset masks for lift-off photoprocessing. *Applied Physics Letters*, 31(5):337–339, 1977.

[128] Julian Kelly. *Fault-tolerant superconducting qubits*. PhD thesis, University of California Santa Barbara, 2015.

[129] T. C. White, J. Y. Mutus, I.-C. Hoi, R. Barends, B. Campbell, Yu Chen, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, J. Kelly, A. Megrant, C. Neill, P. J. J. O'Malley, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, S. Chaudhuri, J. Gao, and John M. Martinis. Traveling wave parametric amplifier with Josephson junctions using minimal resonator phase matching. *Applied Physics Letters*, 106(24):242601, 2015.

[130] Ioan-Mihai Pop, T Fournier, T Crozes, Florent Lecocq, Iulian Matei, Bernard Pannetier, Olivier Buisson, and Wiebke Guichard. Fabrication of stable and reproducible submicron tunnel junctions. *Journal of Vacuum Science & Technology B*, 30(1):010607, 2012.

[131] Daniel Sank. *Fast, Accurate State Measurement in Superconducting Qubits*. PhD thesis, University of California Santa Barbara, 2014.