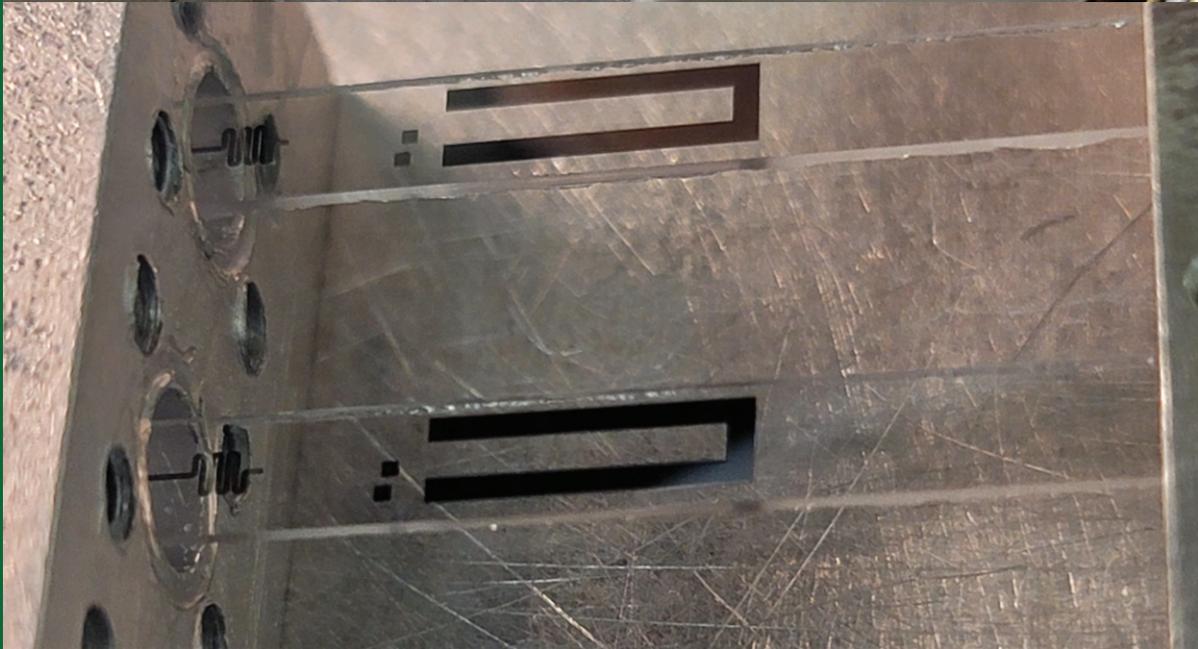
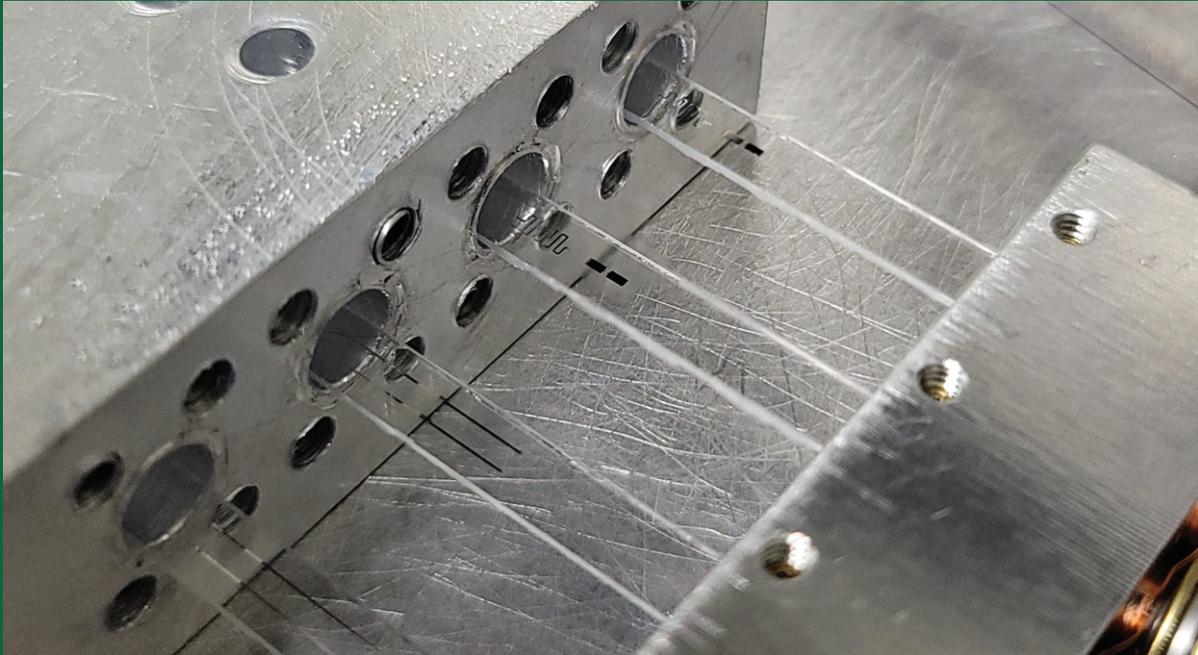


Improving the Coherence of Superconducting Quantum Circuits through Loss Characterization and Design Optimization



Suhas Ganjam

Abstract

Improving the Coherence of Superconducting Quantum Circuits through Loss Characterization and Design Optimization

Suhas Sham Ganjam

2023

Realizing a practical quantum computer with superconducting qubits requires substantially higher gate fidelities, which necessitates further improvements in the coherence of superconducting quantum circuits. Steady improvements have been made over the past two decades, including substantial exploration into energy relaxation mechanisms in superconducting thin films. However, a comprehensive and quantitative understanding of the relative impact of these mechanisms does not yet exist. In this thesis, I utilize a multimode approach to systematically characterize microwave losses in the quantum regime, with the goals of understanding relaxation-limiting loss mechanisms and improving device coherence through materials, process, and circuit design optimization. Using this approach, we measure significant reductions in surface losses by employing a tantalum-based materials platform, and in bulk substrate loss by utilizing high-temperature annealing processes. With this knowledge we predict and experimentally verify the relaxation times of aluminum and tantalum-based transmon qubits. We additionally optimize device geometry to maximize coherence within a coaxial tunnel architecture, and realize on-chip quantum memories with single-photon Ramsey times of 2.0–2.7 ms, limited by their energy relaxation times of 1.0–1.4 ms. This demonstrates an important link between microwave loss characterization and improving coherence in superconducting qubits, and enables a more modular and compact coaxial circuit architecture for bosonic qubits with reproducibly high coherence.

**Improving the Coherence of Superconducting
Quantum Circuits through Loss
Characterization and Design Optimization**

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

by
Suhas Sham Ganjam

Dissertation Director: Professor Robert J. Schoelkopf

May 2024

Copyright © 2023 by Suhas Sham Ganjam

All rights reserved.

Acknowledgments

The marathon that is completing a Ph.D. is not something I could have done on my own. There are so many people I want to thank, including those who have supported me through the last few years and those who have taught me essential hard and soft skills that I will continue to carry and benefit from for the rest of my life. My journey did not start at Yale. My motivation to study physics came about during my time at UC Berkeley, where I was humbled by the students studying there, and even more so by the brilliant professors and scientists I had the opportunity to learn from. In particular, working in the labs of AKITO KUSAKA, MATT PYLE, and BERNARD SADOULET first stimulated my interest in pursuing a career in physics by exposing me to the rich and exciting world of possibilities that could be found there. Of course, the mentorship, optimism, and inspiration I received from CHARLIE HILL and BRUNO SERFASS helped me forge my own path and develop my own interests as I left Berkeley and began graduate school.

I have been advised by ROB SCHOELKOPF for the entirety of graduate school. When I first met him, I was inspired by his optimism and excitement to study quantum information science and by his drive towards building a quantum computer. As I progressed as a graduate student, I was amazed at the depth of his knowledge and his eagerness to spend over ten hours in a single day talking about science with seemingly never-ending stamina. Rob always has an eye on the bigger picture and is able to contextualize your project and convince you that what you are doing is important. The remainder of my thesis committee, MICHEL DEVORET, LEONID GLAZMAN, and PETER RAKICH, have served as foundations of the Yale Quantum Institute and the amazing scientific community that has been fostered here. Through their own personal expertise and the research groups they have developed over the years, I was confident that I could always find someone to turn towards to fill the gaps of my lack of understanding in certain subject areas. In particular, Rob and Michel have fostered an environment here on the fourth floor of Becton that

provides the best place on Earth to study circuit quantum electrodynamics. The amount of resources and knowledge to be found here is second to none, and it has contributed to ground-breaking achievements and developments in the field over the years. I also want to thank my external reader, NATHALIE DE LEON, whose background in quantum information and materials science has supplemented my own understanding on how importance materials are to qubit coherence.

LUIGI FRUNZIO has been integral to the continued proper functioning of the lab and comprises a large part of the historical knowledge that I have benefited from throughout my time here. Luigi has a vast bank of knowledge and experience in experimentation, device fabrication, materials science, and quantum information. Luigi is always excited to help with projects and eager to get his hands dirty. In particular, his experience in tantalum sputtering laid the groundwork for the processes I was able to develop for this thesis. More than that, Luigi was always a source of levity and comforting advice when things do not go as planned.

Over the years, I worked closely with CHAN U LEI, LEV KRAYZMAN, YANHAO WANG, ARCHAN BANERJEE, and NICO ZANI. Chan U and Lev served as early mentors for me as a junior graduate student. They were always willing to answer my questions and showed me the ropes of existing in RSL. Chan U was particularly brilliant and creative, and determined to solve problems no matter how small and aggravating they were. Much of what I learned about device fabrication and process development was due to his mentorship. I enjoyed working with Lev due to his ability to think deeply about the problems we faced in the field and ask difficult and important questions about the work we were pursuing. Lev would often times provide necessary challenges to claims that I make or assumptions that I have to force me to re-examine my thoughts and ideas with a critical eye. Through our conversations, I was able to develop stronger motivations for pursuing the projects I chose and more well-developed justifications for the conclusions

I drew. I enjoyed working with Yanhao, Archan, and Nico over the past two years and am inspired by their dedication and excitement to pick up new projects and explore them deeply. Yanhao in particular worked to develop a more efficient and streamlined way to measure resonators, which allowed us to wrap up and conclude the loss characterization project at the pace we were able to do so. Yanhao's interest in the project also led to new insights and conclusions that would not have happened without him.

In my third year I began working with YAO LU. Yao was a treasure-trove of knowledge in cQED and planar circuit architectures. I learned so much from him about flux modulation, qubit measurements, and general quantum mechanics. Our conversations always lasted so much longer than we originally intended, but often times led to new insights and ways of thinking about quantum circuits, their design, and their control. He was a great friend and confidant who bore with me as I sometimes vented my frustrations to him. My only regret was that I was not able to work more with him, but hope our paths cross again in the future.

I have also had the opportunity to collaborate with many others, which have led to some great accomplishments within the field and close personal friendships. Working with JACOB CURTIS, ALEC EICKBUSCH, VLAD SIVAK, ANIKET MAITI, JOHN GARMON, ALEX READ, and BEN CHAPMAN taught me much about quantum error correction, parametric flux modulation, precision microwave resonator measurements, and cavity-cavity beamsplitters. I have greatly enjoyed my conversations with these colleagues. On the device fabrication front, I benefited greatly from the expertise and knowledge of the cleanroom staff at Yale. YONG SUN, LAUREN MCCABE, KELLY WOODS, and SEAN RINEHART were instrumental in maintaining the cleanroom equipment and in advising me on the development of sapphire annealing tantalum sputtering, and tantalum etching processes. My conversations with NATHALIE DE LEON, MICHAEL HATRIDGE, AVEEK DUTTA, and ALEX PLACE also assisted in my understanding

of various sample cleaning processes that became necessary to achieve high coherence in tantalum-based devices. Additionally, I learned much about materials characterization from our collaborators at Brookhaven National Laboratory. I learned much about TEM, XRD, and XPS from KIM KISSLINGER, MINGZHAO LIU, ANDI BARBOUR, ANDREW WALTER, IGNACE JARRIGE, ASWIN ANBALAGAN, RUOSHUI LI, and CHENYU ZHOU, which has motivated me to explore further what physical materials characterization can do for us in cQED. And of course, I must acknowledge the work and efforts by the YQI and Becton administrative staff who have made life easier for us all and have kept the place running: TERRI EVANGELISTE, NUCH GRAVES, GISELLE MAILLET, MARIA RAO, ALEX BOZZI, FLORIAN CARLE, and RACQUEL MILLER.

Finally, I would like to thank my family. My parents SHASHI and SHAM have always been a source of comfort and emotional stability, and are always willing to drop everything to help me in a time of need. I would not have been able to get through graduate school without their support. My brother SANJEET and sister-in-law VRUSHALI are people I look up to as role models. The life they have built together and the dedication and drive they have in their professions have motivated and inspired me to do the same. I am very fortunate to have them in my life.

The work in this thesis was supported by the U.S. Army Research Office (ARO) under grants W911-NF-18-1-0212, and W911-NF-23-1-0051 and by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Co-design Center for Quantum Advantage (C2QA) under contract No. DE-SC0012704.

Contents

Contents	xi
List of Figures	xiv
List of Tables	xv
1 Introduction	1
1.1 What is Quantum Computing?	3
1.2 Experimental Quantum Computing	6
1.3 The Problem of Decoherence	9
1.4 Thesis Overview	11
1.5 Suggestions for Further Reading	12
2 Circuit Quantum Electrodynamics	14
2.1 Superconducting Circuit Quantization	15
2.1.1 LC Oscillator	15
2.1.2 The Josephson Junction	17
2.1.3 The Cooper-Pair Box	20
2.1.4 The Transmon Qubit	22
2.2 Classical Drives on Superconducting Quantum Circuits	25
2.2.1 Driving a Linear Oscillator	26

2.2.2	Driving a Transmon Qubit	28
2.3	Coupling a Transmon to a Linear Oscillator	32
2.3.1	The Dispersive Limit	36
2.3.2	Transmon State Readout	39
2.3.3	Quantum Information Storage in a Resonator	42
2.4	cQED in an Arbitrary Electromagnetic Environment	46
2.4.1	Black Box Quantization	47
2.4.2	Energy-Participation Quantization	48
2.5	cQED Architectures	50
2.5.1	Planar Architecture	50
2.5.2	3D Cavity Architecture	54
2.5.3	Coaxial Tunnel Architecture	58
2.6	Suggestions for Further Reading	59
3	Decoherence in Superconducting Quantum Circuits	61
3.1	The Participation Ratio Model	64
3.2	External Losses	67
3.2.1	The Purcell Effect	71
3.3	Internal Losses	75
3.3.1	Conductor Loss	75
3.3.2	Dielectric Loss	85
3.3.3	Contact (Seam) Loss	95
3.4	Dephasing	97
3.4.1	TLS Noise	98
3.4.2	Thermal Photon Shot Noise	98
3.4.3	Flux Noise	100

3.5	Coherence Measurement Techniques	101
3.5.1	Device Packaging	101
3.5.2	Measurement Setup	103
3.5.3	Resonator Quality Factor Measurements	105
3.5.4	Transmon Coherence Measurements	114
3.5.5	Quantum Memory Coherence Measurements	118
4	Microwave Loss Characterization	123
4.1	Simple Loss Characterization Methods	125
4.2	Matrix Loss Characterization	127
4.3	The Multimode Approach to Loss Characterization	131
4.4	Forky Whispering-Gallery-Mode Resonator	132
4.4.1	FWGMR device design	134
4.4.2	Participations & measurement sensitivity	136
4.4.3	Device measurement & loss extraction	137
4.4.4	Comparison of loss factors of different materials & processes	138
4.4.5	Conductor loss in bulk Al	142
4.4.6	Predicting losses in 3D cavity resonators	145
4.5	Tripole Stripline	147
4.5.1	Multimode stripline device design	147
4.5.2	TSL/ASL participations & measurement sensitivity	152
4.5.3	Device measurement & loss extraction	154
4.5.4	Comparison of loss factors of different materials & processes	159
4.5.5	Conductor loss in thin-films	165
5	Understanding and Optimizing Losses in Planar Transmon Qubits and Quantum Memories	170

5.1	Validating the Loss Model with Qubit Measurements	171
5.1.1	Transmon Device Design	171
5.1.2	Transmon Participations	173
5.1.3	Characterizing Ta/Al Contact Loss	176
5.1.4	Predicting Transmon Loss	180
5.1.5	Verifying Transmon Coherence Predictions	181
5.2	Optimized Geometry to Maximize Coherence in a Quantum Memory . .	186
5.2.1	Quantum Memory Device Design	187
5.2.2	Hairpin Stripline Participations	189
5.2.3	Predicting Hairpin Stripline Loss	190
5.2.4	Verifying Hairpin Stripline Coherence Predictions	191
6	Future Directions and Conclusion	194
6.1	Open Questions About Microwave Losses	194
6.2	Future Directions in Microwave Loss Characterization	197
6.2.1	Exploring Frequency Dependence with Multimode Resonators . .	197
6.2.2	Measuring Conductor and MA Losses with Flip-Chip Circuits . .	197
6.2.3	Characterizing Losses in Rhenium	202
6.2.4	Rhenium-Based Transmons with Long Relaxation Times	206
6.3	A Flip-Chip Architecture for Bosonic Qubits	208
6.4	Summary	210
	Bibliography	226
	A Device fabrication for Thin-Film Devices	227
	Appendices	227
A.1	Substrate Preparation	227

A.2	Aluminum/Josephson Junction Process	228
A.3	Tantalum Process	229
A.4	Rhenium Process	230
A.5	Flip-Chip Process	232
A.6	Wafer Dicing	233
A.7	Flip-chip Bonding	233
B	Materials Characterization	235
B.1	Sapphire Annealing	235
B.2	Aluminum Film Characterization	236
B.3	Tantalum Film Characterization	237
B.4	Rhenium Film Characterization	238
C	Surface Participation Ratio Simulations for Thin Films	241
D	Copyright Permissions	245

List of Figures

1.1	Bloch sphere representation of a qubit	4
2.1	LC oscillator	16
2.2	Josephson junction	18
2.3	Cooper-Pair Box	21
2.4	Transmon qubit energy spectrum	24
2.5	Driving an LC oscillator and transmon	26
2.6	Transmon Rabi oscillations	30
2.7	Transmon coupled to an LC resonator	33
2.8	Dispersive readout	40
2.9	Simulated photon number statistics in a resonator	43
2.10	Josephson junction embedded in an arbitrary electromagnetic environment	47
2.11	Planar cQED architecture	51
2.12	3D cavity cQED architecture	55
2.13	Coaxial tunnel architecture	57
3.1	External vs internal environment	62
3.2	LC resonator coupled to an external resistor	68
3.3	External coupling schemes in cQED	70
3.4	Purcell effect in a qubit coupled to a readout resonator	72

3.5	Temperature dependence of resonator frequency and Q due to thermal quasiparticles	81
3.6	Lossy dielectric regions	86
3.7	Power and temperature dependence of TLSs	93
3.8	Device Packaging	102
3.9	Fridge Wiring Diagram	104
3.10	Resonator Measurement Configurations	107
3.11	Resonator Fitting	113
3.12	Transmon coherence measurement pulse sequences	115
3.13	Transmon coherence measurements	116
3.14	Quantum memory coherence measurement pulse sequences	119
3.15	Quantum memory coherence measurements	121
4.1	Forky whispering-gallery-mode resonator (FWGMR)	133
4.2	FWGMR modes and measurement sensitivity	135
4.3	Extracted single-photon loss factors for bulk superconductors	139
4.4	TEM of unetched and etched bulk Al	141
4.5	Predicted quality factor and loss budget of 3D cavity resonators	146
4.6	Tripole striplines in the coaxial tunnel architecture	148
4.7	TSL & ASL device design	149
4.8	TSL & ASL measurement sensitivity	155
4.9	TSL loss extraction	157
5.1	3D Transmon qubit device design	172
5.2	Segmented stripline to extract Ta/Al contact loss	177
5.3	Predicting transmon loss	180
5.4	Al- and Ta-based transmon T_1	182

5.5	Temporal fluctuations in transmon coherence & resonator quality factor . . .	183
5.6	Predicted vs measured transmon quality factors	185
5.7	Temporally “stable” transmon	185
5.8	Hairpin stripline quantum memory	188
5.9	Hairpin stripline quantum memory coherence	191
5.10	Temporal fluctuations in quantum memory coherence	192
6.1	Flip-chip circuit architecture	199
6.2	Indium redeposition on rhenium devices	204
6.3	Flip-chip transmon and quantum memory	207
6.4	Rhenium transmon T_1	208
B.1	Atomic force microscopy of sapphire surface	235
B.2	TEM of Al/AlO _x /Al film	236
B.3	TEM of a tantalum film	237
B.4	Tantalum film characterization	237
B.5	TEM of a rhenium film	239
B.6	Rhenium film T_c	240
C.1	Simulating edge-field behavior of thin films	242
C.2	Simulating near-JJ fields	244

List of Tables

4.1	FWGMR Participation Matrix	136
4.2	FWGMR extracted loss factors of 6061 Al	138
4.3	Tripole & adjacent stripline dimensions	150
4.4	TSL and ASL participation matrices	151
4.5	TSL & ASL loss factors at $\bar{n} = 1$	162
4.6	TSL & ASL loss factors at high power	163
4.7	Average single-photon surface & bulk loss factors for thin-film devices . .	164
4.8	TSL & ASL conductor participation	166
5.1	Typical transmon & readout parameters	171
5.2	Transmon participations	175
5.3	Segmented stripline participations	179
5.4	Segmented stripline loss	179
5.5	Typical hairpin stripline parameters	189
5.6	Hairpin stripline participations	190
6.1	Flip-Chip Participation Matrix	202
6.2	Rhenium Device Quality Factors Measured at $\bar{n} \approx 1$	203
6.3	Flip-Chip Loss Factors Measured at $\bar{n} \approx 1$	205

Chapter 1

Introduction

Progress in experimental quantum information science has advanced tremendously over the past six years. When I arrived at Yale, state-of-the-art quantum processors consisted of single-digit numbers of qubits, whereas more recently, industry leaders have created processors made up of hundreds of qubits and are carrying out more complex experiments than ever before. These developments have been preceded by decades of development towards understanding the nature of quantum information, and usefulness of quantum computing, and the experimental control of single quantum systems. Interest in quantum computing arose in the 1980s when Richard Feynman proposed the idea of a machine that exploits the features of quantum mechanics to solve problems that would be otherwise difficult or inefficient to solve[1]. Feynman expressed concern about the inefficiencies of simulating even mid-sized quantum systems using classical computers, and argued that the best way to simulate a quantum system is to use a computer that also behaves quantum mechanically. This observation spurred exploration towards supporting Feynman’s claim, and in 1985 David Deutsch introduced the idea of quantum parallelism as a way to solve particular problems faster than they could be done classically[2], suggesting that the inherent “weirdness” of quantum mechanics could perhaps be exploited to our technological advantage.

In 1994 Peter Shor developed one of the earliest and perhaps most well-known examples of a useful quantum algorithm[3]. Shor showed exponential speedup relative to classical computers through the use of multiple quantum algorithms that could factor and take discrete logarithms of large numbers. These types of problems are hard to solve but easy to verify with a classical computer; that is, factoring a large number takes an amount of time that scales exponentially with the bit-size of the number, but multiplying the two factors together is easy. This has important implications in cryptography, as problems that are easy to verify but hard to solve serve as the basis for modern encryption systems such as RSA[4]. However, quantum computers, due to their ability to easily solve these types of problems, can effectively break the RSA encryption system, which demonstrates a need for the future implementation of quantum encryption schemes that can be made unbreakable[5, 6].

Three years after Shor, Lov Grover introduced a quantum algorithm to search through unstructured data sets quadratically faster than a classical one[7], which would bring much-needed efficiency improvements for any problem that requires exhaustive searches through large data sets. There has since been an explosion of ideas on how to best use a quantum computer and it is becoming more and more clear that its successful implementation can have revolutionary implications to multiple areas of technology, such as cryptography, computer science, materials science, chemistry, and quantum simulation. However, on a more sobering note, we are still far away from having a computer that can make meaningful advances on these fronts. This is partially due to the infancy of the technology (however rapid its advancement), but also partially because of the enormous power and versatility of modern-day classical computing. Since classical computing has had over 70 years to develop, quantum computing is decades behind; the first fault-tolerant fully functional quantum computer will likely not be able to solve useful problems faster than a classical supercomputer. However, the gap would begin to close quite rapidly due

to the improved efficiency that quantum computers would bring, and eventually quantum computers would overtake even the fastest classical supercomputers and deliver new breakthroughs in science and technology that could not have been achieved otherwise.

This introduction is intended to provide a general understanding of quantum computing and quantum speedup, its implementation in experimental physics, and the introduction of the circuit quantum electrodynamics (cQED) platform. I will then introduce the problem of coherence in superconducting quantum circuits, which will motivate the remainder of this thesis.

1.1 What is Quantum Computing?

Modern classical computers work by performing operations on bits, physical objects that can take on one of two possible states. We define these states arbitrarily as either 0 *or* 1, and all information processing and storage involve manipulation of these states using logical operations, or gates. In a quantum computer, the bit is replaced by its quantum analog, a quantum bit, or “qubit”. A qubit is a quantum-mechanical object that can exist in a coherent superposition of 0 *and* 1. More concretely, we can define the state $|\psi\rangle$ of a qubit as[8]

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

where α and β are complex numbers with relationship $|\alpha|^2 + |\beta|^2 = 1$. Like a classical bit, a qubit can be in state $|0\rangle$ (by setting $\beta = 0$) or state $|1\rangle$ (by setting $\alpha = 0$). However, unlike a classical bit, a qubit can also be in a state represented by a linear combination of $|0\rangle$ and $|1\rangle$, which is what we call a superposition. We can also describe the states of a qubit as spanning a 2-dimensional complex vector space (Hilbert space) with $|0\rangle$ and $|1\rangle$ serving as the orthogonal basis vectors.

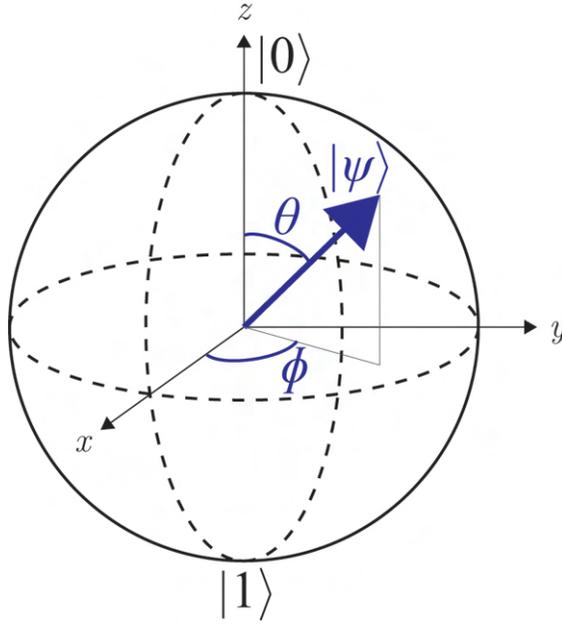


Figure 1.1: **Bloch sphere representation of a qubit.**

A more intuitive description for the states of a qubit is provided by the Bloch sphere representation, where the states of the qubit can be represented by points on the surface of a three-dimensional unit sphere. In such a description, we can re-parametrize the vector description of the qubit state using the a polar coordinate system:

$$|\psi\rangle = \cos \frac{\theta}{2} |0\rangle + e^{i\phi} \sin \frac{\theta}{2} |1\rangle, \quad (1.2)$$

where θ is the polar angle and ϕ is the azimuthal angle, as shown in Fig 1.1. Here, the $|0\rangle$ and $|1\rangle$ states are located at the north and south poles, respectively, and any other point on the sphere is a superposition state. The Bloch sphere representation intuitively highlights one of the main differences between a classical bit and a quantum bit. Classical bits can be either 0 or 1; in other words, classical bits are restricted to lie on the poles of the sphere, and classical logic gates can only transform 0 to 1 or vice versa. Qubits, on the other hand, can reside anywhere on the surface; there is a continuum of states between $|0\rangle$ and $|1\rangle$, and quantum logic gates can allow the qubit state to traverse from any one point of the

sphere to another. This allows N qubits to span a state space of dimensionality 2^N ; the exponential growth of the state space with qubit number provides an important resource for quantum information processing[8].

Another important resource in quantum computing is entanglement, a highly non-intuitive and uniquely quantum phenomenon. Entanglement refers to the correlation between the states of two distinct qubits; as a result, one bit of information can be shared by multiple qubits *regardless* of their physical separation from one another. The canonical example of an entangled pair of qubits is a Bell state:

$$|\Psi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \quad (1.3)$$

For this Bell state, if the measurement of one of the qubits yields $|0\rangle$, the state of the other qubit is immediately known to also be $|0\rangle$ (and vice-versa for $|1\rangle$). In other words, the states of the qubits are correlated; operations performed on one of the qubits affects the other qubit. This is a resource that can be exploited in numerous ways. Entanglement can play a role in quantum measurement, where a qubit state can become entangled with a pointer state of another physical system that is easier to observe. Entanglement can also serve as the basis for quantum cryptographic schemes, enabling quantum encryption protocols that are unbreakable with our current understanding of physics[5]. However, when entanglement is uncontrolled, it can be a significant hindrance, such as when a qubit becomes entangled with its environment, causing decoherence and loss of information.

Superposition and entanglement are the two features of quantum systems that enable quantum speedup in a quantum computer. These properties are preserved as long as the system remains undisturbed. Measurement disturbs the system, and forces the superposition to collapse, yielding a discrete state. For example, measurement of the qubit defined in Eq. (1.1) yields 0 with probability $|\alpha|^2$ or 1 with probability $|\beta|^2$. However, as long as

one does not perform a measurement and the system remains otherwise undisturbed, the superposition and entanglement will be preserved. This allows for the implementations of clever algorithms that allow parallel computation, also known as quantum parallelism[2]. Quantum parallelism allows for the simultaneous evaluation of a function for many different values of its arguments. By initializing multiple qubits in a superposition state, a quantum computer can evaluate a function $f(x)$ for multiple values of x *simultaneously*. This can then be followed by the use of constructive and destructive interference to concentrate the weights of the superposition towards the correct or desired solution. This is in sharp contrast to a classical computer, which needs to perform a distinct evaluation of $f(x)$ for *each* value of x . With properly designed quantum algorithms, a quantum computer can solve certain problems far more efficiently than a classical computer.

1.2 Experimental Quantum Computing

Attempts to build a quantum computer have been underway since the 1990s. Whereas a classical bit can be made up of any physical system that has two states, such as a coin or a transistor, a qubit must be realized using a *quantum* system that can exist in a superposition of two *quantum* states – a “quantum transistor”. However, this is not the only requirement. A proposed quantum computing platform based on a particular quantum system must satisfy five requirements, as defined by David DiVincenzo[9]:

1. *A scalable physical system with well characterized qubits*
2. *The ability to initialize the state of the qubits to a simple fiducial state*
3. *Long relevant decoherence times, much longer than the gate operation time*
4. *A “universal” set of quantum gates*

5. *A qubit-specific measurement capability*

Some of these criteria are straightforward to understand. Large-scale quantum computers require large numbers of qubits that can interact together to preserve superposition and entanglement without leaving the computational space. Additionally, the qubits need to be controllable; on-demand high-fidelity control is required in order to initialize the quantum computer and enact a set of quantum gates that can allow the computer to explore the entire computational space. Finally, in order to extract information from the computer at the end of a computation, the qubits must be measurable in a process called “readout”. However, the above requirements can only be fulfilled if coherence times are long, which will be discussed further in Sec. 1.3.

The DiVincenzo criteria are very broadly defined with good reason: every quantum computing platform has a different method of realizing a qubit, whether it is through ion traps[10], nuclear spins[11, 12], neutral atoms[13, 14], electron spins in quantum dots[15], or superconducting circuits[16–18]. As a result, the paths to fulfill the criteria and the relative importance one should place on specific criteria can vary based on the platform being used. In this thesis, I will focus on the superconducting circuits platform, which has been the specific focus of my graduate studies and forms the basis for the rest of this thesis. The platform utilizes superconducting circuits to realize what is known as an artificial atom that serves as the qubit. Superconducting devices are operated at microwave (\sim GHz) frequencies and cooled to nearly absolute zero in dilution refrigerators ($T \sim 20$ mK) where the circuits become nearly lossless and thermal fluctuations become negligible, allowing quantum effects to dominate. Gate operations are enacted with well-controlled and calibrated microwave pulses. The use of superconducting circuits to realize and control qubits is described under the umbrella of circuit quantum electrodynamics (cQED)[19].

The use of an artificial atom sets cQED apart from the other aforementioned platforms, which use individual atoms or elementary particles as their qubits. By doing so, these plat-

forms ensure that each of their qubits is identical. Moreover, the microscopic size of these qubits and small interaction with their environment results in very slow information loss rates. However, these features also make it difficult to control the qubits, and as a result, the time it takes to perform gate operations is quite long. In contrast, the artificial atom is macroscopic in size and is made up of a large number of atoms whose collective behavior is quantum-mechanical. Their much larger size provides them with intrinsically larger dipole moments, allowing them to couple strongly to other qubits or control elements. This results in much faster gate speeds and more efficient control and readout sequences.

The nature of the artificial atom (or superconducting qubit) being an electrical circuit makes the cQED platform highly versatile. Since circuit parameters are ideally well-defined and easily designed, the qubit can be designed and fine-tuned from a circuit model point-of-view to attain desired couplings within a larger system. Qubit control is realized by applying microwave pulses, for which there have been decades worth of technological development by the telecommunications industry. Furthermore, the platform is straightforwardly scalable, as the circuits themselves are lithographically patterned with standard techniques that have been in use for decades in the fabrication of computer chips. The enormous versatility of superconducting circuits over other platforms have cemented it as one of the leading platforms of choice by industry for the realization of a functional quantum computer. Quantum computers with over 50 qubits have already been demonstrated[20, 21] using the cQED platform. However, despite its success, the cQED platform suffers from a major problem that hinders its further advancement: decoherence, which leads to the loss of quantum information.

1.3 The Problem of Decoherence

Compared to classical bits, qubits are far more sensitive to noise processes that can change or scramble the state. Classical bits are conventionally realized by the voltage of a transistor. In this scheme, small voltage fluctuations can be tolerated, as the distinction between the 0 and 1 states are usually thresholded; for example, the 0 state may correspond to a voltage between 0 – 2.5 V, and the 1 state may correspond to a voltage between 2.5 – 5 V. As long as voltage fluctuations do not exceed the threshold of 2.5 V, the state of the bit will be preserved. However, thresholding requires measurement, which would collapse any superposition in a qubit. Additionally, while classical bits are only subject to bit-flip errors, qubits are subject to decoherence, which is comprised of bit-flips as well as phase-flips that can scramble the phase relationship of a superposition state.

Decoherence is caused by noise processes that originate from uncontrolled degrees of freedom present in the qubit's environment. Types of noise processes can include thermal fluctuations from the control line, amplitude or phase fluctuations from a nearby coupled circuit element, electromagnetic field fluctuations due to defects in the materials used to fabricate the qubit, or thermal fluctuations in nearby dissipative elements that couple to the qubit. These processes result in fluctuations of the probability amplitudes α and β of the qubit states (Eq. (1.1)), eventually resulting in the loss of information as the state evolves into some unknown and unpredictable state. Decoherence is quantified by a characteristic time known as T_2 :

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi} \quad (1.4)$$

Here, T_1 is the energy relaxation time that quantifies the time-scale over which the qubit returns to thermal equilibrium with the environment, and T_ϕ is the dephasing time that

quantifies the extent to which environmental fluctuations cause the transition energy between the $|0\rangle$ and $|1\rangle$ states to fluctuate. In analogy to Eq. (1.2), relaxation leads to a loss of information of the polar angle θ , while dephasing leads to a loss of information of the azimuthal angle ϕ ; by construction decoherence T_2 is the loss of information of the superposition. Importantly, in the limit where no dephasing exists, the decoherence is entirely limited by relaxation; as a result, improving relaxation times is crucial towards improving the ultimate coherence limit in cQED.

Interaction with the environment is in some sense an entangling process; the quantum information becomes lost in the correlations between the qubit and its environment. This can be visualized in the Bloch sphere representation, in Fig. 1.1. A qubit in a closed system with no interaction with its environment can occupy any state corresponding to a point on the surface of the sphere (also known as a pure state), and can traverse the surface of the sphere deterministically through a well-defined gate operation. However, upon interaction with the environment, some of that information becomes lost, and the quantum state corresponds to a point *inside* of the sphere (also known as a mixed state). As a result, the state degrades over time and stochastically evolves into one that was not expected; effectively, the size of the Bloch sphere *shrinks* as $|\alpha|^2 + |\beta|^2 < 1$. Eventually, the Bloch sphere will shrink to a point on the origin; all information is lost to the uncontrolled degrees of freedom in the environment.

The importance of decoherence is highlighted by DiVincenzo's third criterion. The qubits of a good quantum computer must have longer coherence times than the gate operation times. Coherence effectively limits gate fidelity – a measure of how many gates can be applied to a qubit until the evolved state no longer resembles the expected state. This is particularly important for running large quantum algorithms such as those that are required for useful quantum computation. For example, factoring a 768-bit number using Shor's algorithm can require $\sim 10^{11}$ gates performed on $\sim 10^3$ qubits[22]. Coherence and

gate fidelity have improved by orders of magnitude over the last two decades[23]. Currently, the highest single and two-qubit gate fidelities are slightly over 99.99% and 99.9%, respectively[24, 25], which are unfortunately still far too low to faithfully run Shor’s algorithm. However, minimizing sensitivity to noise is not the only way to improve fidelity. Quantum error correction can also be utilized to detect and even correct errors due to unwanted environmental interactions as they occur[26–28]. At the same time, improving intrinsic coherence in quantum circuits improves the performance of error correction protocols[29]. As a result, improvements on both fronts must be made simultaneously in service of the pursuit to realize a useful quantum computer.

1.4 Thesis Overview

In this thesis, I will focus on improving coherence in superconducting quantum circuits. In the cQED platform, sources of decoherence can be attributed to systematic noise in the control pulses required to enact gates, unwanted or excess coupling to the environment due to improper circuit design, and intrinsically dissipative elements or regions in the materials used in circuit fabrication. The first can be alleviated with careful pulse calibration and will not be covered in this thesis. The other two sources of decoherence are more difficult to characterize due to them being difficult to distinguish from each other and will be the focus of this thesis. Obtaining a comprehensive understanding of how materials and fabrication processes limit the coherence of quantum circuits would allow us to understand the current limitations on coherence in state-of-the-art devices. Additionally, this knowledge can be used to optimize devices with regards to circuit design, materials used, and processes employed in order to develop new devices with improved coherence in service of realizing large-scale quantum processors with lower error rates.

I will begin in Ch. 2 by giving a general overview of cQED and describing its fun-

damental elements: quantum harmonic oscillators and Josephson junction-based qubits. I will then describe their interactions with each other and to their control lines, as well as their measurements and experimental implementations. In Ch. 3 I will explain how noise processes result in decoherence in qubits, and describe the participation ratio model that quantifies various energy relaxation mechanisms. I will then discuss the various sources of energy loss in superconducting circuits and how their participation in a quantum circuit are calculated. In Ch. 4 I will describe the main project of my thesis, which is the characterization of microwave losses using multimode microwave resonators. I will lay out the theory of loss extraction through matrix inversion, and describe how errors can be propagated using a least-squares approach. I will then introduce two experimental realizations of this loss characterization approach: the forky whispering-gallery-mode resonator (FWGMR) and the tripole stripline (TSL). I will then discuss in Ch. 5 the usefulness and practicality of loss analysis by applying our newly gained knowledge on intrinsic losses to predict the coherence times of aluminum- and tantalum-based transmon qubits and understand what limits them. I will then demonstrate how this knowledge can be used to improve our current devices by designing compact on-chip quantum memories with greater than millisecond coherence times. Finally, in Ch. 6 I will describe future plans for loss characterization in the flip-chip architecture and the use of new materials that may enable more highly compact, scalable, and versatile quantum circuits with much improved coherence.

1.5 Suggestions for Further Reading

Much of this chapter provided a general introduction to quantum computing, which is explained in greater detail in *Quantum Computation and Quantum Information* by Nielsen and Chuang[8]. Feynman’s keynote, “Simulating Physics with Computers”[1] is also a fascinating read that provides his perspective on quantum computing and quantum simula-

tion. If the reader is curious about quantum parallelism, I would suggest “Quantum-theory, the Church-Turing principle and the universal quantum computer” by David Deutsch[2].

Chapter 2

Circuit Quantum Electrodynamics

In Sec. 1.2 I described cQED as a platform based on superconducting circuits. I introduced the notion of the artificial atom in which a superconducting circuit can be designed to manifest as a two-level system that can encode a qubit. In this chapter, I will describe in detail how a qubit is realized from a superconducting circuit. This involves circuit quantization, where I will start from a basic LC circuit model and show how it can be expressed using a quantum harmonic oscillation Hamiltonian. I will introduce the Josephson junction, the fundamental nonlinear element that allows us to realize a qubit in cQED. I will then discuss how the qubit is controlled with a microwave drive, and how the qubit can be coupled to an oscillator in order to enable either nondestructive readout of the qubit state, or the encoding of a logical qubit in the bosonic states of a linear oscillator. I will finally briefly discuss how to use finite-element simulations to extract Hamiltonian parameters of a quantum circuit, followed by an overview of the various experimental implementations of cQED.

2.1 Superconducting Circuit Quantization

The most basic elements in an electrical circuit are resistors, inductors, and capacitors. In an ideal superconductor at microwave frequencies, the electrical resistance is negligible, effectively resulting in a lossless circuit. In reality, various phenomena cause the circuit to have significantly more resistive loss than predicted, which will be detailed in Ch. 3. However, this loss term is usually very small and its inclusion will not significantly alter the discussion here. Therefore, in this chapter, I will assume the circuits are lossless unless otherwise indicated.

2.1.1 LC Oscillator

The simplest circuit we can make is an LC oscillator (Fig. 2.1a), consisting of an inductor and capacitor in parallel. At the resonance frequency ω_r , energy converts from electrical energy in the capacitor to magnetic energy in the inductor over one oscillation period. When energy dissipation is small, as in the case for a superconducting circuit, the oscillatory behavior can persist for long periods of time, i.e. $T \gg 2\pi/\omega_r$. While oscillators themselves cannot be used for quantum information processing, they are an essential ingredient for doing so. We begin with a classical description of this circuit by writing down the Lagrangian of the circuit in terms of the electrical charge Q stored in the capacitor and the magnetic flux Φ stored in the inductor[30, 31]:

$$\mathcal{L} = \frac{Q^2}{2C} - \frac{\Phi^2}{2L} \quad (2.1)$$

The choice of using the charge basis is made because charge and flux are canonical coordinates, where flux is the canonical position coordinate and charge is the canonical momentum coordinate. This can be seen more clearly when we express the charge in terms

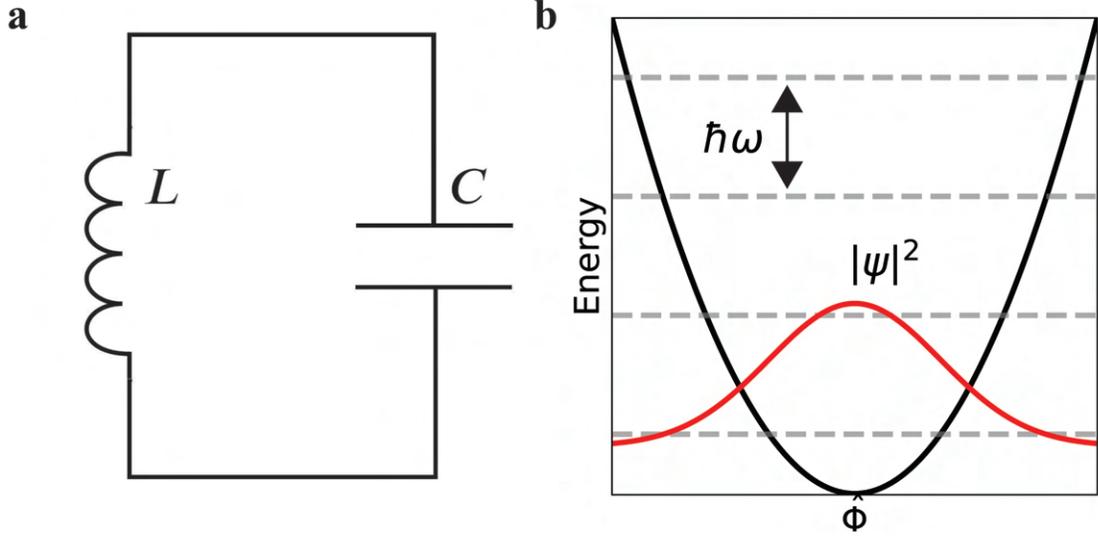


Figure 2.1: **LC oscillator.** **a** Circuit diagram and **b** energy ladder for a quantum LC oscillator. The ground-state wavefunction is shown in red.

of flux, $Q = C\dot{\Phi}$. Comparing flux and charge to the position and momentum of an object attached to a spring, we see that C is analogous to the object's mass and $1/L$ is analogous to the spring constant. Therefore, the potential energy term is the inductive energy and the kinetic energy term is the capacitive energy. This Lagrangian resembles that of a simple harmonic oscillator; as a result, the dynamics are the same. We can derive the Hamiltonian of this system as

$$\mathcal{H} = \dot{\Phi} \frac{\partial \mathcal{L}}{\partial \dot{\Phi}} - \mathcal{L} = \frac{Q^2}{2C} + \frac{\Phi^2}{2L} \quad (2.2)$$

We can now proceed with quantizing the circuit by promoting the charge and flux coordinates to quantum operators with the canonical commutation relation $[\hat{Q}, \hat{\Phi}] = i\hbar$:

$$\hat{Q} = iQ_{ZPF}(\hat{a} - \hat{a}^\dagger) \quad (2.3a)$$

$$\hat{\Phi} = \Phi_{ZPF}(\hat{a} + \hat{a}^\dagger), \quad (2.3b)$$

Here, \hat{a} , \hat{a}^\dagger are the lowering and raising operators, respectively, for the quantum harmonic oscillator, and $\Phi_{ZPF} = \phi_0 \sqrt{\frac{Z}{2\pi R_Q}}$ and $Q_{ZPF} = e \sqrt{\frac{R_Q}{2\pi Z}}$ are the zero-point fluctuations of the flux and charge operators, respectively, where $\phi_0 = h/2e$ and $R_Q = h/2e^2$ are the magnetic flux quantum and resistance quantum, respectively, $Z = \sqrt{L/C}$ is the characteristic impedance of the resonance mode, and e is the charge of the electron. Now we can express the Hamiltonian in terms of the ladder operators,

$$\hat{\mathcal{H}} = \hbar\omega_r \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right), \quad (2.4)$$

where $\omega_r = 1/\sqrt{LC}$. We thus arrive at the Hamiltonian of a simple harmonic oscillator.

The energy diagram for the harmonic oscillator is given in Fig. 2.1b. Due to its quadratic potential, the energy eigenstates are equally spaced by $E = \hbar\omega_r$; we call this a linear system. While this provides at least two quantum states that can exist in superposition, the harmonic oscillator on its own cannot be used as a qubit, because selective transitions between specific energy states cannot occur. If photons of energy $\hbar\omega$ were used to drive the oscillator, transitions between multiple energy levels would occur simultaneously. In order to resolve this issue, we want to design a circuit whose potential energy gives rise to unequal energy spacings between the energy eigenstates. This requires the addition of some form of nonlinearity to the circuit; in cQED, the nonlinearity takes the form of a Josephson junction.

2.1.2 The Josephson Junction

A Josephson junction is a simple superconducting circuit element where a thin (\sim nm) insulating barrier is sandwiched between two superconducting electrodes (Fig. 2.2). Superconducting charge carriers (Cooper pairs) can coherently tunnel through the barrier and

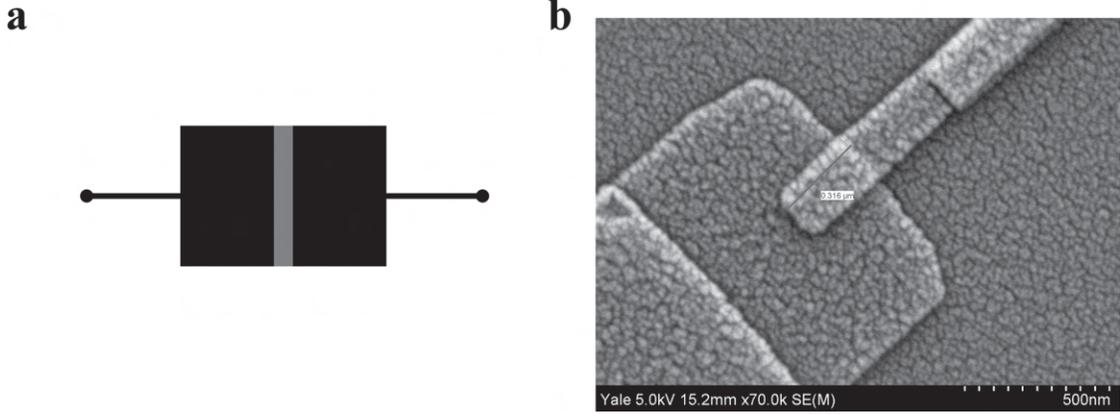


Figure 2.2: **Josephson junction.** **a** Diagrammatic representation of a Josephson junction - a thin insulator (grey) sandwiched between two superconducting electrodes (black). **b** Scanning electron microscopy (SEM) image of an Al/AlO_x/Al junction. The wider aluminum strip is the bottom electrode, the skinnier strip is the top electrode. A thin oxide is buried between the two electrodes.

acquire a phase, described by the Josephson equations:

$$I = I_c \sin \varphi \quad (2.5a)$$

$$\frac{d\varphi}{dt} = \frac{V}{\Phi_0} \quad (2.5b)$$

where I_c is the critical current of the junction, φ is the phase drop across the junction, $\Phi_0 = \hbar/2e$ is the reduced magnetic flux quantum, and I and V are the current through and voltage across the junction, respectively. From Faraday's law, $V = L \frac{dI}{dt}$. Taking the derivative of Eq. (2.5a) and combining it with Eq. (2.5b), we can define a nonlinear inductance of the Josephson junction:

$$L = \frac{L_J}{\cos \varphi} \quad (2.6)$$

where $L_J = \Phi_0/I_c$ is the characteristic Josephson inductance.

The LC circuit in Fig. 2.1a contains a linear inductor. this inductance is geometric; the

inductive energy is stored in the magnetic field. The Josephson junction, on the other hand, has a *kinetic* inductance where the inductive energy is stored in the motion of Cooper pairs. We can compute this energy by integrating $\int IV dt = \int I_c \sin \varphi \Phi_0 \frac{d\varphi}{dt} dt = -E_J \cos \varphi$, where we have defined the Josephson energy $E_J = \Phi_0 I_c = \Phi_0^2 / L_J$, which represents the inductive energy of *one* Cooper pair tunneling through the junction.

While the Josephson equations describe an element that is purely inductive, the physical realization introduces a small capacitance as well. This is due to the finite area of the junction electrodes and the small thickness of the tunnel barrier, which is usually a dielectric with relatively high permittivity. The capacitance can be well-approximated with a parallel-plate approximation, $C_J = \epsilon A / d$, where A is the electrode area, d is the barrier thickness, and ϵ is the barrier's dielectric permittivity. Increasing the area of the electrodes increases the capacitance but also increases the number of conductance channels, decreasing the average kinetic energy of Cooper pairs flowing across the junction, thereby decreasing the inductance.

By far the most common type of Josephson junction used in cQED is an Al/AlO_x/Al junction, with an amorphous aluminum oxide tunnel barrier sandwiched between aluminum electrodes (Fig. 2.2b). As long as the electrodes are superconducting, the junction behaves like a nonlinear inductor; above the superconducting transition temperature, the junction is simply a resistor with normal-state resistance R_n . Conveniently, R_n at low temperature is related to the critical current of the junction by the Ambegaokar-Baratoff relation[32, 33]:

$$I_c = \frac{\pi \Delta_0}{2e R_n^{T \rightarrow 0}} \quad (2.7)$$

Here, $\Delta_0 = \frac{\pi}{e\gamma} k_B T_c$ is the zero-temperature superconducting bandgap (applicable for an ideal superconductor such as aluminum), where $\gamma \approx 0.5772$ is the Euler–Mascheroni con-

stant, k_B is the Boltzmann constant, and T_c is the superconducting transition temperature. The relationship between R_n and I_c provides a simple way to determine L_J by measuring the resistance, a room-temperature parameter. However, to complete this relation, a conversion between room-temperature resistance and low-temperature resistance must be determined. This can be done using physical parameters and is given by the following[34]:

$$R_n^{T \rightarrow 0} = R_n^T \left(1 + \frac{T^2}{\frac{3\hbar^2 \phi_w}{\pi^2 k_B^2 m_e^{eff} d^2}} \right) \quad (2.8)$$

where ϕ_w is the work function of the electrode, m_e^{eff} is the effective electron mass in the insulator, and T is the temperature of the junction upon measurement, which is usually 300 K. By combining Eqs. (2.7) and (2.8), a conversion factor between room-temperature resistance and Josephson inductance can be determined. However, the factor is heavily dependent on fabrication parameters that can determine the barrier thickness and electrode film thickness and T_c . These effects can be combined into process-specific parameters of the junction called the critical current density, J_c , and the specific capacitance, C_{sp} . Junctions for this thesis were fabricated using a process that yields $J_c \approx 50$ A/cm²[35] and $C_{sp} \approx 55$ fF/ μm^2 , and yields a resistance-to-inductance conversion factor of approximately 1.3 nH/k Ω . Typical junction areas used for qubits range from 0.03 – 0.3 μm^2 , giving $C_J = 1.6 - 16$ fF.

2.1.3 The Cooper-Pair Box

It is intuitive to define the Hamiltonian of the Josephson junction in the phase basis rather than the charge basis. By integrating Eq. (2.5b), we see that the phase basis is equivalent to the flux basis, $\hat{\varphi} = \hat{\Phi}/\Phi_0$. This allows us to regard the energy in the junction as the potential energy term of the Hamiltonian, and the motion of Cooper pairs through the junction as the kinetic energy term. We can therefore describe the tunneling of a single

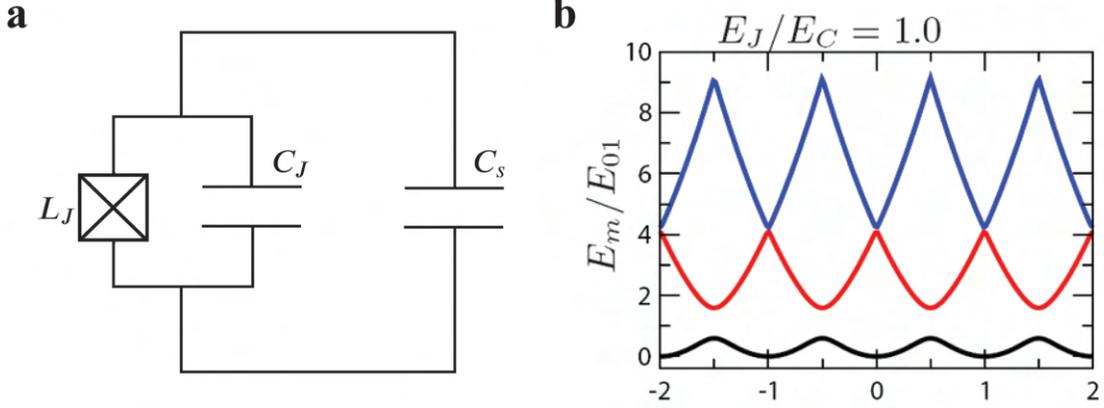


Figure 2.3: **Cooper-Pair Box.** **a** Circuit diagram of a Josephson Junction shunted by a capacitor C_s . The insulator being sandwiched by two electrodes gives rise to a parallel plate capacitance C_J , resulting in an equivalent capacitance for the circuit $C_q = C_J + C_s$. **b** Energy diagram for a CPB for $E_J/E_C = 1$ for the first three levels ($m = 0, 1, 2$ are black, red, blue) as a function of n_g . Figure adapted from Koch et al. [36].

Cooper pair of charge $2e$ across the junction with the following Cooper-pair box (CPB) Hamiltonian, derived from Fig. 2.3a:

$$\hat{\mathcal{H}} = 4E_C(\hat{n} - n_g)^2 - E_J \cos \hat{\Phi}/\Phi_0 \quad (2.9)$$

where $E_C = e^2/2C_q$ is the charging energy due to one electron's worth of charge accumulated in an equivalent capacitor with capacitance C_q , $\hat{n} = \hat{Q}/2e$ is the charge number operator corresponding to one Cooper pair being transferred across the junction, and n_g is a continuous variable that corresponds to an offset charge between the junction electrodes. Compared to the harmonic oscillator Hamiltonian, the quadratic potential is replaced by a cosine potential, which breaks linearity and results in unequal energy spacings between eigenstates. If we are describing the Josephson junction by itself, the charging energy would be entirely determined by the junction capacitance, $C_q = C_J$. However, adding additional shunting capacitance C_s as in Fig. 2.3a modifies the equivalent capacitance to $C_q = C_J + C_s$. A large C_s will be important towards the realization of the transmon qubit in Sec. 2.1.4.

The CPB Hamiltonian contains the necessary nonlinearity to realize a qubit through its energy eigenstates and is well-described in David Schuster’s thesis[37]. In the limit where $E_J \sim E_C$, the energy $\hbar\omega_{ge} = \sqrt{[4E_C(1 - n_g)]^2 + E_J^2}$ [38] between the ground state $|g\rangle$ and first excited state $|e\rangle$ is different from that of all other states when $n_g = 1$, as can be seen in Fig. 2.3b. A microwave drive at frequency ω_{ge} can be applied to *selectively* drive transitions between the ground and first excited state without causing leakage to higher excited states. Effectively, this creates a two-level system; hence, a qubit.

The Cooper-pair box has been demonstrated as a qubit and shown to have high relaxation times of up to $T_1 = 200 \mu\text{s}$ [38]. However, coherence times are typically at most $1 \mu\text{s}$, dominated by dephasing[36]. This is due to stochastic fluctuations in offset charge n_g , caused by locally fluctuating electric fields. This results in fluctuations in ω_{ge} , which dephase the qubit (Fig. 2.3b). To suppress the effects of charge noise, a circuit modification can be made. By increasing the shunting capacitance C_s , the charging energy decreases, thereby decreasing the amplitude of n_g fluctuations. This serves as the basis for the transmon qubit, which is described by the same Hamiltonian as in Eq. (2.9) but in the regime where $E_J \gg E_C$.

2.1.4 The Transmon Qubit

To see how $E_J \gg E_C$ changes the dynamics of the CPB, it is helpful to express the zero-point fluctuations of \hat{n} and $\hat{\Phi}$ in terms of E_J and E_C :

$$n_{ZPF} = \left(\frac{E_J}{32E_C} \right)^{1/4} \quad (2.10a)$$

$$\Phi_{ZPF} = \Phi_0 \left(\frac{2E_C}{E_J} \right)^{1/4}, \quad (2.10b)$$

For $E_J \gg E_C$, $n_{ZPF} \gg n_g$. Achieving this regime generally requires the addition of a large shunting capacitor $C_s \gg C_J$, which further dampens the amplitude of n_g fluctuations. This motivates us to drop the n_g term from the Hamiltonian. The charge dispersion of the m th state due to an offset charge fluctuation of ± 1 is given by[36]

$$\epsilon_m \simeq (-1)^m E_C \frac{2^{4m+5}}{m!} \sqrt{\frac{2}{\pi}} \left(\frac{E_J}{2E_C} \right)^{\frac{m}{2} + \frac{3}{4}} e^{-\sqrt{8E_J/E_C}}, \quad (2.11)$$

which shows that increasing E_J/E_C exponentially dampens the amplitude frequency fluctuations due to charge noise. The transmon qubit is designed to be in the regime where $E_J/E_C > 50$ (Fig. 2.4c), where for typical parameters the estimated dephasing time due to charge noise can approach several tens of milliseconds to seconds[36] while maintaining relaxation times in excess of 400 μ s[39]

In the transmon limit, the zero-point fluctuations in flux become very small; therefore, we can expand the cosine potential in the CPB Hamiltonian up to 4th order and still maintain a good approximation for the system:

$$\hat{\mathcal{H}} \approx 4E_C \hat{n}^2 - E_J \left(1 - \frac{1}{2} \left(\frac{\hat{\Phi}}{\Phi_0} \right)^2 + \frac{1}{24} \left(\frac{\hat{\Phi}}{\Phi_0} \right)^4 \right) \quad (2.12)$$

The first term of the cosine expansion is a constant that does not affect the equations of motion and can be dropped from the Hamiltonian. The second term reproduces the quadratic potential that resembles that of the LC oscillator (Eq. (2.2)). The final term is a nonlinear term that gives rise to the useful dynamics required for operating the transmon as a qubit. Notably, the nonlinear term is small and can be treated as a perturbation to the linear oscillator (Fig. 2.4a). We can therefore express the transmon Hamiltonian in terms of the harmonic oscillator ladder operators:

$$\hat{\mathcal{H}} \approx \sqrt{8E_J E_C} \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right) - \frac{1}{12} E_C (\hat{a} + \hat{a}^\dagger)^4 \quad (2.13)$$

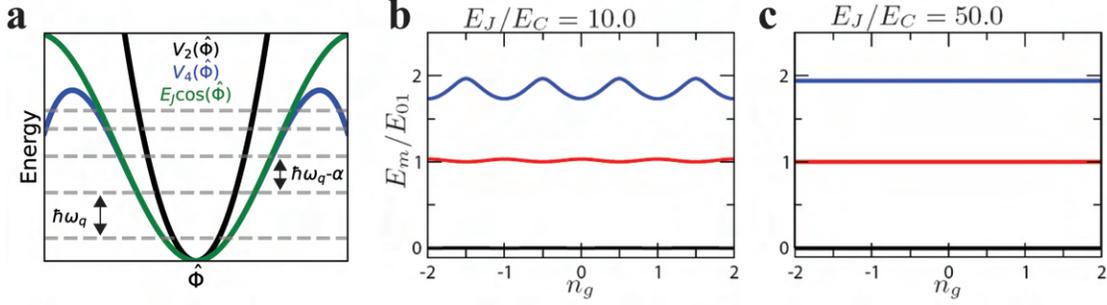


Figure 2.4: **Transmon qubit energy spectrum.** **a** Energy diagram for the transmon. The cosine potential (green) is not well-approximated by the quadratic potential (black). Expanding the cosine potential to the 4th order (blue) provides a good approximation. **b** Energy diagram and charge dispersion for a CPB with $E_J/E_C = 10.0$. Compared to Fig. 2.3b, the eigenenergies are less dependent on n_g . **c** Energy diagram and charge dispersion for a transmon. The eigenenergies are almost completely independent of n_g . Figure adapted from Koch et al. [36].

By expanding the 4th order term and keeping only the energy-conserving terms, we arrive (after simplification) at

$$\begin{aligned}
\hat{\mathcal{H}} &\approx \sqrt{8E_J E_C} \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right) - \frac{1}{2} E_C (\hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a} + 2\hat{a}^\dagger \hat{a}) \\
&= \hbar\omega_q \hat{a}^\dagger \hat{a} - \frac{1}{2} E_C \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a}
\end{aligned} \tag{2.14}$$

where $\omega_q = \omega_p - E_C/\hbar$ is the transition frequency between the ground state $|g\rangle$ and first excited state $|e\rangle$ of the transmon. One effect of the nonlinear term $\hat{\mathcal{H}}_{nl} = -\frac{1}{12} E_C (\hat{a} + \hat{a}^\dagger)^4$ is the addition of a Lamb shift $-E_C/\hbar$ to the plasma frequency $\omega_p = \sqrt{8E_J E_C}/\hbar$ [40]. The other effect is the introduction of anharmonicity to the harmonic oscillator potential. We see that $\hat{\mathcal{H}}_{nl} |k\rangle = 0$ for $k = 0, 1$, i.e. the ground and first excited states, and for $k > 1$, $\hat{\mathcal{H}}_{nl} |k\rangle = \frac{1}{2} E_C (k^2 - k)$. The state-dependent transition energy then becomes $E_{k+1} - E_k = \hbar\omega_q - kE_C$. The slight modification of the harmonic oscillator potential introduces a Kerr-type nonlinearity, where the anharmonicity $\alpha = -E_C/\hbar$ allows selective transitions between the states of the transmon.

Typical transmons used in cQED have frequencies $\omega_q/2\pi = 3 - 7$ GHz and $E_C/h = 100 - 300$ MHz. The anharmonicities are actually quite small, around an order of magnitude smaller than the $g - e$ transition frequency, highlighting that the transmon is in reality simply a weakly anharmonic oscillator. The small anharmonicity limits the effective bandwidth of microwave pulses used to control the transmon; if the bandwidth exceeds the anharmonicity, higher states can be excited, limiting the fidelity of gate operations in the $g - e$ manifold. On the other hand, the existence of these higher states, when properly considered during system design, can be used to implement operations such as error mitigation and error detection. Nevertheless, if we restrict the microwave control pulses on the qubit to have a frequency bandwidth that does not exceed the anharmonicity, we can treat the transmon like a two-level system. The transmon anharmonicity is indeed large enough to achieve this while retaining the ability to implement fast gate operations on timescales of tens of nanoseconds. The simplicity, robustness to charge noise, and ability to implement fast gates has made the transmon the most popular implementation of an artificial atom to date, and will be the main nonlinear element used in this thesis.

2.2 Classical Drives on Superconducting Quantum Circuits

To perform useful operations on qubits and quantum circuits, we must be able to control them. In cQED, circuits are controlled using classical microwave drives, pulse sequences driven at energy scales far higher than the single-photon regime. This is to ensure that the drives remain in a stiff-pump limit; the number of photons absorbed is negligible compared to the total number of drive photons. While many control schemes exist, this thesis will focus on driving a quantum circuit through a drive line that couples to the flux operator $\hat{\Phi} \propto (\hat{a}^\dagger + \hat{a})$ (the drive can also be coupled to \hat{Q} , but this doesn't change the

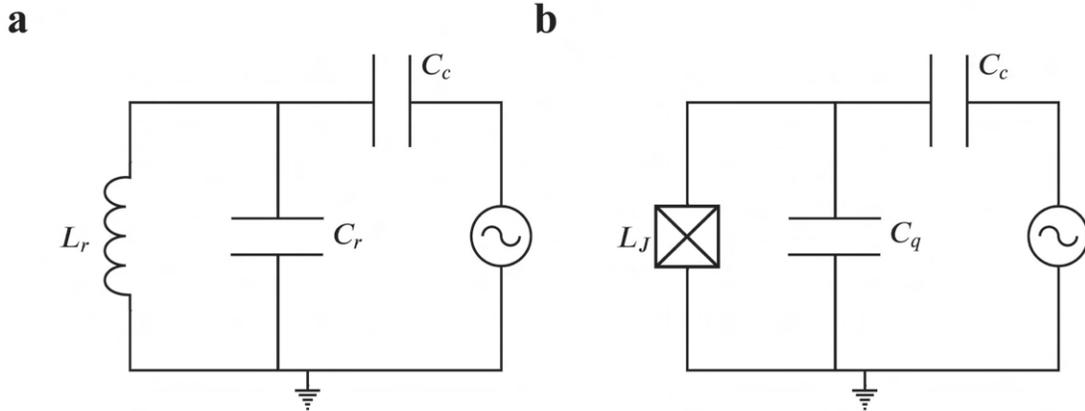


Figure 2.5: **Driving an LC oscillator and transmon.** **a** Circuit diagram of a driven LC oscillator; an AC voltage source is capacitively coupled to the circuit. **b** Circuit diagram of a driven transmon; an AC voltage source is capacitively coupled to the circuit.

overall dynamics aside from the addition of a phase). The physical realization of this drive is a transmission line that is capacitively coupled to the circuit (Fig. 2.5). In order to limit energy leakage through the drive line (external coupling), the effective coupling capacitance C_c is small enough such that its dissipation from external sources is much smaller than its internal dissipation (i.e. dissipation due to sources within the materials that make up the circuit). The drive term Hamiltonian has the form

$$\hat{\mathcal{H}}_d = \hbar\epsilon(t)(\hat{a}^\dagger + \hat{a}) \quad (2.15)$$

where $\epsilon(t)$ represents the time-varying fields of the drive. Here, we will discuss how this drives affect the dynamics of linear oscillators and transmon qubits.

2.2.1 Driving a Linear Oscillator

Since the linear oscillator has equally-spaced energy levels, driving it excites all energy eigenstates. As a result, the drive can only create and induce transitions between coherent

states. The full Hamiltonian of the system is

$$\hat{\mathcal{H}} = \hbar\omega_r \hat{a}^\dagger \hat{a} + \hbar\epsilon(t)(\hat{a}^\dagger + \hat{a}) \quad (2.16)$$

We make a unitary transformation $\hat{\mathcal{H}} = \hat{U}\hat{\mathcal{H}}\hat{U}^\dagger - i\hbar\hat{U}\frac{d}{dt}\hat{U}^\dagger$ and choose $\hat{U} = e^{i\omega_r t \hat{a}^\dagger \hat{a}}$ to move into the rotating frame of the oscillator:

$$\hat{\mathcal{H}} = \hbar\epsilon(t)(\hat{a}^\dagger + \hat{a}) \quad (2.17)$$

where $\hat{a} = \hat{a}e^{-i\omega_r t}$ is stationary in this frame. We can then write down the time-evolution unitary propagator[41]:

$$\hat{U}(t) = e^{-i\frac{1}{\hbar}\int_0^t dt' \hat{\mathcal{H}}(t')} = e^{\beta\hat{a}^\dagger - \beta^*\hat{a}} \equiv \hat{D}(\beta) \quad (2.18)$$

where we define \hat{D}_β to be the displacement operator with displacement amplitude $\beta = -i\int_0^t \epsilon(t')dt' e^{i\omega_r t}$. We can use the Baker-Campbell-Hausdorff formula to rearrange the operator to $\hat{D}(\beta) = e^{-|\beta|^2/2} e^{\beta\hat{a}^\dagger} e^{-\beta^*\hat{a}}$, and operate on the vacuum state to produce a coherent state $|\beta\rangle$:

$$\hat{D}(\beta)|0\rangle = e^{-|\beta|^2/2} \sum_n \frac{\beta^n}{\sqrt{n!}} |n\rangle = |\beta\rangle \quad (2.19)$$

Coherent states are eigenstates of \hat{a} with eigenvalue β , with mean photon number of the state being $\bar{n} = \langle\beta|\hat{a}^\dagger\hat{a}|\beta\rangle = |\beta|^2$. It can further be shown by calculating the expectation value of the canonical position operator that the dynamics of a coherent state resemble very closely the dynamics of a classical simple harmonic oscillator. As a result, non-classical states such as photon-number states require some nonlinearity to prepare and manipulate.

2.2.2 Driving a Transmon Qubit

The nonlinearity of the transmon enables a classical drive to prepare and manipulate a variety of quantum states. We can think of the drive as inducing an amplitude-dependent phase drop across the junction. The full Hamiltonian of the system is

$$\hat{\mathcal{H}} = \hbar\omega_q \hat{a}^\dagger \hat{a} - \frac{1}{2} E_C \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a} + \hbar\epsilon(t)(\hat{a}^\dagger + \hat{a}) \quad (2.20)$$

we will use this Hamiltonian to describe the transmon's behavior subject to on- and off-resonant drives.

Resonant Drive

If we restrict the computational subspace to the lowest two energy levels $|g\rangle$ and $|e\rangle$, we can treat the system as a spin-1/2 system and replace \hat{a} and \hat{a}^\dagger with the ladder operators $\hat{\sigma}_\pm = (\hat{\sigma}_x \pm i\hat{\sigma}_y)/2$, where $\hat{\sigma}_{x,y,z}$ are the Pauli matrices[42]:

$$\hat{\mathcal{H}} = -\frac{\hbar\omega_q}{2} \hat{\sigma}_z + \hbar\epsilon(t)\hat{\sigma}_x \quad (2.21)$$

We then go into the rotating frame of the qubit with the unitary $\hat{U} = e^{-i\frac{\omega_q}{2}t\hat{\sigma}_z}$. The transformed Hamiltonian, after some algebra, is:

$$\hat{\mathcal{H}} = \hbar\epsilon(t)(\hat{\sigma}_x \cos(\omega_q t) + \hat{\sigma}_y \sin(\omega_q t)) \quad (2.22)$$

We assume the drive has form $\epsilon(t) = \Omega_d \cos(\omega_d t + \phi)$, where Ω_d is the drive amplitude, ω_d is the drive frequency, and ϕ is the drive's phase. This simplifies $\hat{\mathcal{H}}$ to

$$\hat{\mathcal{H}} = \frac{\hbar\Omega_d}{2} (\hat{\sigma}_x \cos(\Delta t - \phi) + \hat{\sigma}_y \sin(\Delta t - \phi)) \quad (2.23)$$

after invoking the rotating wave approximation (RWA) to drop fast-rotating terms. So far, we have kept the drive at arbitrary frequency; the detuning $\Delta = \omega_q - \omega_d$ is nonzero. If we now drive resonantly, $\Delta = 0$, and the drive term becomes

$$\hat{\mathcal{H}} = \frac{\hbar\Omega_d}{2}(\hat{\sigma}_x \cos(\phi) - \hat{\sigma}_y \sin(\phi)) \quad (2.24)$$

Therefore, by controlling the phase of the drive, qubit rotations about an arbitrary axis on the $x - y$ plane of the Bloch sphere (defined for the $g - e$ computational space) can be performed. By choosing $\phi = 0$, for example, $\hat{\mathcal{H}} = \frac{\hbar\Omega_d}{2}\hat{\sigma}_x$, and the drive performs rotations about the x -axis. To see the time evolution of x -axis rotations, we can define the time-evolution unitary:

$$\hat{U} = e^{i\frac{\Omega_d}{2}t\hat{\sigma}_x} = \cos\left(\frac{\Omega_d}{2}t\right) + i\hat{\sigma}_x \sin\left(\frac{\Omega_d}{2}t\right) \quad (2.25)$$

which upon operating on $|g\rangle$ we obtain:

$$\hat{U}|g\rangle = \cos\left(\frac{\Omega_d}{2}t\right)|g\rangle + i\sin\left(\frac{\Omega_d}{2}t\right)|e\rangle \quad (2.26)$$

From this, we see that the qubit state undergoes Rabi oscillations between $|g\rangle$ and $|e\rangle$ (Fig. 2.6a). Continuous driving is known as Rabi driving and $\Omega_{Rabi} = \Omega_d$ is known as the Rabi frequency. Since Ω_d is a measure of drive amplitude, the Rabi frequency is linear in drive amplitude (Fig. 2.6b). Driving the qubit from $|g\rangle$ to $|e\rangle$ requires a X_π -pulse, in which the drive is applied for time $T_\pi = 2\pi/\Omega_d$; driving for an intermediate amount of time leaves the qubit in a superposition state. Rotations around the z -axis can be easily done virtually by applying a relative phase ϕ to the drive between two σ_x or σ_y rotations, which effectively rotates the qubit state about the z -axis by ϕ between the two $\sigma_{x,y}$ rotations. Therefore, this drive can be used to apply universal single-qubit gate

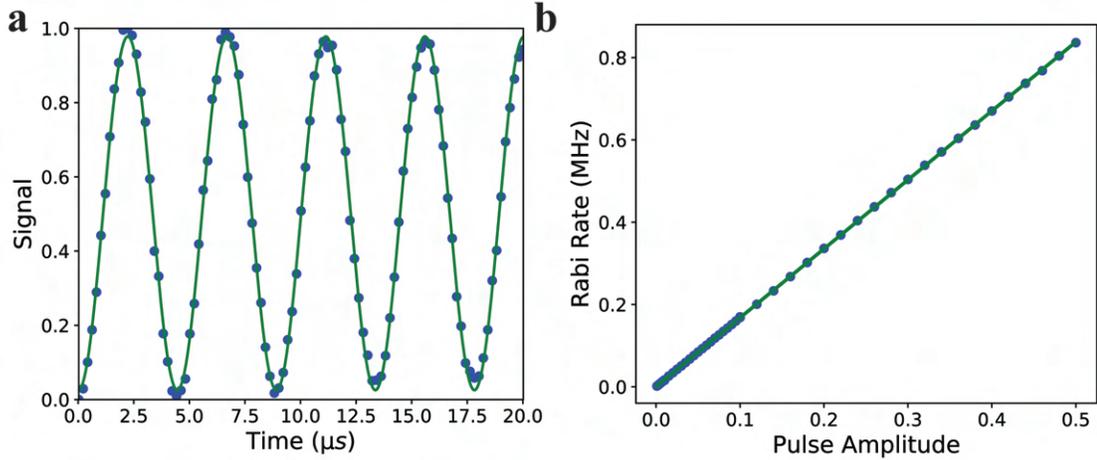


Figure 2.6: **Transmon Rabi oscillations.** **a** Oscillation of the qubit between $|g\rangle$ (Signal = 0) and $|e\rangle$ (Signal = 1) when driven continuously at $\Omega_d = \omega_q$. **b** Rabi frequency scales linearly with drive amplitude.

operations, i.e. arbitrary rotations about the Bloch sphere.

The treatment of the Rabi drive assumes a continuous drive; turning the drive “on” and “off” implies the application of an envelope to the drive, which would make Ω_d explicitly time-dependent. The time-evolution unitary would then become $\hat{U} = e^{i\hat{\sigma}_x/2 \int_0^t \Omega_d(t') dt'}$. This would naturally affect the frequency content of the drive by introducing off-resonant frequency terms that may cause unwanted transitions between higher qubit states outside the computational subspace. To alleviate this, a Gaussian envelope is used, due to its off-resonant frequency terms being damped exponentially. However, attempting to apply faster qubit rotations results in an increased frequency bandwidth for the pulse, resulting in a small drive term that is resonant with transitions to the second excited state $|f\rangle$. Fortunately, pulse engineering methods exist to reduce this leakage to higher states[43, 44].

Fast qubit pulses have an additional drawback in that they have reduced frequency selectivity. The faster the pulse, the larger the bandwidth; as a result, the qubit can be excited even with slightly off-resonant drives. This can pose a problem when attempting to apply photon number-selective pulses to control a coupled resonator state, as will be

discussed in Sec. 2.3.3. For such operations, the pulses have to be deliberately engineered to be slower in order to reduce the bandwidth enough for the pulse to be selective.

Off-Resonant Drive

Driving the qubit off-resonantly results in a frequency shift of the transmon due to the AC Stark effect. This shift, called the Stark shift, also occurs during qubit gate operations, as strong enough drives that are resonant with the $g - e$ transition are off-resonant with transitions between higher states; therefore, their transition frequencies will shift[45]. We begin with adding a drive to the Hamiltonian of Eq. (2.13):

$$\hat{\mathcal{H}} = \hbar\omega_p \hat{a}^\dagger \hat{a} - \frac{1}{12} E_C (\hat{a} + \hat{a}^\dagger)^4 + \hbar(\epsilon(t) + \epsilon^*(t))(\hat{a}^\dagger + \hat{a}) \quad (2.27)$$

Here, we have dropped the constant terms from the Hamiltonian, and redefine $\epsilon(t) = \frac{\Omega_d}{2} e^{-i\omega_d t}$. We then move into the rotating frame of the transmon:

$$\hat{\mathcal{H}}' = -\frac{1}{12} E_C (\hat{a} + \hat{a}^\dagger)^4 + \hbar(\epsilon(t) + \epsilon^*(t))(\hat{a}^\dagger + \hat{a}) \quad (2.28)$$

where $\hat{a} = \hat{a} e^{-i\omega_p t}$. We then move into the displaced frame of the drive using the unitary $\hat{U} = e^{\xi \hat{a}^\dagger - \xi^* \hat{a}}$ for displacement ξ . Using the Baker-Campbell-Hausdorff formula, this transforms \hat{a} as $\hat{a} \rightarrow \hat{U} \hat{a} \hat{U}^\dagger = \hat{a} - \xi$ and leads to the transformed Hamiltonian:

$$\begin{aligned} \hat{\mathcal{H}} = & -\frac{1}{12} E_C (\hat{a} + \hat{a}^\dagger - \xi - \xi^*)^4 + \hbar(\epsilon(t) + \epsilon^*(t))(\hat{a}^\dagger + \hat{a} - \xi - \xi^*) \\ & - i\hbar(\dot{\xi} \hat{a}^\dagger + \xi \dot{\hat{a}}^\dagger - \dot{\xi}^* \hat{a} - \xi^* \dot{\hat{a}}) \end{aligned} \quad (2.29)$$

The effect of the displacement drive is to remove the drive term from the transformed Hamiltonian. There is a value of ξ that accomplishes this. We expand the drive term and

invoke the RWA to drop fast-rotating terms. Then the solution for the value of ξ is given by solving the differential equation

$$\dot{\xi} = -i\epsilon(t) - \left(\frac{\kappa}{2} + i\omega_p\right) \xi \quad (2.30)$$

where κ is a small energy decay rate due to the capacitive coupling to the 50Ω drive line that has been so far neglected. The solution for ξ is

$$\xi(t) = \frac{i(\Omega_d/2)e^{-i\omega_d t}}{\kappa/2 + i(\omega_p - \omega_d)} \quad (2.31)$$

The transformed Hamiltonian is therefore rewritten as

$$\hat{\mathcal{H}} = -\frac{1}{12}E_C(\hat{a} + \hat{a}^\dagger - \xi - \xi^*)^4 \quad (2.32)$$

which, after expanding the term and dropping all constants and rapidly rotating terms, yields

$$\hat{\mathcal{H}} = -\frac{1}{2}E_C(\hat{a}^\dagger\hat{a}^\dagger\hat{a}\hat{a} + 2\hat{a}^\dagger\hat{a} + 4\hat{a}^\dagger\hat{a}|\xi|^2) \quad (2.33)$$

where the second term in the parenthesis is the previously discussed Lamb shift to the transmon plasma frequency, and the third term is the AC Stark shift[31]:

$$\Delta_{Stark} = -2E_C\hat{a}^\dagger\hat{a}\frac{\Omega_d^2}{\kappa^2 + (\omega_p - \omega_d)^2} \quad (2.34)$$

2.3 Coupling a Transmon to a Linear Oscillator

In the previous section I discussed how we can universally control a single transmon qubit using a classical drive. However, entangling operations require the qubit to be coupled to

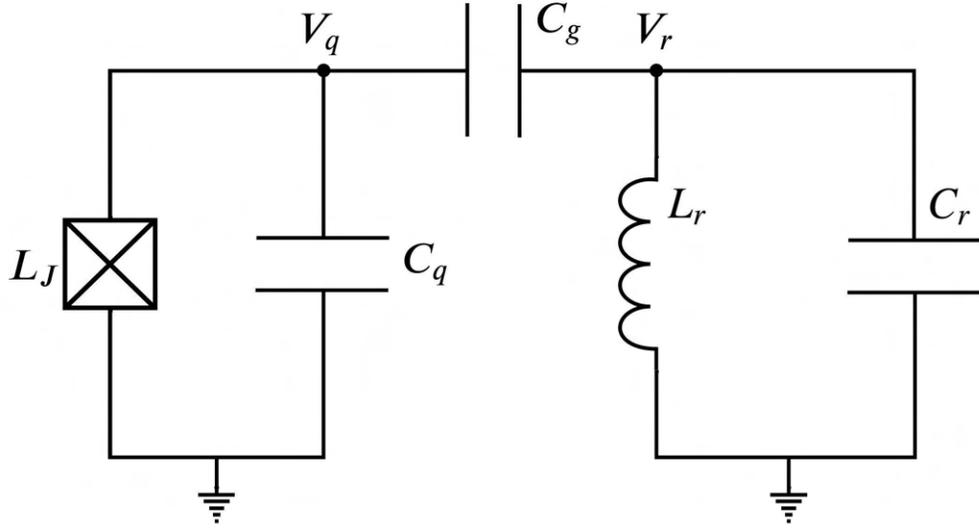


Figure 2.7: **Transmon coupled to an LC resonator.**

other circuits. Coupling a qubit to another qubit enables two-qubit operations, but in this thesis I will focus on the coupling of a transmon qubit to a linear oscillator. This enables two main functions: 1) to perform non-destructive readout of the qubit state, and 2) to encode a logical qubit in the bosonic states of the oscillator; such a qubit is also known as a quantum memory. Both of these functions begin with the same general circuit model and Hamiltonian, where a transmon is capacitively coupled to an LC resonator (Fig. 2.7).

The circuit Hamiltonian can be described in parts, and then combined through an interaction term. It is helpful to first analyze the circuit in the voltage basis, where the capacitive energy is $\frac{1}{2}CV^2$. By defining voltage nodes V_q and V_r , as in Fig. 2.7, the Hamiltonians of the transmon, resonator, and transmon-resonator interaction can be straightforwardly written separately and then combined[42]:

$$\begin{aligned}
\hat{\mathcal{H}}_q &= \frac{1}{2}C_q\hat{V}_q^2 - E_J \cos \hat{\varphi}_q \\
\hat{\mathcal{H}}_r &= \frac{1}{2}C_r\hat{V}_r^2 + \frac{\hat{\Phi}_r^2}{2L_r} \\
\hat{\mathcal{H}}_{int} &= \frac{1}{2}C_g(\hat{V}_q - \hat{V}_r)^2 = \frac{1}{2}C_g(\hat{V}_q^2 - 2\hat{V}_q\hat{V}_r + \hat{V}_r^2) \\
\hat{\mathcal{H}} &= \frac{1}{2}(C_q + C_g)\hat{V}_q^2 - E_J \cos \hat{\varphi}_q + \frac{1}{2}(C_r + C_g)\hat{V}_r^2 + \frac{\hat{\Phi}_r^2}{2L_r} - C_g\hat{V}_q\hat{V}_r \\
\hat{\mathcal{H}} &= 4E_{C_{q,\Sigma}}\hat{n}_q^2 - E_J \cos \hat{\varphi}_q + 4E_{C_{r,\Sigma}}\hat{n}_r^2 + \frac{\hat{\Phi}_r^2}{2L_r} - 4e^2\frac{C_g}{C_qC_r}\hat{n}_q\hat{n}_r \quad (2.35)
\end{aligned}$$

where the substitution $\hat{V}_i = 2e\hat{n}_i/C_i$ was made in the last line, $E_{C_{q,\Sigma}} = \frac{e^2}{2}\frac{C_q+C_g}{C_q^2}$, and $E_{C_{r,\Sigma}} = \frac{e^2}{2}\frac{C_r+C_g}{C_r^2}$. In practical applications, we take the weak coupling limit where $C_g \ll C_q, C_r$; therefore, $E_{C_{i,\Sigma}} \approx E_{C_i} = \frac{e^2}{2C_i}$. We can now express the Hamiltonian in terms of the harmonic oscillator ladder operators[42]:

$$\hat{\mathcal{H}} = \hbar\omega_q\hat{q}^\dagger\hat{q} + \hbar\omega_r\hat{r}^\dagger\hat{r} - \frac{E_C}{2}\hat{q}^\dagger\hat{q}^\dagger\hat{q}\hat{q} - \hbar g(\hat{q} - \hat{q}^\dagger)(\hat{r} - \hat{r}^\dagger) \quad (2.36)$$

where \hat{q} and \hat{r} are the transmon and resonator ladder operators, respectively, ω_r is the resonator frequency, and the cosine potential has been expanded out to 4th order. The transmon-resonator coupling constant g is defined as[40]

$$g \equiv -\omega_r\frac{C_g}{C_q}\left(\frac{E_J}{2E_C}\right)^{1/4}\sqrt{2\pi\alpha_0\frac{Z_r}{Z_{\text{vac}}}} \quad (2.37)$$

where $Z_r = \sqrt{L_r/C_r}$ is the characteristic impedance of the oscillator mode, Z_{vac} is the vacuum impedance, and α_0 is the fine-structure constant[46]. The presence of the fine-structure constant reflects the nature of cQED in that interactions occur between electromagnetic fields. Excitations of the transmon field interact with the resonator field and vice-versa, producing the behavior that we observe. In analogy with cavity QED, the

dipole moment of an atom couples to the electromagnetic fields of a microwave cavity resonator; however, since the atom is so small, the coupling strength is often very weak. In cQED, while $C_g \ll C_q$ in most cases, the coupling g can still be made quite large by increasing Z_r and E_J/E_C ; either can be accomplished by tuning circuit parameters. However, increasing E_J/E_C comes at the expense of reducing the anharmonicity. Typical values for g are around $|g|/2\pi \sim 100 \text{ MHz} \ll \omega_q, \omega_r$.

In the small-coupling limit, we can simplify Eq. (2.36) by invoking the RWA and eliminate $\hat{q}\hat{r}$ and $\hat{q}^\dagger\hat{r}^\dagger$ that rotate quickly enough to average to zero[42]:

$$\hat{\mathcal{H}} = \hbar\omega_q\hat{q}^\dagger\hat{q} + \hbar\omega_r\hat{r}^\dagger\hat{r} - \frac{E_C}{2}\hat{q}^\dagger\hat{q}^\dagger\hat{q}\hat{q} + \hbar g(\hat{q}^\dagger\hat{r} + \hat{q}\hat{r}^\dagger) \quad (2.38)$$

By restricting the computational subspace of the transmon to the two lowest energy states, \hat{q} and \hat{q}^\dagger reduce to $\hat{\sigma}_-$ and $\hat{\sigma}_+$, respectively, and the Hamiltonian becomes

$$\hat{\mathcal{H}} = -\frac{\hbar\omega_q}{2}\hat{\sigma}_z + \hbar\omega_r\hat{r}^\dagger\hat{r} + \hbar g(\hat{\sigma}_+\hat{r} + \hat{\sigma}_-\hat{r}^\dagger) \quad (2.39)$$

This is the Jaynes-Cummings Hamiltonian[19, 40, 47], which has been widely used in cavity QED to describe the interaction between a two-level atom and the resonant electromagnetic fields of the microwave cavity resonator inside which it resides. the atom's dipole moment interacts with the cavity field, resulting in the coherent exchange of excitations between the atom and the resonator. Since the atom is incredibly small compared to the cavity, the dipole moment is also small; therefore, the interaction strength tends to be very weak. In analogy with Eq. (2.37), the size of the transmon's "dipole moment" is proportional to $(E_J/2E_C)^{1/4}$, and the resonator's electromagnetic field strength is related to its characteristic impedance. We have thus reproduced the Jaynes-Cummings interaction within the field of cQED, but here we have the additional advantage of easily tuning the parameters that determine the strength of the dipole moment and resonator field. Inceas-

ing Z_r increases the strength of the electromagnetic field fluctuations; increasing the ratio E_J/E_C increases the size of the transmon's "dipole moment". As a result, much higher interaction strengths can be achieved in cQED, allowing us to access the strong dispersive limit[37], which is integral to both qubit readout and logical encoding of information in an oscillator.

2.3.1 The Dispersive Limit

In the regime of strong coupling or where $\omega_q = \omega_r$, the transmon is highly entangled with the resonator; as a result, the state is not a pure state within the qubit's computational subspace. It is therefore useful to work in the dispersive limit, where the qubit-resonator detuning $\Delta = \omega_r - \omega_q$ is much larger than the coupling, $|g| \ll \Delta$. In this regime, the large detuning prevents coherent exchange of excitations; interactions only occur through virtual photon exchange. These processes also involve higher energy states of the transmon; therefore, the two level approximation of the transmon is insufficient to describe the interaction. Instead, we use Eq. (2.38) and separate the linear and nonlinear parts:

$$\begin{aligned}\hat{\mathcal{H}}_l &= \hbar\omega_q\hat{q}^\dagger\hat{q} + \hbar\omega_r\hat{r}^\dagger\hat{r} + \hbar g(\hat{q}^\dagger\hat{r} + \hat{q}\hat{r}^\dagger) \\ \hat{\mathcal{H}}_{nl} &= -\frac{E_C}{2}\hat{q}^\dagger\hat{q}^\dagger\hat{q}\hat{q}\end{aligned}\quad (2.40)$$

and use the Bogoliubov transformation[40] to diagonalize $\hat{\mathcal{H}}_l$ with the unitary $\hat{U} = e^{\delta(\hat{r}^\dagger\hat{q} - \hat{r}\hat{q}^\dagger)}$. Using the Baker-Campbell-Hausdorff formula, this transforms \hat{r} and \hat{q} as $\hat{U}\hat{r}\hat{U}^\dagger = \cos(\delta)\hat{r} - \sin(\delta)\hat{q}$ and $\hat{U}\hat{q}\hat{U}^\dagger = \cos(\delta)\hat{q} + \sin(\delta)\hat{r}$. The transformed $\hat{\mathcal{H}}_l$ is then[40]

$$\hat{\mathcal{H}}_l = \hat{U}\hat{\mathcal{H}}_l\hat{U}^\dagger = \hbar\tilde{\omega}_q\hat{q}^\dagger\hat{q} + \hbar\tilde{\omega}_r\hat{r}^\dagger\hat{r} + \hbar\left[g\cos(2\delta) - \frac{1}{2}\Delta\sin(2\delta)\right](\hat{q}^\dagger\hat{r} + \hat{q}\hat{r}^\dagger)\quad (2.41)$$

where $\Delta = \omega_r - \omega_q$ is the transmon-resonator detuning, and $\tilde{\omega}_q$ and $\tilde{\omega}_r$ are the renormalized transmon and resonator frequencies, respectively, given by

$$\tilde{\omega}_q = \frac{1}{2} (\omega_q \cos^2 \delta + \omega_r \sin^2 \delta - g \sin 2\delta) \quad (2.42a)$$

$$\tilde{\omega}_r = \frac{1}{2} (\omega_r \cos^2 \delta + \omega_q \sin^2 \delta + g \sin 2\delta) \quad (2.42b)$$

For $\hat{\mathcal{H}}_1$ to be diagonalized, the condition $\tan 2\delta = 2g/\Delta$ must be satisfied. This implies that $\sin 2\delta = 2g/\sqrt{\Delta^2 + 4g^2}$, $\sin^2 \delta = \left(1 - \Delta/\sqrt{\Delta^2 + 4g^2}\right)/2$, and $\cos^2 \delta = \left(1 + \Delta/\sqrt{\Delta^2 + 4g^2}\right)/2$. Substituting these into Eqs. (2.42a) and (2.42b) gives

$$\tilde{\omega}_q = \frac{1}{2} (\omega_q + \omega_r - \sqrt{\Delta^2 + 4g^2}) = \omega_q - \frac{g^2}{\Delta} \quad (2.43a)$$

$$\tilde{\omega}_r = \frac{1}{2} (\omega_r + \omega_q + \sqrt{\Delta^2 + 4g^2}) = \omega_q + \frac{g^2}{\Delta} \quad (2.43b)$$

The vacuum fluctuations of the resonator induce an additional Lamb shift $-g^2/\Delta$ to the transmon frequency and vice-versa for the resonator[48]. Applying the same transformation to $\hat{\mathcal{H}}_{\text{nl}}$ and expanding $\sin \delta \approx \delta$ and $\cos \delta \approx 1$ yields

$$\begin{aligned} \hat{\mathcal{H}}_{\text{nl}} = \hat{U} \hat{\mathcal{H}}_{\text{nl}} \hat{U}^\dagger \approx & - \frac{E_C}{2} \hat{q}^\dagger \hat{q}^\dagger \hat{q} \hat{q} - \delta^4 \frac{E_C}{2} \hat{r}^\dagger \hat{r}^\dagger \hat{r} \hat{r} - 2\delta^2 E_C \hat{q}^\dagger \hat{q} \hat{r}^\dagger \hat{r} \\ & - \delta E_C [\hat{q}^\dagger \hat{r} (\hat{q}^\dagger \hat{q}) + (\hat{q}^\dagger \hat{q}) \hat{q} \hat{r}^\dagger] \\ & - \delta^3 E_C [(\hat{r}^\dagger \hat{r}) \hat{q}^\dagger \hat{r} + \hat{q} \hat{r}^\dagger (\hat{r}^\dagger \hat{r})] \\ & - \delta^2 \frac{E_C}{2} (\hat{q}^\dagger \hat{q}^\dagger \hat{r} \hat{r} + \hat{r}^\dagger \hat{r}^\dagger \hat{q} \hat{q}) \end{aligned} \quad (2.44)$$

The second line in Eq. (2.44) resembles a transmon state-dependent exchange interaction and can be eliminated with another unitary transformation with $U = e^{-\lambda[\hat{q}^\dagger \hat{r}(\hat{q}^\dagger \hat{q}) - (\hat{q}^\dagger \hat{q})\hat{q} \hat{r}^\dagger]}$, with the proper choice of λ :

$$\lambda = \delta \frac{E_C}{\Delta + E_C(1 - 2\sigma^2)} \quad (2.45)$$

The third line in Eq. (2.44) resembles a resonator state-dependent exchange interaction; however, it is preceded by a λ^3 term, which is small and can be eliminated. Finally, the fourth line resembles a two-photon exchange interaction that is dropped due to rapidly-rotating terms. Since $\delta = \frac{1}{2} \arctan(2g/\Delta) \approx g/\Delta$ for small g/Δ , we can therefore express the complete transformed Hamiltonian as

$$\hat{\mathcal{H}} \approx \hbar\tilde{\omega}_q \hat{q}^\dagger \hat{q} + \hbar\tilde{\omega}_r \hat{r}^\dagger \hat{r} + \hbar \frac{\chi_{qq}}{2} \hat{q}^\dagger \hat{q}^\dagger \hat{q} \hat{q} + \hbar \frac{\chi_{rr}}{2} \hat{r}^\dagger \hat{r}^\dagger \hat{r} \hat{r} + \hbar \chi_{qr} \hat{q}^\dagger \hat{q} \hat{r}^\dagger \hat{r} \quad (2.46)$$

where $\chi_{qq} = \alpha$ and $\chi_{rr} = \alpha(g/\Delta)^4$ are the self-Kerr nonlinearities of the transmon and resonator, respectively, and $\chi_{qr} = 2\frac{g^2\alpha}{\Delta(\Delta+\alpha)}$ is the cross-Kerr interaction[42].

The dispersive Hamiltonian is a very useful way of understanding the transmon-resonator interaction. The transformations show that the effect of this interaction is that the resonator inherits nonlinearity from the transmon, and as a result gains a small anharmonicity χ_{rr} . It should be emphasized that this Hamiltonian is approximate; it is only valid for small numbers of excitations or photons in both the resonator and transmon[19], and for small coupling g . In this regime, since g/Δ is small, χ_{rr} is also small. Typical parameters of $\alpha/2\pi = -200$ MHz, $g/2\pi = 100$ MHz, and $\Delta/2\pi = 2$ GHz yield $\chi_{rr}/2\pi = -20$ kHz and $\chi_{qr}/2\pi \approx -1$ MHz. The important parameter in this Hamiltonian is the last term, $\hbar\chi_{qr} \hat{q}^\dagger \hat{q} \hat{r}^\dagger \hat{r}$, which enables both readout and quantum information storage.

2.3.2 Transmon State Readout

The transmon's dispersive coupling to a resonator results in a weak entanglement of the resonator and transmon states. This manifests as the qubit state is mapped onto the frequency of the resonator. To see this, the dispersive Hamiltonian can be rearranged to:

$$\hat{\mathcal{H}} = \hbar\omega_q\hat{q}^\dagger\hat{q} + \hbar(\omega_r + \chi_{qr}\hat{q}^\dagger\hat{q})\hat{r}^\dagger\hat{r} + \hbar\frac{\chi_{qq}}{2}\hat{q}^\dagger\hat{q}^\dagger\hat{q}\hat{q} + \hbar\frac{\chi_{rr}}{2}\hat{r}^\dagger\hat{r}^\dagger\hat{r}\hat{r} \quad (2.47)$$

which is interpreted as a transmon state-dependent frequency shift of the resonator (from here onwards, I will be omitting the tilde and making the substitution $\omega_i = \tilde{\omega}_i$). As a result, measurement of the transmon state can be done by measuring the frequency, a classical property, of the resonator. It should be noted that the inclusion of higher states of the transmon is important for the dynamics of this Hamiltonian. To understand this more clearly, we write down the Hamiltonian after restricting the transmon to the first two levels[40]:

$$\hat{\mathcal{H}} = -\frac{\hbar\omega_q}{2}\hat{\sigma}_z + \hbar\left(\omega_r - \frac{\chi_{qr}}{2}\hat{\sigma}_z\right)\hat{r}^\dagger\hat{r} \quad (2.48)$$

where small self-Kerr of the resonator is neglected. In the two-level approximation, the dispersive shift is $\mp \chi_{qr}/2$ relative to ω_r ; in the multi-level picture, exciting the transmon causes the resonator frequency to only decrease. Since the transmon is a multi-level “atom”, the appropriate Hamiltonian to use when describing the dispersive interaction is Eq. (2.47).

Rather than measure the resonator frequency directly, the state-dependent phase response is measured. A microwave pulse at a probe frequency ω_{RO} is applied through a capacitively-coupled drive line and the phase response is measured (Fig. 2.8b). Because the resonator's center frequency changes based on the transmon state, so too will

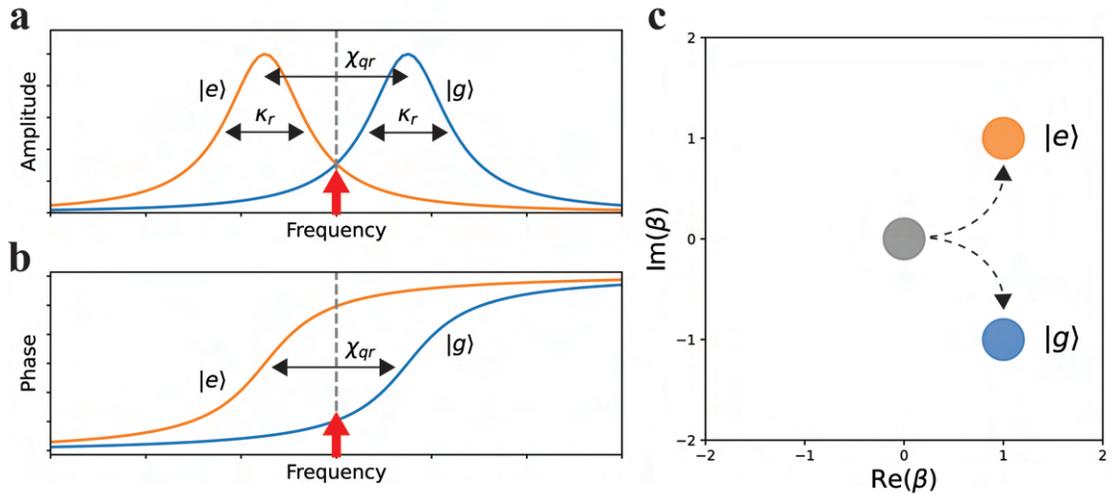


Figure 2.8: **Dispersive readout.** **a** Amplitude vs frequency of readout resonator when qubit is in $|g\rangle$ (blue) or $|e\rangle$ (orange). The Lorentzian peaks are shifted by χ_{qr} . For a readout tone at ω_{RO} (red arrow), there is no change in amplitude. **b** Phase response vs frequency of the readout resonator when the resonator is in $|g\rangle$ (blue) and $|e\rangle$ (orange). For a readout tone at ω_{RO} , the measured phase response is different, resulting in the ability to distinguish between qubit states. **c** Phase space evolution of the readout coherent state. The state travels in different trajectories depending on the qubit state; if the “blobs” that correspond to the quadrature uncertainty overlap, state distinguishability is reduced.

the phase response. To enable fast measurement, the readout resonator is strongly coupled to the drive line; typically, readout resonators used in this thesis have frequencies $\omega_r/2\pi \approx 9$ GHz and linewidths $\kappa_r/2\pi \sim 1$ MHz, allowing readout times $T_r \sim 1$ μ s.

This method of dispersive readout is useful because the transmon’s state can be probed using a readout drive that can populate the readout resonator by many photons, resulting in higher signal-to-noise ratio (SNR) on the measurement outcome compared to more traditional methods of directly measuring the absorption or emission of a single excitation. The latter approach not only has much lower SNR, but also results in a change of state. Dispersive readout is also a quantum non-demolition (QND) experiment, meaning that to good approximation measurement of the resonator does not alter the state of the transmon[49, 50]. As a result, successive measurements can be made in order to further improve measurement SNR, provided that the transmon does not relax to equilibrium. It is also typical to utilize a first-stage quantum-limited amplifier to amplify readout signals before being digitized and processed by the signal acquisition electronics. For this thesis, a SNAIL parametric amplifier (SPA) is used[51].

Readout optimization must be done in the pulse control space in order to maximize measurement outcome fidelity. For distinguishing between the transmon $|g\rangle$ and $|e\rangle$ states, the readout tone may be resonant with the resonator “ $|g\rangle$ ” (ω_r) or “ $|e\rangle$ ” ($\omega_r + \chi_{qr}$) peaks; however, a more optimal frequency would be in-between, $\omega_{RO} = \omega_r + \chi_{qr}/2$. In phase space, this displaces the vacuum state in the resonator and causes it to traverse in mirrored trajectories with respect to the real axis that depend on the transmon state (Fig. 2.8c). The distance of the displaced state from the origin is equal to its displacement; as a result, increasing the readout pulse amplitude increases the separation of the two states, resulting in improved distinguishability. However, the readout amplitude cannot be too high, as the number of photons in the resonator would exceed the limits required for the dispersive approximation to hold. An alternative approach to increasing the readout fidelity then

would be to increase the readout time; integrating the phase information for a longer period of time extracts more information about the qubit state, effectively resulting in the same behavior in the resonator.

In conjunction with pulse optimization, design of the readout resonator and its coupling can also be tuned to improve readout fidelity. Large κ_r allows for faster measurement due to its stronger coupling to the measurement line. The resonator can be populated with photons faster, resulting in faster acquisition of phase information. Additionally, designing $\chi_{qr} \approx \kappa_r$ results in optimal state distinguishability[52]. However, if $\kappa_r \gg \chi_{qr}$, the ability of the drive to distinguish between the two states is greatly diminished. This can be qualitatively seen in the frequency-dependent amplitude behavior of the resonator, where the linewidth of the resonator would be so large that the dispersive shift results in a barely noticeable shift in the Lorentzian peak. In phase space this can be interpreted as the overlap of the states (representing the displacement and phase of the state using polar coordinates and the uncertainty in both quadratures using a filled circle, or “blob”) that correspond to $|g\rangle$ and $|e\rangle$. Conversely, if $\kappa_r \ll \chi_{qr}$, the readout tone would be too off-resonant to sufficiently displace the resonator, resulting in the same behavior in phase space. In such a regime, it is preferable for ω_{RO} to be resonant with either the resonator “ $|g\rangle$ ” peak *or* the “ $|e\rangle$ ” peak, where the measured phase difference between the two peaks would be larger.

2.3.3 Quantum Information Storage in a Resonator

Readout resonators draw their utility from their large coupling $\kappa_{\text{ext}} = \kappa_r$ to the external environment; information about the resonator’s phase leaks out through the readout transmission line at a rapid rate. However, if the resonator is instead weakly coupled to the control line (i.e. $\kappa_{\text{ext}} \ll 1/T_1$), the resonator can be used for quantum information storage, also known as a quantum memory[53]. In such an implementation, the resonator

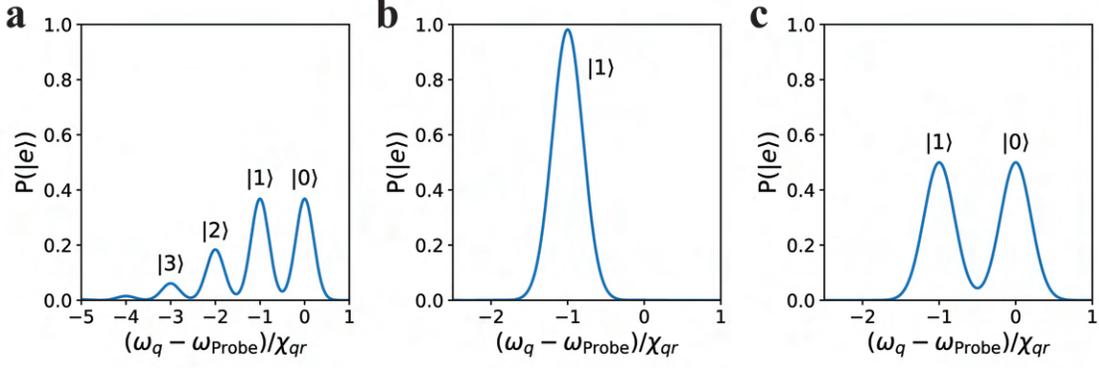


Figure 2.9: **Simulated photon number statistics in a resonator.** **a** Simulated number-resolved spectroscopy on a storage resonator displaced by $\beta = 1$. An X_π pulse was applied using a probe tone frequency ω_{Probe} , resulting in peaks weighted by the photon number probability distribution of the resonator state. The widths of the peaks correspond to the selectivity of the X_π pulse, determined by the pulse's bandwidth. **b** Simulated Fock $|1\rangle$ state preparation using a SNAP gate interleaved between two displacements. Number-resolving spectroscopy reveals that the resonator is found in $|1\rangle$ with close to unity probability. **c** Simulated Fock $(|1\rangle + |2\rangle)/\sqrt{2}$ state preparation.

is used to encode a logical qubit and for quantum information processing rather than the transmon. This has several advantages including improved coherence in an appropriately design resonator (i.e. low intrinsic loss), low intrinsic dephasing, and access to the large Hilbert space of the oscillator to implement hardware-efficient error correction[29, 54]. The transmon then plays the role of an ancilla, whose nonlinearity allows preparation and control of nonclassical bosonic states, which, in addition to displacements, provides universal quantum control of the resonator state[55]. A more comprehensive discussion will be had in Ch. 5, but for now, we can rearrange the Hamiltonian of Eq. (2.47) and interpret it in a different way:

$$\hat{\mathcal{H}} = \hbar (\omega_q + \chi_{qr} \hat{r}^\dagger \hat{r}) \hat{q}^\dagger \hat{q} + \hbar \omega_r \hat{r}^\dagger \hat{r} + \hbar \frac{\chi_{qq}}{2} \hat{q}^\dagger \hat{q}^\dagger \hat{q} \hat{q} + \hbar \frac{\chi_{rr}}{2} \hat{r}^\dagger \hat{r}^\dagger \hat{r} \hat{r} \quad (2.49)$$

Here, we only consider the coupling of the transmon with the storage resonator. In practice, the transmon is also coupled to another resonator for readout. In this interpretation,

the transmon inherits a resonator photon-number dependent frequency shift. This gives rise to number-splitting, where different photon number states will shift the transmon frequency by a different amount. For example, when the resonator is displaced, a coherent state which is a superposition of Fock states is created. This results in the transmon acquiring multiple transition frequencies upon repeated spectroscopic measurement[31], which reproduces the probability amplitudes of the Fock state superposition (Fig. 2.9a).

The ability for a qubit or transmon to probe the photon number statistics of the resonator is very powerful; for one, it provides a sensitive probe for the photon number of the resonator; for another, it allows for single-Fock state preparation. For example, a resonator displacement of $\beta = 1$ followed by an X_π rotation of the qubit at frequency $\omega_q + \chi_{qr}$ will flip the qubit from $|g\rangle$ to $|e\rangle$ if the resonator is in Fock state $|1\rangle$ and leave the qubit in $|g\rangle$ otherwise. Since the resonator is in state $|1\rangle$ with probability $e^{-1} \approx 37\%$, the qubit will also successfully flip with the same probability. If the qubit is now read out with result $|e\rangle$, the resonator's superposition state immediately collapses into $|1\rangle$. Post-selection on successful qubit flips verified through readout therefore forces the resonator state into $|1\rangle$. This method can be used to create any Fock state and relies on the qubit's entanglement with the storage resonator and the use of readout to measure the qubit state. However, because measurement collapses superposition, this method cannot be used to create superpositions of Fock states; a more powerful method based on selective qubit-controlled phase gates is used.

Consider driving a qubit using an $X_{2\pi}$ pulse; a rotation about the x -axis by an angle 2π . This fully revolves the qubit around the axis; the qubit will return to the point where it started. Ostensibly, this does nothing; however, this trajectory taken by the qubit imparts a global Berry phase of π to the system. In general, any closed trajectory taken by the qubit imparts a Berry phase on the system[31, 56]. Because the qubit is coupled to the resonator, the Fock states as well will accrue a phase. However, if the pulse that performs this oper-

ation is number-selective $T_{2\pi} < \pi/\chi_{qr}$, then only the selected Fock state will accumulate a phase. For example, an $X_{2\pi}$ pulse resonant with the $|g, 0\rangle \rightarrow |e, 0\rangle$ (the qubit $g - e$ transition conditioned on the resonator being in vacuum) will impart a π phase on the $|0\rangle$ Fock state on the resonator, while leaving the other Fock states unchanged. This is known as the selective number arbitrary phase (SNAP) gate[55]. With interleaved displacements, SNAP can perform universal control of the quantum states of an oscillator[57].

SNAP can be used to straightforwardly prepare single Fock states and complex arbitrary superpositions of Fock states. In this thesis, SNAP is used to prepare both Fock $|1\rangle$ and Fock $(|0\rangle + |1\rangle)/\sqrt{2}$ states. Both of these states require an initial displacement, a qubit rotation that encloses a phase on Fock $|0\rangle$, and a final displacement that interfere with the probability amplitudes of the Fock $|0\rangle$ and/or $|1\rangle$ states, resulting in the desired state. For example, for Fock $|1\rangle$, a small displacement of $\beta_1 = 1.14$ is performed to produce a coherent state with high probability amplitude in both Fock $|0\rangle$ and $|1\rangle$. An $X_{2\pi}$ gate is performed on the qubit conditioned on the Fock $|0\rangle$ state, which imparts a π phase-shift on Fock $|0\rangle$. Then a final displacement with amplitude $\beta_2 = -0.58$ destructively interferes with the Fock $|0\rangle$ state, causing its probability amplitude to fall to nearly zero. The result is the creation of the Fock $|1\rangle$ with high fidelity, as shown in Fig. 2.9b. In contrast to the post-selection method, which has a success probability of only $\approx 37\%$, this method can deterministically generate the desired state. To generate the superposition state $(|0\rangle + |1\rangle)/\sqrt{2}$, the same sequence of gates is used, but with $\beta_1 = 0.56$ and $\beta_2 = -0.24$ (Fig. 2.9c)[53].

While SNAP can be used to generate arbitrary states in an oscillator, more complex states that what were described in the previous paragraph could require many combinations of phase gates and displacements, which can rapidly become inefficient and take significantly long times to finish state preparation. While this was not an issue for the experiments carried out in this thesis, it should be noted that another, more numerical

method can be used to optimally define a control pulse for state generation and manipulation. Rather than sequentially performing the displacements and phase gates, they are performed simultaneously. The result is a numerically optimized pulse that efficiently generates the desired state. The numerical optimization method is called gradient ascent pulse engineering (GRAPE) and was developed for spin control in NMR[58].

2.4 cQED in an Arbitrary Electromagnetic Environment

In this chapter the discussion has thus far been limited to describing the properties of simple lumped-element circuits involving either one distinct resonator mode and/or one transmon mode. In practice, because cQED operates in the microwave regime where the wavelength $\lambda = c/f = 1 - 10\text{cm}$, the spatial dependence of the electromagnetic fields must be taken into account. Moreover, practical circuits are dielectrically loaded by the substrates on which they are fabricated, which reduces the wavelength even further by the dielectric constant. Each capacitor has a stray geometric inductance and each inductor has stray geometric capacitance[59]; their nonzero sizes give rise to standing wave modes. Moreover, the use of distributed-element transmission line resonators and microwave cavities is ubiquitous in cQED; their field behavior gives rise to fundamental harmonic oscillator modes and an infinite set of higher order modes that must be properly understood in order to describe how circuit elements couple to each other. Additionally, these circuits are embedded into a greater electromagnetic environment that includes the package in which the device is housed and the materials that comprise the device. This environment is subjected to geometric conditions that support their own standing waves, resulting in “package” modes and their higher-order harmonics. As a result, a distributed-element approach must be adopted in order to more precisely model and design these circuits and the effects of the environment.

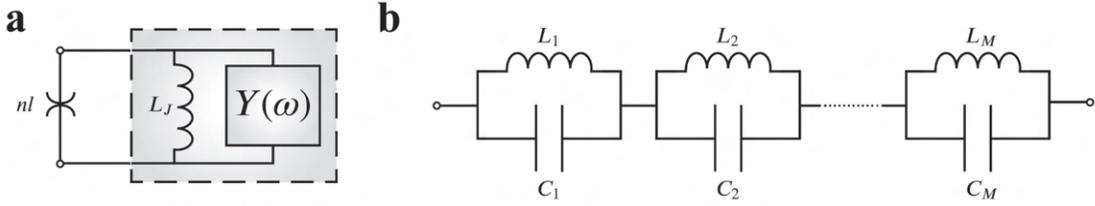


Figure 2.10: **Josephson junction embedded in an arbitrary electromagnetic environment.** **a** Josephson junction connected to the port of a black box whose admittance describes the linear electromagnetic environment. The junction is separated into its linear inductance L_J , which is included in the linear network, and its nonlinear component, which is treated separately. **b** Foster decomposition of the linear microwave network under the lossless assumption.

Precise modeling of a quantum circuit and its environment involves solving Maxwell’s equations under geometrically defined boundary conditions. This is made easier by treating the nonlinearity in a perturbative manner; Josephson junctions are treated as linear lumped inductors in parallel with a nonlinear element (the junction is the only element whose size is orders of magnitude smaller than the microwave wavelength; it can be treated as lumped). The system is often diagrammed as a nonlinear element that is connected to a black box that contains an arbitrary linear electromagnetic embedding network. The problem then reduces to solving the linear network to find the eigenmodes, their frequencies, and their coupling to the linear element (Fig. 2.10a). There are two main approaches to doing this, and both utilize finite-element simulation to solve the linear network.

2.4.1 Black Box Quantization

The first method is called black box quantization (BBQ)[60]. In this method, the linear network is treated as an arbitrary N -port microwave network, where N is the number of nonlinear elements present in the circuit. The network can therefore be fully described by an admittance matrix $Y(\omega) = Z(\omega)^{-1}$ which can be decomposed using Foster’s theorem

(Fig. 2.10b). For the single-junction case $N = 1$, the decomposition is

$$Z(\omega) = \sum_{k=1}^M \left(i\omega C_k + \frac{1}{i\omega L_k} \right)^{-1} \quad (2.50)$$

where M is the number of eigenmodes in the diagonalized system. For each mode, whether it is defined by the existence of a lumped-element subcircuit or a distributed-element network inside the embedding network, has an effective capacitance C_k and inductance L_k in parallel. The zeros of $\text{Im}Y(\omega)$ give the mode frequencies $\omega_k = 1/\sqrt{L_k C_k}$, and $C_k = \text{Im}Y'(\omega_k)/2$, where $\text{Im}Y'(\omega)$ is the first derivative of $\text{Im}Y(\omega)$. With these relations, the black box network is effectively renormalized to the eigenmode basis. Because coupling to the Josephson junction amounts to a mode eigenfield inducing a phase across the junction, the total phase across the junction is a sum of phase contributions from all the modes:

$$\hat{\Phi} = \sum_{k=1}^M \sqrt{\frac{\hbar}{2} Z_k^{\text{eff}}} (\hat{a}_k + \hat{a}_k^\dagger) \quad (2.51)$$

where $Z_k^{\text{eff}} = \sqrt{L_k/C_k}$ is the effective characteristic impedance of the mode. With the flux operator defined, the nonlinear term in the Hamiltonian $\hat{\Phi}^4/(24\Phi_0^2 L_J)$ can be reintroduced and treated perturbatively, yielding the modes' inherited self- and cross-Kerr couplings. Higher order terms can be introduced as well; the perturbation can be treated numerically to arbitrary order.

2.4.2 Energy-Participation Quantization

The second method is called the energy-participation ratio (EPR) method and was more recently introduced[61]. This approach is similar in that it treats the nonlinearity as a perturbation to a linear system. However, the EPR method does not diagonalize the system through a Foster decomposition; instead, the junction is treated as a lumped-element linear

inductor in finite-element simulation. The system then becomes entirely linear and can be numerically diagonalized in analogy with the Bogoliubov transformation employed in Eq. (2.41):

$$\hat{\mathcal{H}}_1 = \sum_k^M \hbar\omega_k \hat{a}_k^\dagger \hat{a}_k \quad (2.52)$$

within the finite-element simulation, this diagonalization is accomplished by numerically solving Maxwell's equations and calculating the spatial field distribution of each mode. As a result, calculating the voltage induced by each mode's eigenfield is straightforward to do. The EPR method is based around calculating the fraction p_k of total magnetic energy stored in the junction for each mode:

$$p_k = \frac{\frac{1}{2}L_J I_J^2}{U_{\text{tot}}} \quad (2.53)$$

For the transmon, p_J is very close to unity. For weakly coupled or uncoupled modes, the junction participation is close to zero. The degree to which the inductive energy of a mode is stored in the junction determines the strength of the zero-point fluctuation of the junction's flux operator:

$$\hat{\Phi} = \sum_{k=1}^M \sqrt{p_k \frac{\hbar\omega_k}{2E_J}} (\hat{a}_k + \hat{a}_k^\dagger) \quad (2.54)$$

which, like with BBQ, can be used to numerically calculate the Kerr matrix.

While the two methods are very similar, in practice they involve different workflows. BBQ is advantageous in that the admittance matrix provides not only the mode frequencies, but also their impedances. However, determination of the admittance matrix requires full simulation of the system as a function of frequency; effectively this requires a resimulation of the system over the entire desired range of frequencies. While some

finite-element simulators like Ansys HFSS have adaptive methods that significantly reduce simulation time, BBQ often times requires multiple rounds of simulation in order to yield workable results: an initial coarse simulation to find the mode frequencies, and a second simulation where the drive frequencies close to the mode frequencies are more finely swept. The EPR method, on the other hand, only requires one well-meshed eigenmode simulation in which the fields for all the eigenmodes are calculated; as a result, a single simulation is sufficient to extract relevant circuit parameters. One drawback of the eigenmode approach is that the simulation eigenspace must be truncated to prevent the simulation from becoming too computationally intensive; as a result more complex systems which contain many modes are more difficult to simulate.

2.5 cQED Architectures

The experimental implementation of practical quantum circuits in cQED has a history of around two decades, and can be traced back to work performed here at Yale in the early 2000s. The field has matured significantly since then; initially, planar circuits were fabricated on a chip that was embedded in a copper box. As time progressed, that circuit architecture was refined and expanded upon, and new architectures with different features were developed. In this section, I will provide an overview of commonly used modern cQED architectures and describe how transmon qubits and resonators are implemented in planar circuits and in 3D cavities.

2.5.1 Planar Architecture

In a planar architecture, circuits lithographically are patterned using thin-film superconductors deposited on a sapphire or silicon substrate or chip. The chip are typically wire-bonded to a printed circuit board (PCB) and enclosed in a copper box. The PCB acts

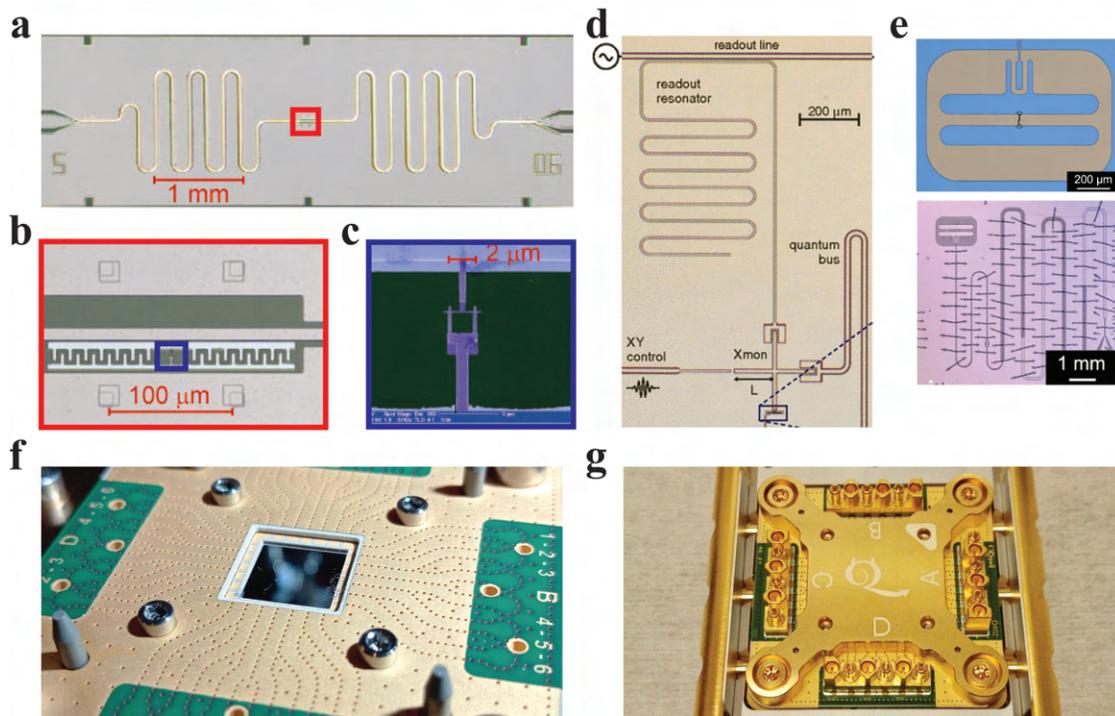


Figure 2.11: **Planar cQED architecture.** **a** Early circuit design. Planar transmission lines act as a waveguide to couple to a qubit, made using an **b** interdigitated capacitor shunted by a nonlinear element made up of **c** Josephson junction(s). Figures **a**, **b**, and **c** are adapted from Schuster et al. [62]. **d** Modern circuit design. X-shaped transmon qubits (“Xmon”) couple to transmissions lines for control, and to quarter-wave CPW resonators for readout. Figure adapted from Barends et al. [63]. **e** New planar transmon design (“barbell”) optimized for high coherence. Large capacitor pads separated from each other and the ground plane by 50 – 100 μm , resulting in reduced sensitivity to surface losses. The transmon is coupled to a CPW resonator and controlled via impedance-matched transmission lines that continue off-chip. Figure adapted from Place et al. [64]. **f** Circuit-containing chip (center) wire-bonded to a PCB that contains signal-carrying transmission lines. The wire-bonds short the on-chip ground plane with the ground of the PCB, which is in turn connected to the ground of the fridge. Wire-bonds are also used to connect the signal-carrying transmission lines to the on-chip transmission lines. **g** The chip is then enclosed within a copper or gold-plated copper box to shield the circuit from the external environment. Launchers on the PCB connect the signal-carrying transmission lines to coaxial cables that carry the signals into and out of the device. Figures **f,g** are adapted from Andresen [65] and are representations of a commercial planar cQED package called QCage.24 by QDevil APS.

as a routing intermediary between the classical, normally conducting control lines that originate from room-temperature microwave signal generators and the superconducting circuits that are located on the chip. The control lines are $50\ \Omega$ transmission lines that continue on the chip to deliver signals to specific elements. This allows for a large amount of versatility in routing control lines on the chip to deliver signals to specific parts of the chip. The presence of an on-chip co-planar ground plane is a distinguishing feature of the planar architecture and effectively allows for the implementation of compact filtering on control lines as well as the flexibility to realize both capacitive and inductive (flux) coupling to qubits and/or resonators. A description of the architecture and the technical design considerations required is given in Huang et al. [66].

The planar architecture is named as such because the circuits are all defined on one plane (the substrate surface). It is important to note, however, that the terminology of “planar” refers to the physical implementation of the circuits that define the useful modes for computation and not the behavior of electromagnetic fields, which always exist in three dimensions. The architecture is also often times referred to as “2D” or “on-chip” within the field, in order to contrast it with the 3D architecture which will be discussed in the Sec. 2.5.2. However, these features do not unambiguously distinguish this architecture from the other two architectures that will be discussed in Secs. 2.5.2 and 2.5.3; these other architectures also utilize planar circuits fabricated on a chip. In addition, modern large-scale circuits[20, 21, 66–69] implement flip-chip architectures where the circuits are not all defined on two separate planes; this can in some sense be described as a 2.5-dimensional circuit.

The uniquely distinguishing feature behind this architecture seems to be the presence of a ground plane on-chip; circuits are therefore embedded within co-planar waveguide (CPW) transmission lines. The presence of a ground plane in close proximity to circuit elements allow electromagnetic field lines to terminate there, resulting in highly compact

electromagnetic profiles. This allows circuits to be highly compact, resulting in the increased circuit density required to implement hundreds of qubits on a \sim cm-scale chip. However, the drawback of having compact electromagnetic fields is that they tend to produce high field densities near lossy regions on substrate and superconducting surfaces, which leads to increased susceptibility to decoherence mechanisms. Additionally, the normally-conducting copper enclosure will define a package mode resonance at ≈ 15 GHz whose nonlocalized field can couple to all qubits in the device, resulting in increased inherited loss and unwanted cross-talk that can lead to correlated error and reduced gate fidelity. Nevertheless, development of the planar architecture has come a long way. 14 years ago state of the art transmons had relaxation times $\sim 1 \mu\text{s}$ [70]; improvements in transmon design, materials, and fabrication processes have all played a role in bringing steady improvements to several $\sim 10 \mu\text{s}$ [20, 63] and in more recent years, to $\sim 100 \mu\text{s}$ [39, 64, 71]. Examples of various planar circuit designs are shown in Fig. 2.11.

Resonators in the planar architectures are implemented using shorted transmission lines or pseudo-lumped-element LC circuits and are commonly used for readout, filtering, and coupling multiple qubits together. Quarter-wave CPW resonators are short-circuited transmission line to ground on one end and open-circuited on the other. The resonance frequencies are determined by their length which can be tuned with lithographic precision. Lumped-element circuits can also be implemented using interdigitated capacitors and meandering inductors. These elements tend to have low self-resonant frequencies (~ 1 pF capacitors or ~ 1 nH inductors have enough geometric inductance or capacitance to achieve self-resonant frequencies of $\sim 15 - 20$ GHz) due to their nonzero size. As a result they can be thought of as pseudo-lumped elements. Both quarter-wave resonators and “lumped” element resonators have higher-frequency modes that require appropriate modelling with a distributed-element treatment.

The design of a transmon qubit has evolved over the years; notably, the shunting ca-

capacitor’s design has been considered to be the dominant source of decoherence. Early shunting capacitors were interdigitated with capacitor electrode separations $\sim 1 \mu\text{m}$; this allowed for highly compact transmon designs at the cost of reduced coherence[72] for reasons mentioned earlier. More modern designs such as the “Xmon”[20, 63] increased the physical size of the capacitor with separation from the ground plane $\sim 10 \mu\text{m}$, resulting in less spatially confined electric fields and therefore lower loss. The “barbell” design takes this to the extreme[39, 64, 71], where the capacitor pads are floating and are separated by far larger length scales (50 – 100 μm). This also allows for even greater separation of the capacitors pads from the ground plane. The “barbell” design in fact begins to resemble the design of the “3D” transmon[73, 74] which is used in 3D cavity architectures and will be discussed next. Notably, the evolution of transmon design has also been paired with improvements in materials and fabrication processes to achieve relaxation times that are in the hundreds of μs range. To my knowledge, Place et al. [64] has reported the highest “Xmon” relaxation times, approaching 300 μs (time-averaged), while almost 400 μs (median) relaxation times have been reported for the “barbell” design by Gordon et al. [71].

2.5.2 3D Cavity Architecture

In the early 2010s when transmons in the planar architecture were struggling to achieve coherence times beyond a few microseconds, the 3D architecture was developed. To address the problem of transmon coherence, larger transmon designs were employed in which the capacitance was defined using large area superconducting pads that were separated from each other by as much as 50 – 100 μm ; the self-capacitance of these pads dominate the shunting capacitance of the transmon C_s . The transmons were then embedded inside a superconducting enclosure whose walls served as the ground. The enclosure itself supported standing wave modes, the most fundamental of which was used as a readout resonator for

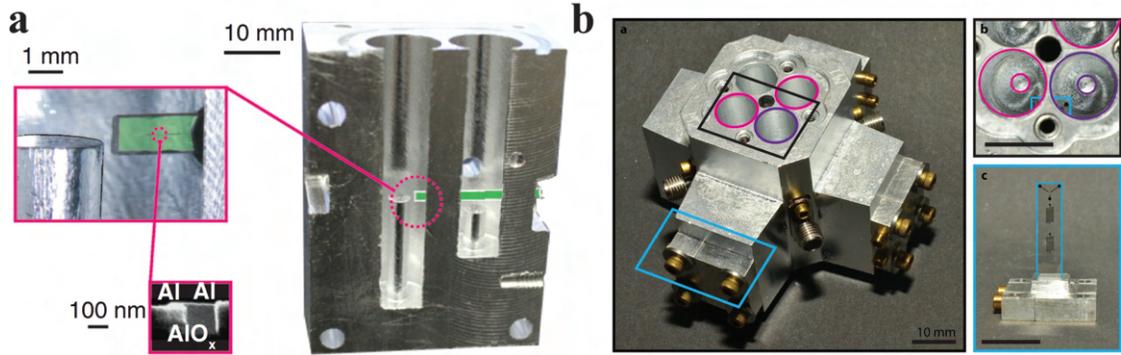


Figure 2.12: **3D cavity cQED architecture.** **a** Coaxial post cavity with transmon inserted through a side tunnel. Figure adapted from Reagor et al. [53]. **b** Multicavity experiment, where side-tunnels are used to couple adjacent cavities together using a transmon-based mixing element. Figure adapted from Chou et al. [75].

the transmon[73]. This served as the basis for the 3D cavity architecture. The electromagnetic fields of the enclosure are mostly in vacuum, which is lossless. As a result, the transmon is embedded in an environment that has low loss. In addition, the transmon does not have a nearby ground plane; as a result, its electric fields are far more spatially diffuse and localized in vacuum, resulting in less sensitivity to surface losses.

Modern 3D cavity architectures are based around the observation that cavity resonators, due to their high energy participation ($U_{\text{vac}}/U_{\text{tot}} \approx 1$) in vacuum, have coherence times ~ 1 ms[53]. Recent experiments using surface processing techniques from the particle accelerator community have shown that niobium cavity resonators can be made with coherence times approaching ~ 10 ms[76] to ~ 1 s[77] (although coupling to a qubit has not yet been demonstrated for the ~ 1 s cavity). These coherence times are an order of magnitude higher than those of modern transmon qubits; in addition, they seem to have little to no dephasing aside from what is inherited from the coupled transmon. As a result, the 3D cavity architecture is a promising platform to implement quantum information storage in the highly coherent 3D cavity mode.

The most common implementation of a 3D cavity for quantum memories is known as the $\lambda/4$ coaxial post cavity (Fig. 2.12a)[31, 53]. It is effectively a shorted quarter-

wave transmission line that is machined out of a block of high-purity bulk superconductor (most commonly, 5N Aluminum), resulting in a monolithic cavity. After applying the proper surface treatments that remove damage and contaminants introduced during the machining process, these cavities have coherence times exceeding 1 ms. The transmon is introduced by inserting the chip into a narrow tunnel that is machined into the side of the cavity; the design of the 3D transmon here is important not only for improved coherence (modern tantalum-based 3D ancilla transmons have been demonstrated to have relaxation times approaching 300 μ s[29]), but also to increase the size of the qubit’s dipole moment, resulting in stronger coupling to the storage mode. This approach is somewhat modular in the sense that the transmon chip can be easily removed and replaced if the transmon has unexpected performance. Additionally, the cavity forms an enclosure around the quantum circuit, allowing each bosonic qubit or “module” to be segregated from others, reducing cross-talk and unwanted interaction. In this configuration, all the operations discussed in Sec. 2.3.3 are possible to implement.

There are a few drawbacks to the 3D cavity architecture. Because there is no ground plane on chip, complex on-chip filtering and delivery of magnetic flux is difficult. The implementation of the latter is even more difficult due to the enclosure being a superconductor; external flux delivery is severely attenuated due to the Meissner effect. While some solutions exist to address this issue[25, 78, 79], they increase the overall footprint of the device and introduce a fair amount of hardware overhead. Additionally, while several multiqubit systems have been demonstrated in which a single qubit or multiple qubits are coupled to small numbers (< 10) of cavities (Fig. 2.12b)[75, 80, 81] using cross-tunnels, the path towards a system with hundreds of bosonic qubits encoded in 3D cavities is not clear, as such a system would require a large volume and would fill up the available space in a dilution refrigerator. Moreover, because the cavities are machined into a block of high-purity aluminum, they are not modular; one defective or unexpectedly low coherence

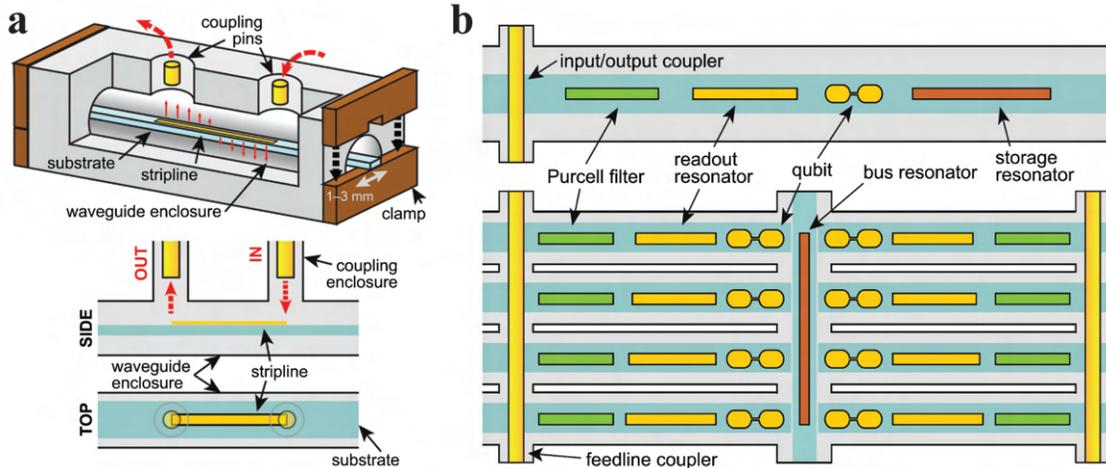


Figure 2.13: **Coaxial tunnel architecture.** **a** Striplines and transmons (3D transmon design from Paik et al. [73]) are patterned on chips and inserted into cylindrical tunnels made of high-purity aluminum. **b** The tunnel architecture can support filter resonators, readout resonators, qubits, and stripline-based storage resonators for bosonic qubits. Adjacent tunnels can couple together using a bus resonator, providing a blueprint for multiqubit systems. Figures adapted from Axline et al. [84].

cavity can ruin the entire system, requiring it to be remade. While it may be possible to make each cavity independently and attach them together, doing so will introduce significant and likely prohibitive contact resistance at the joints between cavities, degrading coherence.

A more compact approach is available in the form of micromachined cavities, in which silicon micromachining techniques are used to define a microwave resonating cavity on a chip [82, 83]. This is a flip-chip device in the flavor of planar architectures and has been successfully demonstrated to be highly coherent while remaining compact. However, the path towards scaling up to multiple qubits is not clear. A promising alternative approach comes in the form of a hybrid architecture called the coaxial tunnel architecture.

2.5.3 Coaxial Tunnel Architecture

The coaxial tunnel architecture is not a new concept; it was first introduced in Axline et al. [84] as an approach to seamlessly couple planar devices to 3D post cavities without introducing additional loss. Here, I will discuss a subset of the coaxial tunnel architecture that is exclusively based on using the modes of planar circuits on a chip that is inserted into a cylindrical tunnel waveguide made of high-purity aluminum. In a sense, the architecture amounts to removing the cavity resonator in the post-cavity module and using an on-chip stripline resonator patterned on the chip as the storage mode (Fig. 2.13b). Each tunnel contains the circuitry for the storage mode, the ancilla transmon, and a readout resonator, defining a “module”. In this way, a bosonic qubit can be realized in a far more compact way than the conventional 3D approach, as multiple tunnels can be placed adjacently to realize multiqubit systems. The tunnels themselves provide low-loss enclosures in which the circuits patterned on chips reside. This results in a well-defined ground plane defined by the walls of the tunnel. Effectively, the circuit with the ground resemble a coaxial transmission line, with the patterned circuit being the center pin; impedance-engineering the center pin by the patterned circuits’ dimensions or including Josephson junctions define the standing-wave resonance modes used for quantum information processing.

This architecture takes cues from both the 3D and planar architectures. For one, the qubit, the storage resonator, and the readout resonator along with optional filtering are all lithographically patterned on a chip (Fig. 2.13a); as a result, their frequencies and couplings can be controlled with high precision. In similarity with a 3D cavity architecture, the tunnels are fully enclosed but for cross-tunnels that can be used to couple circuits from separate tunnels together, as can be seen in Fig. 2.13b. This, combined with the tunnels’ high cutoff frequencies of > 18 GHz, circumvent the problem of nonlocal package modes that than inadvertently couple multiple qubits together. However, in contrast to the 3D

cavity architecture, the tunnel modes do not play the role of the quantum memory; as a result, defects in the surface quality of a particular tunnel do not limit the coherence of the storage mode. The modularity of the system is therefore enhanced, as faulty storage resonators can easily be replaced by replacing individual chips. Finally, the continued lack of an on-chip ground plane allows for spatially diffuse fields that reduce loss. The coaxial architecture adapts the best of both systems, resulting in a hybrid architecture.

Despite this architecture existing for almost seven years at the time of writing, it has not been an attractive approach for implementing information processing with bosonic qubits. This is due to the historically low coherence of on-chip storage modes[84, 85], the most coherent of which had relaxation times of at most 250 μ s. Moreover, such coherence times were unreliably achievable and made them an unattractive platform choice when compared to the cavity-based approach. However, in Chs. 4 and 5 I will show that the origins of this decoherence and its variations can be traced to inconsistency in losses due to the tunnel packages themselves in addition to dielectric losses in the substrate and various interfaces between the superconductor, the substrate, and vacuum (which I will often use interchangeably with “air”). As a result, with the proper choices of materials, processes, and design, storage mode resonators can be implemented on-chip with coherence times that reliably exceed 1 ms[86], rivaling those of post cavities. Before that discussion, it will be important to delve deeper into intrinsic sources of dissipation in superconducting circuits and how to quantify them, which I will do in the next chapter.

2.6 Suggestions for Further Reading

For a review of cQED I suggest the reader to refer to *A quantum engineer’s guide to superconducting qubits* by Krantz *et. al.*[42], and *Circuit quantum electrodynamics* by Blais *et. al.*[40]. Additionally, I have found David Schuster’s thesis[37] to be a great

resource for a contemporary introduction to the field.

Chapter 3

Decoherence in Superconducting Quantum Circuits

In Ch. 1.3 I briefly discussed how noise in a qubit's environment can lead to decoherence. The existence of noise is fundamental; even a qubit in free space can spontaneously emit. This is due to the existence of a continuum of modes whose vacuum fluctuations stimulate the qubit to emit a photon, causing it to decay. The photon can then be re-emitted and absorbed by any of the infinite number of modes in the continuum; it is unlikely for the qubit to re-absorb it. The rate at which emission occurs can be altered by enclosing the atom or transmon in a cavity or package with a well-defined and discrete set of electromagnetic modes. The discrete set of modes dramatically reduces the density of states into which a qubit can emit a photon[87]. The emission rate can then be controlled by engineering the modes of the enclosure to be off-resonant with the qubit[70]. As discussed in Ch. 2.5, the enclosure or package can take on many forms, such as a copper box containing a PCB and circuit-containing chip or a superconducting waveguide cavity, both of which have modes that are engineered to be far detuned in frequency to the modes of a quantum circuit.

The existence of an enclosure divides the universe with respect to a quantum circuit, which can be as basic as a single qubit or as extensive as a multi-qubit systems with cou-

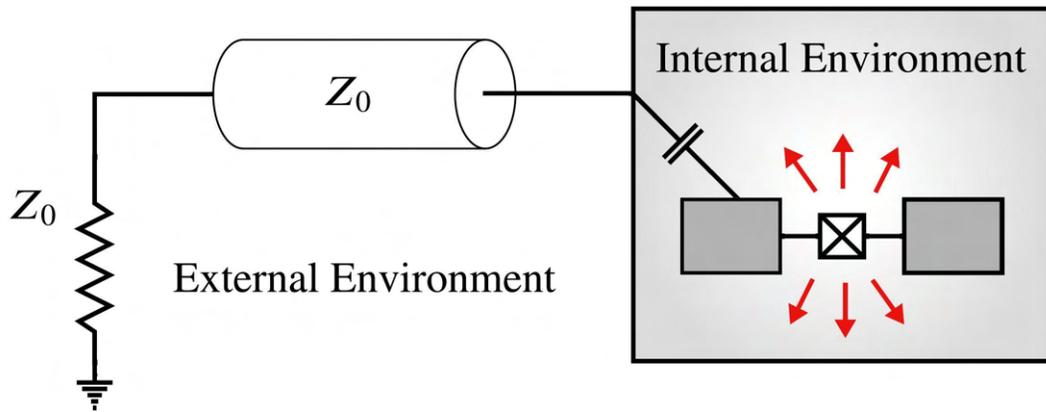


Figure 3.1: **External vs internal environment.** Circuits are separated from the external environment that hosts a continuum of electromagnetic modes. Transmission lines carry signals back and forth between the internal and external environments to measure and control the circuit. Qubits in the internal environment can radiate and leak energy into the transmission line through its capacitive coupling.

pling elements, readout resonators, etc. The environment inside the enclosure will be referred to as the internal environment; the environment outside will be called the external environment. In order to control and measure the circuit, it will be necessary to form a connection between these two environments; as a result, there will always exist a coupling between the circuit and the external world. These couplings are physically implemented through control lines or drive lines, which are transmission lines that carry electrical signals from room-temperature signal generation equipment into the dilution refrigerator and to the quantum device. Since control lines provide a pathway for the external environment to interact with the internal environment, qubits in the circuit can emit photons into the control lines' continuum of electromagnetic modes. Additionally, thermal noise in the lines can both stimulate the qubits to absorb photons, resulting in their heating, and cause frequency fluctuations, resulting in their dephasing (rate $\Gamma_\phi = 1/T_\phi$). The detailed balance

between decay and heating defines the relaxation rate[42]:

$$\Gamma_1 \equiv \frac{1}{T_1} = \Gamma_{1\downarrow} + \Gamma_{1\uparrow} = (1 + e^{-\hbar\omega_q/k_B T}) \Gamma_{1\downarrow} \quad (3.1)$$

where ω_q is the qubit frequency, k_B is the Boltzmann constant, T is the operating temperature, and $\Gamma_{1\downarrow}$ and $\Gamma_{1\uparrow}$ are the decay and heating rates, respectively. The heating rate is exponentially attenuated by temperature, which necessitates operating temperatures $T \approx 20$ mK, where the ratio of heating rate to decay rate is suppressed to $\Gamma_{1\uparrow}/\Gamma_{1\downarrow} \sim 10^{-6}$ for a typical transmon frequency of 5 GHz. This also places a constraint on how low the frequency of a qubit or resonator can be, as modes with frequencies < 3 GHz will have relative heating rates above 1%. However, ambient infrared radiation and excess thermal noise in control lines due to inadequate thermalization results in an additional heat load on the quantum circuits; they can often times be at an equilibrium temperature that is significantly higher than the base temperature of a dilution refrigerator. At Yale, it is not unusual for qubits to have an effective temperature 40 – 80 mK. As a result, heating can be a significant limiter to operational fidelity if proper thermalization, radiation shielding, and RF filtering of the control lines are not implemented.

Decoherence due to external coupling can be controlled through the microwave engineering of the control lines and the modes of the enclosure. In principle, the effects of external coupling can be minimized by reducing the thermal noise temperature and by reducing the coupling capacitance C_c of the drive line to the quantum circuit and to the modes of the enclosure. Decoherence due to the internal environment, on the other hand, is far more difficult to control. This is because there are numerous other decoherence channels present within the materials used and altered by device fabrication processes employed to create the enclosure and the quantum circuit. These channels can couple to all electromagnetic modes of the internal environment; creating additional pathways for de-

phasing or relaxation to occur. Finally, because these channels all have the same effect of reducing T_1 and T_2 , they are difficult to characterize and distinguish from each other.

In this chapter, various decoherence mechanisms in superconducting quantum circuits will be discussed. Importantly, the same sources of decoherence that affect transmon qubits also affect microwave resonators. This is because if we restrict the transmon to the lowest two states and operate in the low drive-strength limit where leakage to higher computational states is negligible, we can approximate the transmon as a harmonic oscillator with frequency ω_q . To simplify the discussion, relaxation mechanisms will therefore be discussed in the context of resonators, with the same treatment being applicable to transmons. I will begin by introducing the participation ratio model as a classical field-based approach to quantifying energy losses in circuits, followed by a description of the relevant sources of loss and how a resonator's sensitivity to them is calculated. I will then discuss dephasing mechanisms in resonators and qubits, and how they can be mitigated. Finally, I will detail how coherence is measured in qubits, resonators, and quantum memories.

3.1 The Participation Ratio Model

The participation ratio model is a useful way to describe sources of energy loss in microwave resonators and transmon qubits. When an electromagnetic mode in the internal environment is excited, the spatial distribution of electric and magnetic fields will overlap with particular dissipative regions in the environment and to whatever pathway exists to the external environment. This leads to a fraction of mode energy being lost every period of oscillation. we can write this more concretely by defining the quality factor Q (or loss $1/Q$) of the resonance mode:

$$\frac{1}{Q} = \frac{1}{\omega T_1} = \frac{P_{\text{diss,tot}}}{\omega U_{\text{tot}}} = \frac{1}{Q_{\text{ext}}} + \frac{1}{Q_{\text{int}}} \quad (3.2)$$

where U_{tot} is the total energy stored in the mode and $P_{\text{diss,tot}}/\omega$ is the total power dissipated in one period of electromagnetic field oscillation. The quality factor is related to the energy decay time T_1 by $Q = \omega T_1$, where ω is the frequency of the mode. For a qubit, T_1 is the characteristic time it takes for a qubit in $|e\rangle$ to lose an excitation and return to $|g\rangle$. For a resonator, it is the characteristic time it takes for a coherent state $|\beta = 1\rangle$ to decay to the vacuum state, or the time it takes for a Fock $|1\rangle$ state to decay to Fock $|0\rangle$ (which is also the vacuum state). Here, I am using T_1 to describe a decay time where it is actually a relaxation time; for the rest of this chapter, I will assume that the heating rate is very small unless otherwise noted, i.e. $\Gamma_{1\uparrow}/\Gamma_{1\downarrow} < 10^{-2}$ (qubit effective temperature < 50 mK), so $1/T_1 \approx \Gamma_{1\downarrow}$.

In Eq. (3.2) I have also expressed the total loss $1/Q$ as a sum of losses from the internal and external environments, $1/Q_{\text{ext}}$ and $1/Q_{\text{int}}$. The external quality factor quantifies how much energy is lost through the connection between the internal and external environments; in an otherwise closed package this energy loss is exclusively through emission into the control lines. Q_{ext} is also known as the coupling quality factor Q_c , as it is directly related to the capacitive or inductive coupling of the circuit to the drive line. The loss of the internal environment, on the other hand, is given by $1/Q_{\text{int}}$ and is determined by the spatial field distribution of the resonance mode and the intrinsic sources of dissipation located within the materials that comprise the circuit. The participation ratio model is used to describe the various sources of dissipation. We can describe the total internal loss as

$$\frac{1}{Q_{\text{int}}} = \sum_i \frac{1}{Q_i} = \sum_i \frac{P_{\text{diss},i}}{\omega U_{\text{tot}}} = \sum_i \frac{U_{\text{diss},i}}{U_{\text{tot}}} = \sum_i p_i \Gamma_i \quad (3.3)$$

where $1/Q_i$ is the loss due to the i th loss channel and $U_{\text{diss},i} = P_{\text{diss},i}/\omega$ is the energy dissipated in a cycle by the i th channel. Of the energy stored in a lossy region, only a fraction of it is lost every cycle. Therefore, we can relate the energy stored U_i to the power

dissipated by $U_{\text{diss},i} = U_i \Gamma_i$, where Γ_i is called the intrinsic loss factor corresponding to the i th loss channel. Finally, we define a participation ratio $p_i = U_i/U_{\text{tot}}$. The participation ratio here is a ratio of energies and is identical in nature to the EPR of the Josephson junction discussed in Ch. 2.4.2; the only difference is in their application.

The power of the participation ratio model lies in its ability to distinguish between intrinsic losses and resonator geometry. The geometric design of a resonator determines the spatial distribution of its electromagnetic field; therefore, the field overlap with a lossy region, given by p_i , is a geometric quantity that can be calculated analytically for simple geometries or by finite-element simulation for arbitrary geometries. The loss factor, on the other hand, is intrinsic to the quality of the materials that make up the device and is modified by whatever processes and treatments are applied during fabrication. For example, a particular thin-film deposition process can impact how the film grows, thereby affecting its superconducting properties or surface chemistry. Or, a particular chemical cleaning may remove contaminants that would have otherwise contributed to excess loss. The materials used in a device are inseparable from the processes used to make it; they must be considered together.

There can be multiple loss channels present in a lossy region that all dissipate energy through a variety of mechanisms. For a well-defined mechanism, we can express the loss factor as a quantity that better describes the nature of the energy loss. Dielectric loss, for example, is described by a loss tangent, $\tan \delta_{\text{delta}} = \Gamma_{\text{diel}}$; conductor loss can be defined by a conductor quality factor $q_{\text{cond}} = 1/\Gamma_{\text{cond}}$. Regardless, loss factors are largely phenomenological quantities; they must be empirically measured for a particular set of materials and fabrication processes. The difficulty in doing this lies in the limited information we can obtain about the loss factors. In a resonator, the main measurement we have is the measurement of the quality factor, or a linear combination of loss factors; measurement of a single loss factor therefore cannot be done with one resonator or resonance

mode alone. In Ch. 4 I will discuss how we can gain a better understanding of or even measure loss factors.

It is important to note that the participation ratio model is only so useful in describing resonator losses as the set of loss channels assumed to be important. We assume that there exists a set of loss channels that constitute most of if not all of the total loss in a resonator. If that assumption is violated, we have an incomplete understanding of what limits a particular resonator. It is often the case where a particular source of loss can be safely neglected due to its insignificance; however, at times there may be loss channels that are not considered but could have a significant impact. It is therefore crucial to have a good understanding of the circuit architecture, the circuit geometry, and the materials used for a system under study, before writing down the participation ratio model.

In the next few sections I will detail the various loss channels we consider to be important in cQED.

3.2 External Losses

Quantum circuits are usually very weakly coupled to the control lines, i.e. $Q_{\text{ext}} \gg Q_{\text{int}}$, to ensure that the external environment does not interact with and collapse the superposition of the quantum state. The few exceptions are readout resonators, which rely on a strong coupling ($Q_c \sim 10^3$) to rapidly extract information about the resonator state which is used to infer a qubit's state (see Ch. 2.3.3), and first-stage quantum-limited amplifiers, whose efficiency is in part determined by how much of the signal propagates to the amplifier mode. Even in these cases, however, $Q_c \gg 1$ meaning that these devices are not overdamped; only a small fraction of energy is lost through the control lines per cycle. Under these conditions, the electromagnetic fields of the modes are well-localized within the internal environment.

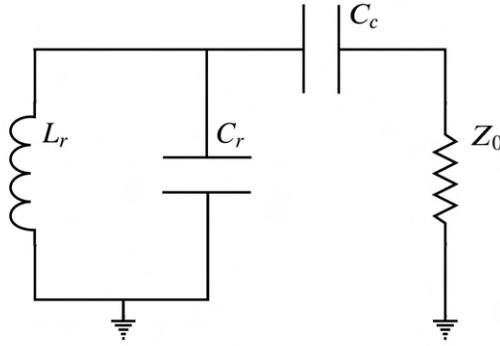


Figure 3.2: **LC resonator coupled to an external resistor.** The resistor represents a source impedance of a transmission line. assuming the transmission line is matched to the source impedance, the total input impedance as seen from the resonator is just Z_0 , which is equivalent to this circuit.

Control lines are generally microwave transmission lines (although in some devices that require the application of DC current, these lines may carry DC or low-frequency signals as well) that carry signals from room-temperature signal generators or to room-temperature signal acquisition electronics (analog-to-digital converters, or ADCs). This equipment is often terminated with a $50\ \Omega$ real impedance Z_0 and is impedance-matched to the transmission line; therefore, energy that leaks into the transmission line will eventually be dissipated by the terminating impedance.

To see how the coupling to a transmission line impacts the quality factor, we can calculate the Q_c of a basic LC resonator capacitively coupled to a transmission line that is matched with its source impedance; in this situation, the transmission line can be omitted (see Eq. (2.44) in Pozar [59]), and the circuit becomes equivalent to what is shown in Fig. 3.2. The impedance of this circuit is

$$Y(\omega) = i\omega C_r + \frac{1}{i\omega L_r} + \frac{i\omega C_c}{1 + i\omega Z_0 C_c} \quad (3.4)$$

The resonance frequency of this circuit is given by the zero of the imaginary part of $Y(\omega)$:

$$\text{Im}Y(\omega) = \frac{\omega^2 L_r C_r - 1}{\omega L_r} + \frac{\omega C_c}{1 + \omega^2 Z_0^2 C_c^2} = 0 \quad (3.5)$$

in the weak coupling limit, $\omega^2 Z_0^2 C_c^2 \ll 1$, so the resonance frequency becomes approximately

$$\omega_r \approx \frac{1}{\sqrt{L_r(C_r + C_c)}} \quad (3.6)$$

The dependence on C_c reflects the capacitive interaction with the transmission line. Assuming for the time being that there are no other losses in the system, we can calculate the total $Q = Q_c$ [60]:

$$Q_c = \frac{\omega_r \text{Im}Y'(\omega_r)}{2 \text{Re}Y(\omega_r)} \approx \frac{1}{Z\omega_r^2 Z_0 C_c^2} \quad (3.7)$$

where $\text{Im}Y' = \frac{d}{d\omega}(\text{Im}Y)$, and $Z = \sqrt{L_r/(C_r + C_c)}$ is the characteristic impedance of the mode. This implies that increasing the mode's characteristic impedance or the mode's frequency reduces the Q_c . The former amounts to reducing the ratio C_r/C_c , resulting in a higher voltage dropped across the coupling capacitor. The latter reduces the impedance across the coupling capacitor, resulting in higher current flow towards the resistor. Regardless, the common way to tune external couplings is to alter C_c .

Coupling capacitance is experimentally realized differently in planar versus the 3D cavity or coaxial tunnel architectures (see Fig. 3.3). In planar circuits, transmission lines from the PCB are wire-bonded to impedance-matched on-chip transmission lines. Interdigitated or planar capacitances are used to tune Q_c . In 3D or coaxial architectures, coupling is achieved through a coupling pin that continues the center pin of a coaxial cable. Inductive coupling is also possible and is shown in Fig. 3.3b,d. Of course, in practical

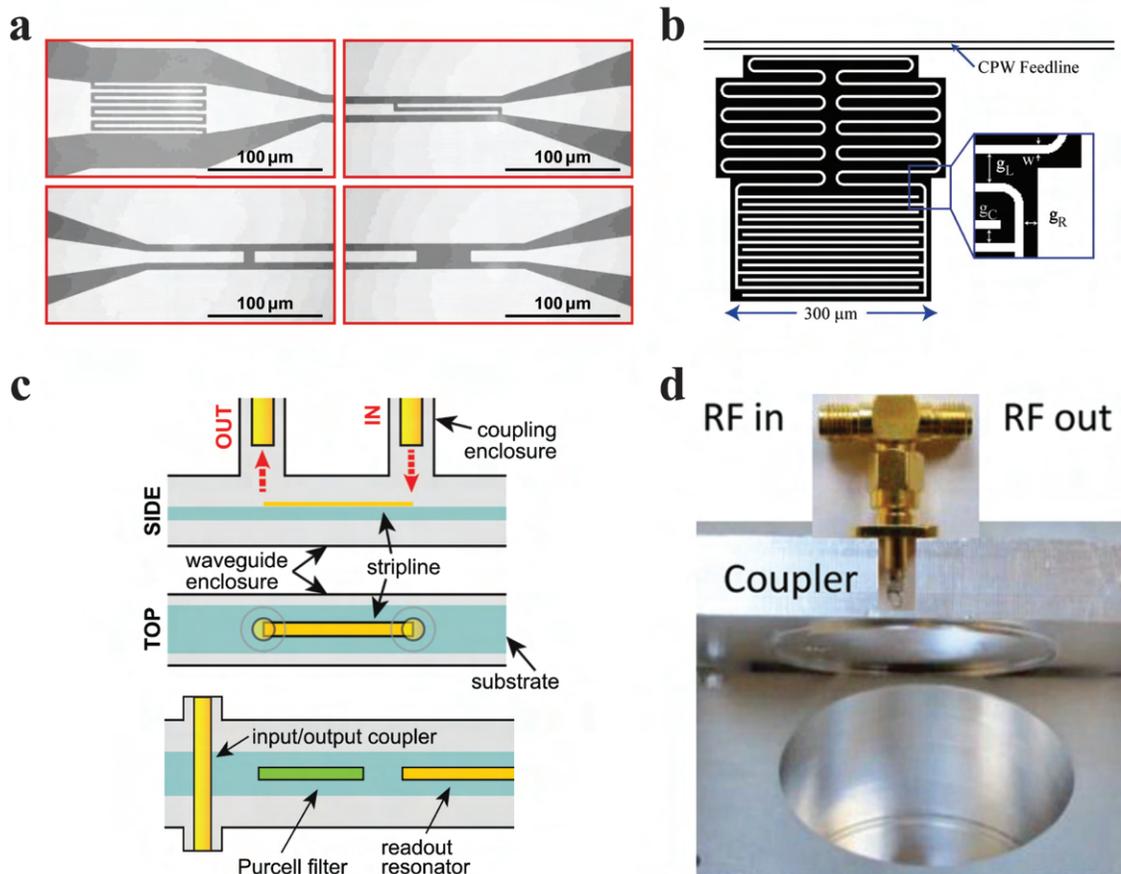


Figure 3.3: **External coupling schemes in cQED.** **a** External capacitive coupling in planar architectures. Capacitances are realized using interdigitated (upper left, upper right) or planar capacitors (lower left, lower right). Figure adapted from Göppl et al. [88]. **b** Inductive coupling in planar architectures. A lumped-element meandering inductor is placed close to the CPW transmission line such that their currents run in antiparallel, allowing the feedline magnetic field to induce a current in the inductor and excite the mode. Figure adapted from Geerlings et al. [89]. **c** Capacitive coupling in 3D and coaxial architectures. A coupling pin protruding from an external coaxial connector (top) extends into the package; the fields of the coupling pin are excited in a cutoff waveguide, resulting in their rapid attenuation in the direction towards the stripline resonator. Coupling can also be done using a transverse feedline (bottom), where a stripline Purcell filter is used to enhance readout κ while protecting the transmon (not shown) from external coupling loss. Figure adapted from Axline et al. [84]. **d** Inductive coupling in a 3D architecture. The coupling pin is bent backwards and connected to the ground, creating a loop coupler. Driving the coupler induces a magnetic field that couples to the magnetic field of the mode. Figure adapted from Reagor et al. [90].

quantum circuits the simple description of a single LC resonator with a coupling capacitor is not enough to fully describe the external coupling of the circuit. Doing so using finite-element simulation is quite straightforward; a $50\ \Omega$ resistive boundary is specified at the end of a $50\ \Omega$ transmission line that then couples to the rest of the circuit. By simulating the transmission line with the circuit, the external coupling can be estimated.

3.2.1 The Purcell Effect

The wide range of couplings required for different elements in a quantum circuit can lead to issues with regards to excess unwanted couplings in some elements. An important example of this is in the requirement for a qubit to be undercoupled ($Q_c \approx 4 \times 10^7$ or higher) to all control lines while its readout resonator is overcoupled ($Q_c \sim 10^3$) to the readout control line. Since the qubit is also coupled to the readout resonator, the qubit will inherit some loss through dissipation into the readout line (Fig. 3.4a). This enhancement of emission due to a nearby resonance mode is known as the Purcell effect[70, 87, 91]. In the language of electromagnetism, this can be described by the hybridization of the transmon mode with the readout mode, resulting in some field participation in the coupling capacitor C_c . Alternatively, in the language of microwave engineering, the resonator changes the transmon's admittance to the transmission line. Because this is a linear effect, it can be calculated using the same methods as before with Eq. (3.7). However, it should be emphasized that this effect is present for *every* mode that couples to the transmon, including environmental package modes. Fortunately, for more complex distributed-element networks, finite-element simulation can capture this effect, as the mode hybridization is simply how the diagonalized eigenfields manifest. As a result, the Purcell limit for a transmon or resonator embedded in an arbitrary electromagnetic environment is already included in the simulation of Q_c .

For the basic circuit shown in Fig. 3.4a, the Purcell relaxation rate γ_p is given by[70,

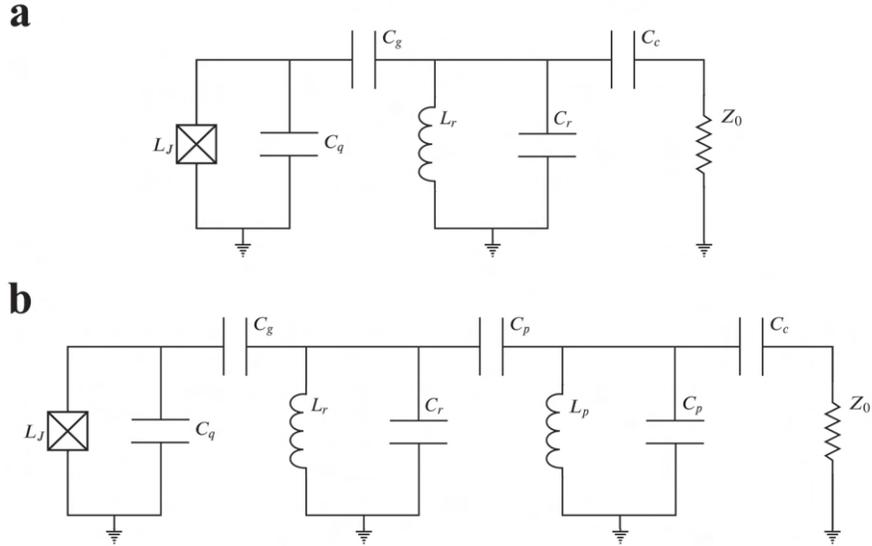


Figure 3.4: **Purcell effect in a qubit coupled to a readout resonator.** **a** The qubit inherits a coupling to the transmission line due to its coupling with the readout resonator, resulting in the enhancement of spontaneous emission into the external environment. **b** A bandpass Purcell filter suppresses the qubit's coupling to the control line while enhancing the readout resonator's coupling, effectively preserving the qubit's relaxation time.

84]

$$\gamma_p \approx \left(\frac{g_{rq}}{\Delta_{rq}} \right)^2 \kappa_r \quad (3.8)$$

where g_{rq} is the linear coupling constant between the transmon and readout resonator, defined in Eq. (2.37), $\Delta_{rq} = \omega_r - \omega_q$ is the detuning between the transmon and resonator, and $\kappa_r = \omega_r/Q_{r,c}$ is the decay rate of the resonator, which is dominated by its coupling quality factor $Q_{r,c}$. For $g_{rq}/\Delta_{rq} = 0.05$ and $\kappa_r/2\pi = 1$ MHz, we arrive at a Purcell relaxation rate of $\gamma_p \approx 1/64 \mu\text{s}^{-1}$. With state-of-the-art transmons having relaxation rates exceeding $400 \mu\text{s}$, this would be a dominating loss channel.

To alleviate this issue, we introduce a Purcell filter between the readout resonator and the control line, as shown in Fig. 3.4. The Purcell filter acts as an impedance transformer; the qubit sees a reduced admittance to the control line. The readout resonator and qubit

now see a filtered environment, suppressing the qubit's emission while preserving the readout resonator's strong coupling. To accomplish this, the Purcell filter is engineered to be close to the readout resonator in frequency. Additionally, the filter is strongly coupled to the environment ($Q_c = 10^2 - 10^3$), giving it a broad 10 – 100 MHz passband to accommodate the readout resonator and its dispersive shifts. An example of a Purcell filter in the coaxial architecture can be seen in Fig. 3.3c (bottom). The modified Purcell decay rate of the qubit is now

$$\gamma_p \approx \left(\frac{g_{rq}}{\Delta_{rq}} \right)^2 \left(\frac{g_{pr}}{\Delta_{rq} + \Delta_{pr}} \right)^2 \kappa_p \quad (3.9)$$

where g_{pr} is the coupling between the readout resonator and the Purcell filter, $\Delta_{pr} = \omega_p - \omega_r$ is the detuning between the Purcell filter and the readout resonator, and κ_p is the decay rate of the Purcell filter. The Purcell filter also modifies the decay rate of the readout resonator by $\kappa_r \approx (g_{pr}/\Delta_{pr})^2 \kappa_p$. Typically, $\Delta_{pr}/2\pi \sim 100$ MHz, which is much smaller than $\Delta_{rq}/2\pi \sim 2$ GHz; therefore, $\gamma_p \sim 1/\Delta_{rq}^4$, a much larger suppression of the Purcell effect. With typical parameters for $g_{pr} = 40$ MHz, $\Delta_{pr} = 100$ MHz, and $\kappa_p/2\pi = 10$ MHz the modified Purcell relaxation rate becomes $\gamma_p \approx 1 \text{ s}^{-1}$, and the readout loss rate becomes $\kappa_r \approx 1.6$ MHz. As a result, the Purcell loss of the transmon is now insignificant, while the readout coupling is preserved if not enhanced slightly.

This filtering approach is convenient because all of these parameters are tunable by modifying the design of the circuit. This allows tremendous flexibility with respect to achievable κ_r and γ_p . The circuit can be engineered to use the readout drive line for qubit readout as well as qubit control by tuning γ_p to be $\sim 1 \text{ ms}^{-1}$ while preserving $\kappa_r \approx 1$ MHz. This allows more efficient driving of the qubit while maintaining its coherence. This is precisely what is done for the devices measured in Ch. 5 to enable multiplexed qubit control and readout.

A regime exists called the resonant filter regime, where the Purcell filter is resonant with the readout resonator. In this case, both modes couple equally strongly with the qubit, effectively giving two readout modes, which is undesirable due to the potential qubit dephasing caused by the additional mode due to its thermal photon shot noise (see Sec. 3.4.2). It is best to avoid this case by ensuring that the filter and readout modes do not strongly hybridize, a condition that is achieved when $g_{pr} < \Delta_{pr}$.

The existence of the Purcell effect also implies that another source of loss is present in ultra high-Q resonance modes, such as those used for quantum information storage. The logical qubit encoded in a quantum memory can also spontaneously emit, and if it is coupled to a relatively lossy mode, this emission leads to an “inverse” Purcell effect[53]. This can happen in the case where a transmon with a particularly low T_1 is coupled strongly to a storage resonator of particularly high Q. In this case, the storage mode is Purcell limited by $\gamma_s \approx (g_{sq}/\Delta_{sq})^2\Gamma_1$, where γ_s is the inverse-Purcell decay rate of the storage mode, g_{sq} is the coupling between the storage mode and the transmon, and Δ_{sq} is the detuning between the storage mode and the transmon. For a typical value of $g_{sq}/\Delta_{sq} = 0.1$, the storage mode can be significantly Purcell-limited if the storage mode’s decay time is over 100 times larger than the transmon’s relaxation time.

Finally, I want to illustrate that the above circuit models are used for instructive purposes and do not necessarily describe the external couplings of a practical quantum circuit. This is due to the presence of additional modes such as package modes and even higher order harmonics of the readout and Purcell filter modes that modify the filtering behavior and the qubit’s Q_c . Proper calculation of the Purcell limit for an arbitrary electromagnetic environment can be done using the “BBQ” approach, where the admittance parameter is calculated from the perspective of the qubit. In a finite-element simulation where the complete environment (including the coupling to a 50 Ω transmission line) is defined, $Y(\omega)$ contains all the information about how the qubit or other modes of the cir-

circuit interact with each other and with the environment's stray modes. The Purcell loss rate is then

$$\gamma_p = 2 \frac{\text{Re}Y(\omega_r)}{\text{Im}Y'(\omega_r)}. \quad (3.10)$$

3.3 Internal Losses

While external losses are entirely determined by circuit design and microwave engineering, internal losses have both a geometric and material component. In this section, I will review the various sources of energy loss in superconducting quantum circuits.

3.3.1 Conductor Loss

The basis of cQED lies in superconductivity. Superconductivity is a quantum phenomenon that manifests when some metals are cooled to extremely low temperatures. At these temperatures, thermal fluctuations are low enough that high-energy phonon scattering is forbidden, and a net attractive force develops between electrons of same momentum and opposite spin due to a virtual phonon interaction mediated by the lattice[30, 33, 92]. This leads to the formation of Cooper pairs with a large superconducting band gap Δ , allowing them to flow with zero resistance. However, unpaired electrons, called quasiparticles, can still flow with finite resistance. Additionally, ambient magnetic fields present at the point of transition (also known as a cooling magnetic field) to the superconducting state can stimulate the formation of superconducting vortices, or islands of normal metal through which magnetic fields are threaded. The motion of these vortices under an AC electromagnetic field results in vortex flow dissipation. Together, the presence of quasiparticles and vortices yield an effective surface resistance or R_s of the superconductor, which leads to conductor loss $1/Q_{\text{cond}}$.

We can calculate Q_{cond} for an arbitrary electromagnetic mode by calculating the total power dissipated by the surface resistance[31, 59, 83]:

$$\frac{1}{Q_{\text{cond}}} = \frac{P_{\text{diss}}}{\omega_r U_{\text{tot}}} = \frac{R_s \frac{1}{2} \int_s |\vec{J}_s|^2 ds}{\omega_r \frac{1}{2} \int_{\text{all}} \mu_0 |\vec{H}|^2 dv} = \frac{R_s \int_s |\vec{H}^{\parallel}|^2 ds}{\omega_r \int_{\text{all}} \mu_0 |\vec{H}|^2 dv} \quad (3.11)$$

Where the integration in the numerator is done over the surface s of the superconductor and the integration in the denominator is done over the volume v of the mode. In the perfect conductor approximation, there is no magnetic field inside the conductor, and therefore the boundary condition $\vec{J}_s = \hat{n} \times \vec{H} = \vec{H}^{\parallel}$ [59] can be used to express the integral in terms of magnetic field. By rearranging some terms we arrive at

$$\frac{1}{Q_{\text{cond}}} = R_s \times \frac{\int_s \mu_0 |\vec{H}^{\parallel}|^2 ds}{\mu_0 \omega_r \int_{\text{all}} \mu_0 |\vec{H}|^2 dv} = \frac{R_s}{G} \quad (3.12)$$

where the numerator is now a “surface” magnetic energy. Here, we have separated the intrinsic material loss term R_s from the geometric component, the ratio of integrals which we call the geometric factor G , which is commonly used in the superconducting radiofrequency (SRF) cavity community[77, 93] and can be calculated analytically for simple geometries or by using finite element simulation. It is important here to highlight that R_s and G both have units of ohms; G is not a ratio of energies and therefore is technically not a participation ratio.

We can express the conductor loss in another way:

$$\frac{1}{Q_{\text{cond}}} = \frac{R_s}{\mu_0 \omega_r \lambda} \frac{\lambda \int_s \mu_0 |\vec{H}^{\parallel}|^2 ds}{\int_{\text{all}} \mu_0 |\vec{H}|^2 dv} = \frac{p_{\text{cond}}}{q_{\text{cond}}} \quad (3.13)$$

where $q_{\text{cond}} = 1/\Gamma_{\text{cond}} = X_s(\omega, T = 0)/R_s = \mu_0 \omega_r \lambda / R_s$ is the conductor quality factor. $X_s(\omega, T = 0)$ [94] is the surface reactance of the superconductor at zero temperature and is analogous to the reactance of an inductor. The conductor participation $p_{\text{cond}} = \mu_0 \omega_r \lambda / G$

is the corresponding energy participation ratio, which is now dimensionless. p_{cond} is also called the kinetic inductance ratio $\alpha = L_k/(L_k + L_g)$ of the mode, where L_k is the kinetic inductance of the Cooper pairs and L_g is the geometric inductance that measures how much energy is stored in the magnetic field.

I have also introduced a term λ to Eq. (3.13) which is known as the magnetic penetration depth. The penetration depth is the characteristic length scale over which magnetic fields attenuate inside the superconductor due to the presence of screening supercurrents (the Meissner effect). However, this screening behavior only applies when the thickness t of the superconductor is much larger than the penetration depth; $t \gg \lambda$. In thin films, $t \sim \lambda$, and the induced supercurrents cannot fully attenuate the magnetic field; as a result, the film is completely penetrated[95]. Importantly, since p_{cond} is now dependent on λ , which is a material parameter, the conductor participation is not a purely geometric quantity. For this reason, it is often times more convenient to describe conductor participation using G ; although it is not dimensionless, it is exclusively a geometric quantity.

In normal conductors, λ is replaced by the skin depth δ ; however, the surface resistance and reactance both have the same dependence on skin depth. As a result, for normal conductors $q_{\text{cond}} = 1$ [59]. This limits the Q_{int} of even a cavity resonator made with normal conductor to $\sim 10^3$. For superconductors, $q_{\text{cond}} \gg 1$ due to the exceptionally small R_s . Additionally, the penetration depth $\lambda \approx 50 - 100$ nm is normally around an order of magnitude smaller than the skin depth of a normal conductor, resulting in significantly reduced conductor participation. This allows superconducting microwave resonators to achieve quality factors in the range of $10^6 - 10^9$, although depending on geometry and materials used, conductor loss due to residual resistance can still be limiting. Thin-film resonators, for example, have $p_{\text{cond}} \approx 10^{-2} - 10^{-3}$ [91] due to their relatively high current densities on their surfaces, while 3D post cavities have $p_{\text{cond}} \approx 3 \times 10^{-5}$ [31] due to their significantly more dilute magnetic fields. However, it seems to be the case that thin films

have much lower residual resistance than bulk superconductors, as we will see in Ch. 4. In the following subsections, I will describe how quasiparticle populations and vortices can give rise to surface resistance and therefore a finite q_{cond} .

Thermal Quasiparticles

Thermal quasiparticles arise from thermal fluctuations in the environment that can break Cooper pairs. this results in an equilibrium population of quasiparticles, and can be minimized if the temperature is low. The complex conductivity $\sigma(\omega, T) = \sigma_1(\omega, T) - i\sigma_2(\omega, T)$ of a superconductor is described by the Mattis-Bardeen equations and derived from Bardeen-Cooper-Schrieffer (BCS) theory[94, 95]. It is calculated by integrating the distribution function for unpaired electrons (which obey Fermi-Dirac statistics) weighted by a temperature-dependent superconducting bandgap $\Delta(T)$, which approaches a limiting value Δ_0 when $T = 0$. The complex conductivity gives rise to a complex surface impedance[94]:

$$Z_s(\omega, T) = \sqrt{\frac{i\mu_0\omega}{\sigma(\omega, T)}} = R_s(\omega, T) + iX_s(\omega, T) \quad (3.14)$$

where μ_0 is the permeability of free space and X_s is the surface reactance of the superconductor. When $T \ll T_c$ and $\hbar\omega \ll \Delta_0$, the surface reactance reduces to the zero-temperature limit $X_s(\omega, 0) = \mu_0\omega\lambda$ as mentioned earlier, and the surface resistance is calculated to be[93]

$$R_s = \mu_0^2\omega^2\lambda^3\sigma_1(\omega) = \frac{2\mu_0^2\omega^2\lambda^3\Delta_0}{\rho_n\hbar\omega} \sinh\left(\frac{\hbar\omega}{2k_B T}\right) K_0\left(\frac{\hbar\omega}{2k_B T}\right) e^{-\Delta_0/k_B T} \quad (3.15)$$

where ρ_n is the normal-state resistivity of the superconductor, and $K_0(x)$ is the modified Bessel function of the second kind. The important thing to note about this expression

is that $R_s \propto e^{-\Delta_0/k_B T} \approx e^{-1.764T_c/T}$ ($\Delta_0 = \pi e^{-\gamma} k_B T_c \approx 1.764 k_B T_c$ in the ideal case only; it is quite accurate for aluminum, but for niobium, for example, $\Delta_0 \approx 1.9 k_B T_c$ [93]); therefore, as long as the operating temperature is sufficiently low, the surface resistance should be negligible. For an aluminum superconductor with $T_c = 1.2$ K, operating temperature $T = 20$ mK, $\lambda = 50$ nm, $\omega/2\pi = 5$ GHz, and $\rho_n = 2.65 \mu\Omega\text{cm}$, the surface resistance becomes $\sim 10^{-50}$, which is an unimaginably small number. In other words, at these temperatures it is likely that only 0 – 1 thermal quasiparticles exist in this regime. Materials with higher T_c may have other beneficial properties with regards to minimizing other sources of loss, but they should not be chosen for their ability to suppress thermal quasiparticle populations. Regardless, there are many viable superconductors for cQED, such as aluminum (Al), niobium (Nb), niobium nitride (NbN), tantalum (Ta), indium (In), titanium nitride (TiN), and rhenium (Re), all of which have $T_c > 1$ K. Importantly, however, the superconducting properties such as the transition temperature and magnetic penetration depth are dependent on fabrication processes; the deposition conditions for thin films affect how they grow, which will affect their ability to conduct supercurrent and result in an excess population of quasiparticles. These effects are not captured with BCS theory, but may result in a residual “nonequilibrium” quasiparticle population.

Conductor loss is the only loss channel whose participation ratio can be measured in experiment rather than needing to rely on simulations or analytical calculations. This is possible due to the temperature dependence of the surface impedance[31, 94, 95]. The derivation of this temperature dependence is quite complicated but is detailed very nicely in Lev Krayzman’s thesis[83]. An intuitive explanation is as follows: the temperature dependence of surface reactance implies that the penetration depth is temperature-dependent; $X_s(T) = \mu_0 \omega \lambda(T)$. In the low-dissipation limit, the frequency of the resonator is determined entirely by the capacitance and inductance, $L_r = L_g + L_k$. Due to the temperature dependence of λ , the kinetic inductance and therefore the frequency

$2\pi f = 1/\sqrt{C_r(L_g + L_k)}$ also has a temperature dependence:

$$\frac{\delta f}{f} = -\frac{1}{2} \frac{\delta L_k}{L_g + L_k} = -\frac{\alpha \delta L_k}{2 L_k} = -\frac{\alpha \delta \lambda}{2 \lambda} = -\frac{\alpha \delta X_s(T)}{2 X_s(0)} \quad (3.16)$$

As the temperature increases, the number of Cooper pairs decreases. As a result, the average kinetic energy per Cooper pair increases in order to sustain the same current. As a result, the kinetic inductance increases, resulting in a drop in frequency. The kinetic inductance ratio α at zero temperature appears as a proportionality constant; as a result, fitting the frequency dependence of the resonator with the Mattis-Bardeen equations for surface reactance allows one to extract α . X_s is also implicitly dependent on Δ_0 ; therefore, a fit to this data allows one to extract the T_c of the superconductor as well. An example of the temperature dependence of frequency is given in Fig. 3.5a. The quality factor of the resonator also follows a similar trend and can be used to extract the same information[83, 95]:

$$\delta \frac{1}{Q} = \frac{1}{Q(T)} - \frac{1}{Q(0)} \approx \alpha \frac{R_s(T) - R_s(0)}{X_s(0)} = \alpha \frac{R_s(T)}{X_s(0)} \quad (3.17)$$

where $R_s(0) = 0$ from Eq. (3.15) and $X_s(T) \approx X_s(0)$ is assumed in the case of $\delta f/f \ll 1$ [90, 93, 95]. In practice, the resonator is limited by other loss mechanisms not predicted by BCS theory; as a result, fitting the temperature dependence of Q often requires the addition of a constant zero-temperature Q_0 which effectively causes the Q to plateau at low temperatures (see Fig. 3.5b).

The usefulness of determining α from measurement should not be understated. The geometric ratio can be simulated for 3D cavities because we can safely use the assumption that the magnetic field that impinges on the superconducting surface attenuates exponentially; as a result, measuring α allows one to determine λ , which is highly dependent on the superconductor and the processes used to fabricate or clean the resonator. λ for

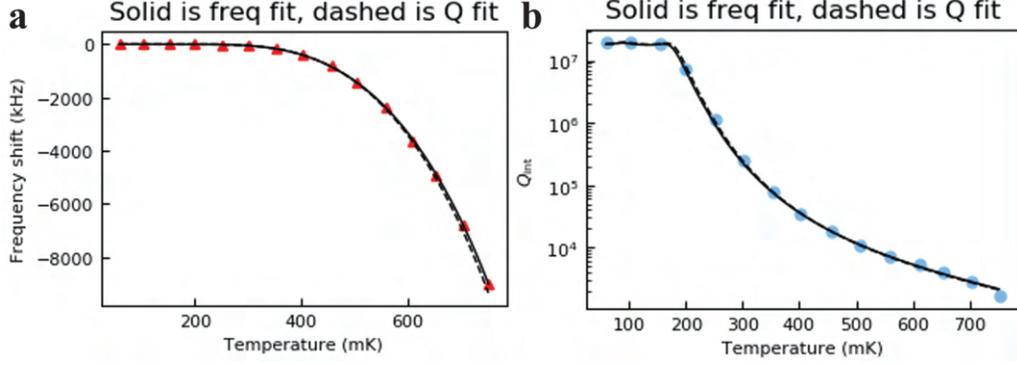


Figure 3.5: **Temperature dependence of resonator frequency and Q due to thermal quasiparticles.** **a** Temperature dependence of frequency of a thin-film (90 nm) aluminum resonator and BCS fit (solid line). The extracted fit parameters are $T_c = 1.26$ K and $\alpha = 0.021$. Data from quality factor fit reproduces the temperature dependence on frequency (dashed line). **b** Temperature dependence of the resonator's quality factor and BCS fit (solid line). The extracted fit parameters are $T_c = 1.31$ K and $\alpha = 0.025$. Data from frequency fit reproduces the temperature dependence on quality factor (dashed line). The fit parameters from both fits are slightly different, likely due to poor resolution of Q_{int} at high temperatures ($Q_c \gg Q_{\text{int}}$; see Sec. 3.5.2), and the temperature range not extending closer to T_c . Since at high temperatures the frequency changes more drastically, data taken at temperatures closer to T_c can better constrain the fit.

bulk-machined high-purity aluminum has been measured to change from ≈ 65 nm to ≈ 50 nm after an acid-based etching treatment has been performed to remove contaminants and machining damage[90]. Relatively thick ($\approx 1 \mu\text{m}$) thermally-evaporated indium films in micromachined cavities have been measured to have $\lambda \approx 30 - 60$ nm[82]. Measurements of λ for thinner films becomes more difficult because the magnetic field fully penetrates a thin-film superconductor; therefore, the expression for G which only considers the surface magnetic field does not capture full field behavior. The extracted λ then reflects the field within the film and therefore picks up a geometric dependence on the film thickness; in some cases, it can be even larger than the film thickness. An example of the temperature dependence of a thin-film (This particular film is a trilayer of 30 nmAl/2 nm AlO_x/60 nmAl) aluminum resonator is in Fig. 3.5, where the simulated geometric ratio is $G = 0.34 \Omega$. Using the extracted α from either the frequency fit or the

Q fit, we can extract $\lambda \approx 110 - 130$ nm, which is larger than the film's thickness and over two times larger than what was extracted for bulk Al in Reagor et al. [90]. In a sense, this implies that λ is heavily dependent on the materials used and the fabrication processes involved—bulk Al is not the same as thin-film Al, and 90 nm-thick Al might not be the same as 50 nm-thick (or 20 nmAl/2 nm AlO_x/30 nmAl) Al.

Nonequilibrium Quasiparticles

Nonequilibrium quasiparticles refer to an excess quasiparticle population that is not described by a thermal population. This is described by a quasiparticle fraction $x_{qp} = n_{qp}/n_{cp}$ which is the ratio of the quasiparticle density n_{qp} to the Cooper pair density n_{cp} . The quasiparticle population is directly related to the surface resistance[95]:

$$R_s = \frac{2\mu_0^2\omega^2\lambda^3\Delta_0^2}{\rho_n\hbar\omega\sqrt{2\pi k_B T}\Delta_0} \sinh\left(\frac{\hbar\omega}{2k_B T}\right) K_0\left(\frac{\hbar\omega}{2k_B T}\right) x_{qp} \quad (3.18)$$

For a thermal population of quasiparticles, $x_{qp} \propto e^{-\Delta_0/k_B T}$ which reproduces Eq. (3.15). The prefactor that relates R_s and x_{qp} is related to the normal-state and superconducting-state properties of the superconductor that makes up the circuit and is of order $\sim 10^{-3}$. Therefore, knowledge of x_{qp} determines Γ_{cond} due to excess quasiparticles.

An excess quasiparticle population results in a residual resistance R_{res} that contributes to the total surface resistance, $R_s = R_{\text{res}} + R_{\text{th}}$, where R_{th} is the surface resistance due to thermal quasiparticles (which is negligible for our choice of superconductors and operating temperatures). Excess quasiparticles may exist endogenously in the superconductor due to defects or nonidealities such as surface inhomogeneity of the BCS potential or Cooper pair scattering due to grain boundaries or impurities such as molecular hydrogen, carbon, and oxygen (and others; see Gurevich [93]). These mechanisms are thought to effectively broaden the density of states near the gap, giving rise to a residual sur-

face resistance. Unfortunately, evidence for how these phenomena translate to loss in the microwave regime has been inconclusive[93], partly due to the difficulty in isolating conductor loss from other sources of loss such as dielectric loss. Excess quasiparticles can also be generated by above-gap radiation[95–98], where an incident photon has enough energy to break a Cooper pair into two quasiparticles, or by cosmic rays[99, 100] that can collide with the substrate, releasing high energy phonons that are absorbed by the superconductor. These quasiparticles are “hot”, and will eventually relax and recombine; however, with a steady-state flux of high energy photons or phonons, an equilibrium will develop between the quasiparticle generation and recombination rates, resulting in an excess x_{qp} . Both pair-breaking processes and endogenous quasiparticle populations have been observed[96–98] by measuring quasiparticle tunneling through Josephson junctions in transmon qubits. Additionally, it has been demonstrated that by suppressing quasiparticle generation by high-energy radiation through improved radiation shielding, RF filtering, and thermalization techniques, the excess quasiparticle populations can be reduced in Al-based transmon qubits to as low as $x_{qp} \sim 10^{-10}$ [98].

Vortices

Residual ambient magnetic fields at the point of superconducting transition can lead to the formation of vortices, which are islands of trapped flux that are normally conducting. In an ideal type 1 superconductor, the Meissner effect leads to complete expulsion of the cooling magnetic field B_0 within the superconductor as long as the cooling field does not exceed the critical magnetic field B_c ; no vortices can form. In a type 2 superconductor, two regimes exist that are bounded by B_{c1} and B_{c2} , where $B_{c1} < B_{c2}$. If the cooling field is below B_{c1} , a complete Meissner effect occurs; however, if $B_{c1} < B_0 < B_{c2}$, a partial Meissner effect occurs where the magnetic flux is not necessarily expelled but pinned within vortices of radius equal to the coherence length ξ [33, 92, 93]. With this description,

only type 2 superconductors like niobium or the superconducting nitrides can host vortices. However, it has been shown that type 1 superconductors can have vortices of potentially macroscopic size[101]. Moreover, thin film type 1 superconductors with $t \sim \lambda$ are fully penetrated by magnetic fields, and manifest type 2 behavior[102]. Under an AC field, these vortices can move, resulting in vortex flow dissipation, where an effective current flows due to normal charge carriers. In a type 1 superconductor, the microscopically-sized vortices can be penetrated by impinging magnetic fields. This results in a vortex flow resistance[101]:

$$R_{\text{vortices}} = \frac{B_0}{B_c} \left(\frac{\mu_0 \rho_n \omega}{2} \right)^{1/2} \quad (3.19)$$

where $B_c = h/(2^{5/2} \pi \lambda \xi)$ is the critical magnetic field (B_{c2} for type 2 superconductors). In practice a Cryoperm magnetic shield is used to suppress the cooling field experienced by the quantum circuit. Room-temperature measurements using a Hall probe have shown that the residual magnetic field inside the shield is $\sim 10^{-7}$ T; however, the permeability of Cryoperm shields increase change with temperature, thereby changing the cooling field at the point of transition to the superconducting state. Additionally, B_c is determined by the penetration depth and coherence length of the superconductor, which is material-dependent. A rough estimate for R_{vortices} assuming $B_0 = 10^{-7}$ T gives $R_{\text{vortices}} \approx 250$ n Ω for Al (assuming $\rho = 2.65$ $\mu\Omega\text{cm}$ for Al), $R_{\text{vortices}} \approx 70$ n Ω for Ta (assuming $\rho = 14.2$ $\mu\Omega\text{cm}$ for Ta), and $R_{\text{vortices}} \approx 30$ n Ω for Nb (assuming $\rho = 15.0$ $\mu\Omega\text{cm}$ for Nb). The higher critical field of tantalum results in a greater resiliency to vortex formation; niobium performs even better. It should be noted, however, that the room-temperature resistivities were used here in this calculation. In practice, the cryogenic resistance just before transition should be used, which would decrease R_{vortices} . Additionally, impurities or grain boundaries can reduce the mean free path l of the superconductor, which will

reduce the effective coherence length and increase the effective penetration depth, thereby changing B_c by a small amount. A more informed estimation of R_{vortices} will be made in Ch. 4 for Al and Ta after measuring some of their material properties.

Finally, vortices have been indirectly observed in aluminum transmon qubits (thin-film superconductors) and 3D cavities (made of either aluminum[101] or niobium[93]). While they can contribute to dissipation, they can also trap excess quasiparticles, thereby *reducing* the total loss[102]. This occurs when quasiparticle loss dominates over vortex loss; if x_{qp} were to be reduced, the total conductor loss would then be more dominated by the remaining vortices. In the case of transmons, trapped quasiparticles cannot tunnel through the junction, thereby reducing relaxation rates.

3.3.2 Dielectric Loss

Dielectric loss is probably the most heavily explored loss channel in cQED. This is because there is substantial evidence that shows that surface dielectric loss dominates in thin-film devices. In general, the dielectric loss factor can be described by a loss tangent, $\Gamma_{\text{diel}} = \tan \delta_{\text{diel}} = \epsilon''/\epsilon'$, where ϵ' and ϵ'' are the real and imaginary parts, respectively, of the dielectric permittivity ϵ of the dielectric region in question. Dielectric participation can be significant in thin-film devices because they have to be fabricated on a crystalline substrate such as silicon or sapphire, both of which have been shown to have low loss tangents of $\approx 10^{-8} - 10^{-7}$ [103–106]. Additionally, surface dielectric losses in various interfaces between the superconductor, the substrate, and air (vacuum) can host defects and contaminants that introduce additional dielectric loss. These “surfaces” are usually considered to be thin (\sim nm) dielectric regions whose physical nature is largely influenced by the materials and fabrication processes used to make the circuit, and seem to have much larger loss tangents of $\approx 10^{-4} - 10^{-3}$, assuming a region thickness of 3 nm and dielectric constant of 10 [86, 105, 107–110]. 3D cavities generally have very little field overlap with a substrate

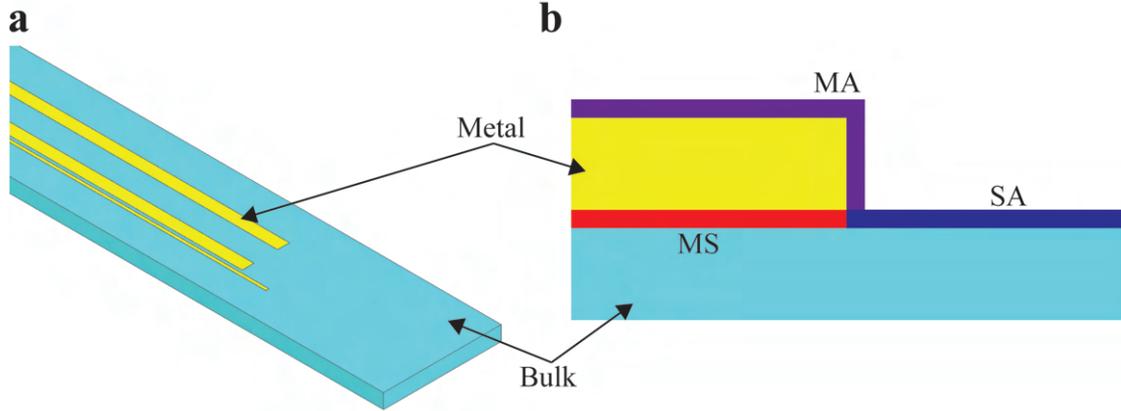


Figure 3.6: **Lossy dielectric regions.** **a** Example stripline resonator patterned on a substrate. **b** Cross-section of the film deposited on the substrate, showing the metal-substrate (MS) interface (red), the metal-air (MA) interface (purple), and the substrate-air (SA) interface (dark blue).

and as a result has low substrate participation. However, the oxide that grows on the surface of the superconductor is dielectric in nature; these oxides are amorphous and can be significantly lossy. For planar and coaxial architectures, since all useful modes are defined by thin-film circuits that are patterned on the substrate, all modes have significant dielectric participation. In 3D cavity architectures, the transmon has significant dielectric participation and the storage mode, due to its fields being mostly in vacuum, only acquires a small dielectric participation due to the insertion of the transmon-containing chip.

Dielectric loss is quantified with the participation ratio model using the following expression:

$$\frac{1}{Q_{\text{diel}}} = p_{\text{diel}} \Gamma_{\text{diel}} = p_{\text{diel}} \tan \delta_{\text{diel}} \quad (3.20)$$

For dielectric loss, four distinct regions are specified which are considered separately and are shown in Fig. 3.6. Bulk dielectric loss refers to the losses associated with the bulk of the substrate on which thin-film devices are fabricated. Common substrates used in cQED are silicon and sapphire; both of them are manufactured to be crystalline and with

low defect density; however, residual defects or dislocations in the crystal structure can introduce a small amount of loss. Bulk dielectric participation is calculated by integrating the electrical energy density over the volume of the substrate:

$$p_{\text{bulk}} = \frac{\int_{\text{bulk}} \epsilon |\vec{E}|^2 dv}{\int_{\text{all}} \epsilon |\vec{E}|^2 dv} \quad (3.21)$$

where ϵ is the real part of the dielectric permittivity of the substrate, and the denominator is the total electric energy of the mode.

Bulk dielectric loss is usually distinguished from surface dielectric loss, because we expect surface regions to be damaged, contaminated, or otherwise modified due to the various fabrication processes to which the wafer is subjected. We additionally distinguish surface dielectric regions by their location and boundaries (Fig. 3.6b): the metal-substrate (MS) interface, the metal-air (MA) interface, and the substrate-air (SA) interface. In practice the device is measured under vacuum, so “air” refers to vacuum. Loss in the MS region is attributed to surface residues, crystalline defects or vacancies, or materials incompatibility between the substrate and the superconductor that can lead to deformations in the substrate’s crystal structure, potentially leading to anomalous piezoelectricity. Loss in the MA region is mostly attributed a lossy oxide on the conductor that grows upon exposure to air; this oxide is amorphous in nature and may play host to a large number of defects and adsorbates from the atmosphere. Additionally, contaminants and residues left over from the fabrication process can add to the loss. Finally, the SA region is lossy due to the presence of left over contaminants, surface adsorbates, crystalline defects, or the presence of an oxide in the case of silicon.

Calculation of $p_{\text{SA,MS,MA}}$ is more complicated than it is for p_{bulk} , mostly due to the small assumed thickness of these regions and large aspect ratio of the thin films (i.e. the width/length of the film is hundreds of times greater than the thickness). Thin-film con-

ductors are approximated in a 3D electromagnetic simulation as perfectly conducting 2D sheets. Field behavior at the edges of the thin-films are approximated using a heavily meshed 2D cross-sectional electrostatic simulation (see Appendix C)[74] with explicitly-defined surface dielectric regions of assumed thickness $t_{\text{surf}} = 3$ nm and relative permittivity $\epsilon_r = 10$ to maintain consistency with other works[105, 110, 111]. It should be noted that the true thickness and relative permittivity of these regions are unknown. We can attempt to measure the thickness with nanometer-scale microscopy; however, this can only yield qualitative information about the interfaces and cannot definitively determine the dielectric properties or the presence or absence of physical signatures of loss. We therefore treat the true surface region thickness and relative permittivity as material/process parameters that re-scale the surface loss tangents and thereby redefine them as $\tan \delta_k = \frac{t_{k0}}{t_{\text{surf}}} \frac{\epsilon_{r0}}{\epsilon_r} \tan \delta_{0,k}$, where $\tan \delta_{0,k}$, t_{k0} , and ϵ_{r0} are the true dielectric loss tangent, thickness of the surface regions, and dielectric constant of region $k = \text{SA, MS, MA}$, respectively[112]. The surface dielectric participations are therefore given by the following integral expressions:

$$p_{\text{SA,MS}} = \frac{t_{\text{surf}} \int_{\text{SA,MS}} \epsilon_r \epsilon_0 |\vec{E}|^2 ds}{\int_{\text{all}} \epsilon |\vec{E}|^2 dv} \quad (3.22)$$

$$p_{\text{MA}}, p_{\text{pkGMA}} = \frac{t_{\text{surf}} \int_{\text{MA}} \epsilon_0 |\vec{E}_{\text{vac}}|^2 ds}{\epsilon_{r,\text{MA}} \int_{\text{all}} \epsilon |\vec{E}|^2 dv} \quad (3.23)$$

The expression for $p_{\text{SA,MS}}$ is slightly different from that of p_{MA} . To calculate $p_{\text{SA,MS}}$, we integrate over a surface located 3 nm below the 2D conducting sheet. Since the thickness of the surface regions are so small, we assume that the electric field is constant over the thickness; therefore, we can simply multiply the surface integral by the thickness to determine the energy stored in the surface. To calculate p_{MA} , the same approach is used, only the integration is done over a surface located 3 nm above the conducting sheet. However,

in this region the electric field is that of vacuum. Since the MA region is not explicitly included in the 3D simulation, the electric field has to be rescaled by the dielectric constant ϵ_r [31]. To do this, we use the continuity of the displacement field, $\epsilon_r \epsilon_0 E_{\text{MA},\perp} = \epsilon_0 E_{\text{vac},\perp}$. From this, we can see that $E_{\text{MA},\perp} = (1/\epsilon_r) E_{\text{vac},\perp}$, which results in a factor of ϵ_r in the denominator. Eq. (3.23) is used to calculate p_{MA} for both thin-film circuits and bulk superconductors.

There is a fifth dielectric region that so far has not been discussed: the Josephson junction. Since the tunnel barrier is an amorphous oxide, one would expect to be as lossy as a surface dielectric region. However, this does not seem to be the case, as measurements on Cooper-pair boxes, which have nearly 100% dielectric participation in the junction, have bound the loss tangent of the junction oxide to below 4×10^{-8} [38]. Since the transmon's capacitance is dominated by the shunting capacitor, its junction oxide participation is $\sim 2\%$; therefore, the junction oxide would limit the transmon to $Q \approx 10^9$, or T_1 of several tens of milliseconds. As a result, junction oxide losses are rarely considered.

The quantification of dielectric loss with the above definitions implicitly assumes certain things about the nature of the dielectric regions under study. For one, these participations as defined assume that the dielectric constant is spatially homogeneous. Modifications can be made to account for anisotropy; for example, sapphire does have a dielectric anisotropy; for this reason, we use C-plane sapphire, which is oriented such that the dielectric constant is equal along the lateral axes. The dielectric constant along the perpendicular axis is slightly different; in practice, however, this does not change the calculation of the bulk dielectric participation by more than $\approx 5\%$. Additionally, localized defects within the bulk dielectric could result in spatial inhomogeneity of ϵ . These effects are assumed to be small and effectively apportioned to the loss factor. For surface dielectric participations, the assumed thickness of 3 nm is done out of convention; the true thickness is treated as an intrinsic parameter of the formation of the interfacial region. While this allows us to tol-

erate small variations between t_{surf} and t_{k_0} , if t_{k_0} is too large, then the approximation that the electric field is constant in thickness for this region no longer holds. Thankfully, this assumption does approximately hold as long as the true thickness of the region is much smaller than the dimensions of the patterned film.

The definition of the dielectric loss factor also assumes frequency independence. It is unclear how much the loss factor changes with frequency; however, it may be possible to probe the frequency dependence using multimode resonators, which will be discussed in Ch. 4. For now, we will assume that the loss factors have no frequency dependence or very weak frequency dependence over the 3 – 10 GHz range. Independent measurements of the dielectric loss tangent of HEMEX-grade sapphire have shown similar numbers whether it is measured at ≈ 4.5 GHz[106] or ≈ 13 GHz[103], which indicates that this may be a good assumption. Another assumption made with substrate loss is that it is purely dielectric in origin; this is at first glance reasonable since the substrate has a large dielectric constant; however, magnetic impurities may give rise to an effective magnetic loss tangent as well. Estimates of magnetic loss show that it must be a subdominant, almost insignificant contributor to the total loss factor[106], which again makes the assumption valid.

In the following two subsections I will discuss the two main mechanisms behind dielectric loss: phonon loss and TLS loss.

Phonon Loss

Phonon loss refers to some form of energy relaxation through the vibrational modes of the dielectric. This requires some form electrostriction or piezoelectricity, where the electromagnetic modes can couple to acoustic modes; the excitations of which (phonons) can scatter and thermalize with the bath. In a model described by Dunne et al. [113], structural disorder in the crystal lattice introduces anharmonic interactions between vibrational modes of the lattice. The higher order nonlinear terms can then interact off-resonantly with

electromagnetic fields, resulting in an anomalous coupling that allows for the conversion of microwave photons into phonons, which thermalize with the bath, resulting in energy relaxation. In an amorphous dielectric such as what can be found in interface regions, the lattice structure is ill-defined; as a result, these vibrational modes are far more spread out in frequency, resulting in stronger coupling and therefore increased loss. This implies that greater crystalline order leads to lower loss, consistent with dielectric loss tangent measurements on different grades of sapphire[106]. It also implies that amorphous oxides that may be present in surface regions may be necessarily more lossy than the crystalline dielectric substrate on which the superconducting circuit is fabricated; i.e $\Gamma_{\text{diel}} < \Gamma_{\text{SA,MS,MA}}$.

TLS Loss

TLS loss is arguably a more pernicious and also more well-studied loss mechanism, due to its ubiquity in cQED. It refers to the presence of anomalous two-level systems in the environment of a quantum circuit; the electric fields of resonators and transmons can couple to the dipole moment of the TLS, resulting in additional decay channels. TLS loss is commonly thought to be the dominant source of loss in planar circuits, due to its presence in surface regions. A detailed description of TLSs is given in Jiansong Gao's thesis[95]; here, I will provide an intuitive explanation. Defects or vacancies in dielectric regions can manifest a double-well potential energy landscape, where tunneling from one well to another results in the emission of energy. These defects are thought to be localized in surface regions, since the bulk of the substrate is highly crystalline and thought to have very low defect density. The TLS model assumes that TLSs are distributed uniformly in space and frequency and can stimulate relaxation as well as dephasing. The TLS model is described by the following expression:

$$\frac{1}{Q_{\text{diel}}} = \frac{1}{Q_0} + \frac{p_{\text{TLS}} \tan \delta_{\text{diel}}}{\sqrt{1 + (\bar{n}/n_c)^\beta}} \tanh\left(\frac{\hbar\omega}{2k_B T}\right) \quad (3.24)$$

where $1/Q_0$ is the power-independent contribution to the total internal loss, p_{diel} is the participation of the particular dielectric region under study, $\tan \delta_{\text{TLS}}$ is the ensemble TLS loss tangent, \bar{n} is the average number of photons in the mode (determined by the input drive power and the excited electric field of the mode), n_c is the critical photon number beyond which the TLSs begin to saturate, and β is an empirical parameter that accounts for the spatially varying electric field of the resonance mode, which inhomogeneously saturates the TLSs (if the electric field is constant in the dielectric region, $\beta = 1$) [110, 114–118]. Importantly, the n_c is inversely related to the electric field-weighted mode volume and the intrinsic T_1 and T_2 of the TLSs. A device with lower surface dielectric participation therefore tends to have a higher n_c ; since the electric field amplitude per photon is smaller, more photons are required to saturate the TLSs to which the device is coupled. The average TLS T_1 follows a thermal distribution and is therefore temperature dependent; $T_1 \propto \tanh\left(\frac{\hbar\omega}{2k_B T}\right)$ [95, 110]. TLSs interacting with other TLSs result in their dephasing due to their thermally-fluctuating energy splittings; therefore, the TLS T_2 is also temperature dependent, $T_2^{-1} \propto k_B T$ (this linear dependence on temperature is derived in Gao [95]; however, it should be noted that Crowley et al. [110] modelled this dependence as $\propto T^{\beta_1}$, where β_1 is an empirical parameter). In our operating regime where $\hbar\omega \gg k_B T$, the TLS T_1 is effectively temperature-independent; however, T_2 is still temperature dependent, implying that n_c reduces with temperature.

Interactions with TLSs produce both a temperature and a power dependence in the resonator; the mechanism for this dependence is related to TLS saturation. TLSs can be saturated both by power and temperature. At low temperatures and low drive powers, the TLSs are maximally coupled to the resonance modes. As the drive power increases, the average photon number in the mode increases; above the critical photon number, the TLSs begin to saturate; they are constantly being driven between their two states by the large number of photons circulating in the mode. As a result, the TLSs effectively de-

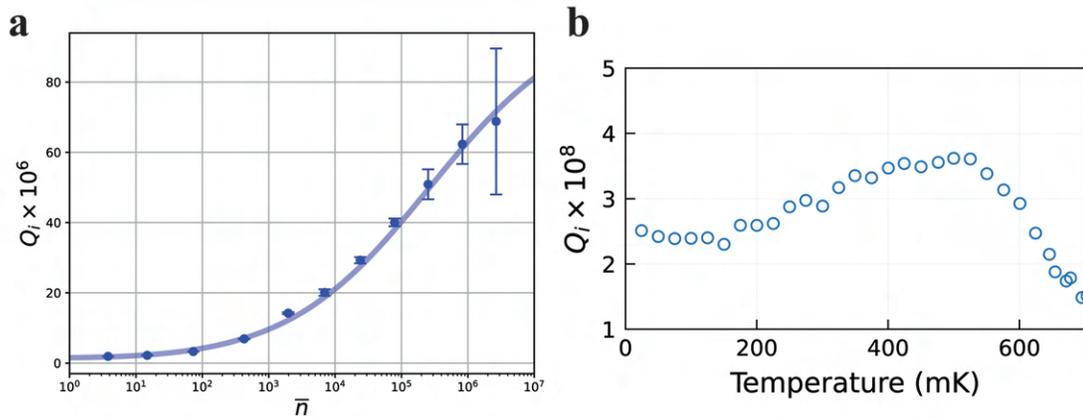


Figure 3.7: **Power and temperature dependence of TLSs.** **a** Power dependence of an aluminum thin-film stripline resonator (circles) and TLS fit (line). The TLSs saturate at high powers, causing Q_{int} to increase dramatically from $Q_{\text{int}} \approx 1.5 \times 10^6$ at $\bar{n} = 1$ to $Q_{\text{int}} \approx 60 \times 10^6$ at $\bar{n} = 10^6$, a 40-fold increase. This large increase is due to the resonator's high surface dielectric participation. **b** Temperature dependence of an indium micromachined cavity[82]. Q_{int} rises by a factor of two as temperature is increased due to thermally saturated TLSs and maximizes at around 500mK. This particular device had orders of magnitude less participation than the device in **a**, which is why the increase in Q with temperature is much less. Additionally, at higher temperatures thermal quasiparticle loss begins to dominate, causing the Q to decrease.

couple from the mode, reducing the dielectric loss and increasing the quality factor (Fig. 3.7a). At fixed drive powers and changing temperatures, the same effect occurs; thermal noise saturates the TLSs, resulting in lower loss (Fig. 3.7b). As temperatures continue to increase, however, thermal quasiparticle loss begins to dominate, and conductor loss reduces the total internal Q of the mode.

Power dependence in superconducting microwave resonators is a hallmark of TLS loss. Generally, thin-film CPW resonators in planar architectures tend to be very sensitive to TLSs due to their high surface participation, with quality factors changing by over an order of magnitude from $\bar{n} = 1$ to $\bar{n} = 10^6$. 3D cavities, on the other hand, show at most a factor of two change in Q_{int} . A special case is the micromachined cavity[82], which is a 3D cavity made with thin-film superconductors which has as little power dependence as 3D cavities. This is because both of these types of cavity resonators have much lower surface participation.

There have tremendous efforts over the years to characterize, understand, and mitigate TLS loss[114, 115, 117, 119, 120]; however, the physical signature or source of TLSs have not been identified. Moreover, much of the focus has been on TLSs that couple to the electric fields of resonators and transmons. However, TLSs could potentially couple to magnetic fields as well; to my knowledge exploration into magnetically-coupled TLSs is lacking.

TLSs in transmons can have some unique effects. For one, the Josephson junction is a region of very high electric field density due to the small size and presence of large electric fields induced from the parasitic capacitance of the junction electrodes. As a result, transmons can strongly couple to discrete TLSs in the region of the junction. This results in transmons being particularly sensitive to relaxation via TLSs. However, there is strong evidence to show that TLSs can fluctuate in frequency over timescales of several hours[100, 121]; this results in temporal fluctuations in transmon coherence as the

coupling between the TLS and the transmon can change over long periods of time. This also implies that the loss tangent of the junction oxide can fluctuate due to the temporal fluctuations of TLSs. There is evidence to show that cosmic rays are responsible for this effect, known as “TLS scrambling”[100].

3.3.3 Contact (Seam) Loss

Contact loss, or seam loss, is resistive loss that occurs at the interface between two contacting metals. This can most commonly be seen at the joint of the enclosure that defines the package that separates the internal from the external environment. An enclosure must be made of two distinct pieces that are joined together. Contaminants, surface oxides, or surface roughness in these interfaces can contribute to imperfect metal-to-metal contact, resulting in some resistive loss when current flows through the seam[122–124]. Contact loss can also exist in planar circuits. In the case where two overlapping thin films deposited and patterned in two separate lithography steps, residues from the fabrication process or surface oxides can lead to a contact resistance[86], as in hybrid transmon qubits made with Nb or Ta capacitor pads and Al/AIO_x/Al junctions. In the case of indium bump-bonding, seam loss can exist when two indium bumps are compressed against each other[82].

The description of seam loss is a phenomenological one; we make no assumptions about the microscopic mechanism, only that there exists some resistance through which current flows and dissipates energy. The seam is modeled as a one-dimensional line around the perimeter of the joint. The fields of the resonance mode induce currents to flow along the superconductor and through the seam of resistance R_{seam} , resulting in power dissipa-

tion:

$$\frac{1}{Q_{\text{seam}}} = \frac{P_{\text{diss}}}{\omega_r U_{\text{tot}}} = \frac{R_{\text{seam}} \frac{1}{2} I_{\text{seam}}^2}{\omega_r \frac{1}{2} \int_{\text{all}} \mu_0 |\vec{H}|^2 dv} = \frac{1}{g_{\text{seam}}} \left[\frac{\int_{\text{seam}} |\vec{J}_s \times \hat{l}|^2 dl}{\omega_r \int_{\text{all}} \mu_0 |\vec{H}|^2 dv} \right] = \frac{y_{\text{seam}}}{g_{\text{seam}}} = p_{\text{seam}} \Gamma_{\text{seam}} \quad (3.25)$$

We have introduced a seam conductance per unit length $g_{\text{seam}} = (LR_{\text{seam}})^{-1}$, where L is the length of the seam. g_{seam} is the intrinsic loss that corresponds to seam loss, and the seam admittance per unit length y_{seam} is the geometric contribution that is determined purely by the current flow into the seam. We can once again invoke the boundary condition $\vec{J}_s = \hat{n} \times \vec{H}$ to write y_{seam} in another way[31]:

$$y_{\text{seam}} = \frac{\int_{\text{seam}} |\vec{H}_{\parallel}|^2 dl}{\omega_r \int_{\text{all}} \mu_0 |\vec{H}|^2 dv} \quad (3.26)$$

which is operationally more convenient to compute in Ansys HFSS. Importantly, g_{seam} and y_{seam} are not dimensionless; they both have units of $(\Omega\text{m})^{-1}$. However, we can make them dimensionless by multiplying both by $(\omega_r \epsilon_0)^{-1}$. Then we can define the seam participation $p_{\text{seam}} = y_{\text{seam}} (\omega_r \epsilon_0)^{-1}$ and seam loss factor $\Gamma_{\text{seam}} = \omega_r \epsilon_0 / g_{\text{seam}}$.

For 3D cavity architectures, seam loss can have a large device-to-device variation due to machining imprecision and lack of control over interface contamination. Additionally, the seam quality is influenced by the force used to clamp the two parts together. This is mostly accomplished by screwing the parts together. However, threaded taps in high-purity aluminum strip quite easily due to the softness of the metal; as a result, it is difficult to apply large amounts of clamping force. There are ways to circumvent this, but it highlights the need to carefully engineer the enclosure in order to reduce the unpredictability of seam loss. The 3D post cavity completely avoids this issue through a “seamless” design; seams are placed far away from the storage and transmon modes, and cutoff waveguide sections are placed between the seam and the important high-Q modes. This results in

sharp attenuation of the electromagnetic fields of the modes as they propagate towards the seams, rendering the modes insensitive to seam loss despite any potential variation in seam quality. Another approach to improve seam loss is to utilize indium gaskets. Since indium is a very soft metal, with enough clamping force the indium can compress within the seam, improving metal-to-metal contact. Indium is also used in flip-chip architectures, where indium bump-bonds are used to electrically connect the the two chips together. Here, the proper use of surface treatments and bonding force can result in very high seam conductances ($g_{\text{seam}} \sim 10^{10}$)[82], orders of magnitude higher than the mating of two bulk superconductors in an enclosure ($g_{\text{seam}} \sim 10^3$, but with a well-compressed indium gasket, seam conductance can be as high as $g_{\text{seam}} \sim 10^6$)[122].

3.4 Dephasing

In Sec. 3.3 I covered energy relaxation mechanisms; here, I will discuss dephasing mechanisms. While relaxation is caused by the qubit or resonator coupling to noise at the resonance frequency, dephasing is caused by off-resonant noise. Linear resonators tend to have no intrinsic dephasing; any coupling of a degree of freedom that could make it dephase turns the system into a nonlinear one. An example is ambient vibrational noise, most commonly associated with the pulse-tube cooler of the dilution refrigerator. Such a noise source causes the circuit to mechanically jitter in its package, potentially causing its capacitance to the walls of the package to fluctuate. Vibrational noise can also drive mechanical resonance modes on the chip itself; however, these are usually high in frequency and beyond the major low-frequency vibrations created by the pulse-tube cooler. Moreover, relaxation in a transmon can cause the frequency of a dispresively coupled resonator to shift stochastically, causing it to dephase. Transmons, on the other hand, couple nonlinearly to other degrees of freedom, even those that are uncontrolled. As a result, noise

sources and fluctuations within those degrees of freedom can cause the transition frequencies of the transmon to fluctuate, leading to dephasing. In this section, I will describe the effects of TLS noise, thermal photon shot noise, and flux noise on transmon qubits.

3.4.1 TLS Noise

As discussed in Sec. 3.3.2, TLSs can provide additional relaxation pathways for electromagnetic resonance modes. In transmons, the fluctuating electric dipoles of TLSs near the junction produces polarization noise that can couple to the transmon and induce charge fluctuations. The derivation of the resulting noise spectrum is derived in Constantin et al. [125], and it is shown that a single TLS can dephase the transmon by introducing a noise spectrum that is Lorentzian ($S_P(\omega) \propto \frac{1}{1+\omega^2}$) at low frequencies due to the intrinsic relaxation of the TLS and has a peak at the TLS resonant frequency. An ensemble of TLSs, however, produce a $1/f$ noise spectrum at low frequencies and a white noise spectrum at high frequencies. Low-frequency dephasing noise from TLSs can potentially be abated through dynamical decoupling, which will be discussed in Sec. 3.5.3.

While TLSs can cause dephasing in a transmons, their effect is substantially reduced in resonators. Since transmons are far more sensitive to TLSs interacting in the region of the junction, they are more sensitive to TLS-induced decoherence than resonators, which have a far larger area and more dilute electric field. The resonator therefore couples to far more TLSs than a transmon; however, it couples to each one of them very weakly. As a result, TLS fluctuations have a lesser impact on resonator coherence.

3.4.2 Thermal Photon Shot Noise

Thermal photon shot noise usually originates from the transmon's dispersive coupling with the readout resonator. Because the readout resonator is strongly coupled to the readout

control line, it will be well-thermalized to it; high thermal photon occupation in the readout line results in photon number fluctuations in the readout resonator, which induce an off-resonant AC stark shift. This results in a transmon dephasing rate of [126]

$$\Gamma_{\phi}^{th} = \frac{\bar{n}\kappa_r\chi_{qr}^2}{\kappa_r^2 + \chi_{qr}^2} \quad (3.27)$$

where \bar{n} is the equilibrium photon occupation of the readout resonator. For large enough \bar{n} and χ_{qr} , photon number-resolved spectroscopy can determine the average photon number of the resonator (see Ch. 2.3.3). However, even for low photon number occupations of $\bar{n} = 0.1$ and $\chi_{qr} = \kappa_r = 2\pi \times 1$ MHz, the dephasing time drops to $T_{\phi} \approx 3$ μ s. As a result, it is very important to ensure that the readout line is well-thermalized to reduce the equilibrium photon occupation under a noise temperature load that originates from the room-temperature control electronics.

Noise temperatures of signal generators and FPGAs tend to be very high, upwards of 10^5 K. This noise can be reduced by using attenuators that are themselves well-thermalized. However, the attenuators themselves add noise at their equilibrium temperature, which necessitates using a cascaded array of attenuators at each temperature stage of the dilution refrigerator until the photon noise in control line can no longer significantly dephase the qubit.

The thermal photon noise of a noise source a temperature T can be calculated using Bose-Einstein statistics:

$$\bar{n} = \frac{1}{e^{\hbar\omega/k_B T} - 1} \quad (3.28)$$

The photon noise is attenuated by an attenuator at a particular temperature stage; the at-

tenuator itself adds noise \bar{n}_a

$$\bar{n}_{\text{attenuated}} = \bar{n}_a(1 - 10^{A/10}) + \bar{n}_{\text{in}}10^{A/10} \quad (3.29)$$

where A is the attenuator's attenuation in dB (attenuation is negative). Eq. (3.29) can be cascaded for additional attenuators and successive temperature stages. It is important to note that the attenuators themselves are playing the role of thermalizers of the signal in the control line. As a result, the attenuators themselves must be well-thermalized, as they are subjected to a constant heat load by the drive.

3.4.3 Flux Noise

Thus far I have only been discussing fixed-frequency transmons. A transmon with a SQUID loop, or two identical junctions in parallel in a closed loop, is frequency-tunable with magnetic flux. This provides another way to control the qubit, as transmons can be quickly tuned with fast DC flux pulses, and AC modulation of flux introduces other new dynamics[25]. However, such a device is now sensitive to flux noise, which cause dephasing. Sensitivity to flux noise is dramatically suppressed around flux sweet spots where the frequency change is to first order insensitive to changes in flux. When the transmon is operated outside the sweet spot, it will dephase due to $1/f$ flux noise, whose origin is hypothesized to be magnetic two-level system defects residing on surface oxides on the SQUID loops[127]. Such a noise model implies that the flux noise may be dependent on the geometry of the loop; $S(\omega) \propto R/W$, where R is the radius of the SQUID loop and W is the width of the superconducting trace that forms the loop. Experimental evidence consistent with this model was shown[128], lending credence to the theory that random flips of electronic spins are the source of the flux noise.

3.5 Coherence Measurement Techniques

In this section, I will cover the experimental techniques to measure coherence in transmons and resonators. Resonators can be measured in frequency domain or time domain to extract Q_c , Q_{int} , and ω_r . Transmons are measured in time domain to extract T_1 and T_2 . Quantum memory T_1 and T_2 is also measured in time domain by inferring the memory state using the transmon's dispersive coupling to it. I will first describe how devices are packaged in the coaxial architecture in order to highlight how the control line drives transmons and resonators. I will then describe the measurement setup within the fridge and detail the cryogenic and room-temperature electronics that make up the measurement chain. Finally, I will describe how resonators, transmons, and quantum memories are measured.

3.5.1 Device Packaging

All on-chip striplines, transmons, and quantum memories for this thesis were designed in the coaxial architecture. Cylindrical tunnel packages that contained four independent tunnels were made out of conventionally-machined high-purity (5N5) aluminum (See Fig. 3.8). To reduce surface losses associated with the packages, they were etched in a mixture of phosphoric and nitric acid (Transene Aluminum Etchant Type A) heated to 50 °C for two hours[90, 112]. The tunnels were approximately 34 mm long and 5 mm in diameter. Coupling was accomplished by a transverse feedline, allowing for devices in the four tunnels to be measured in a multiplexed hanger configuration[84]. For the measurements that will be detailed in Ch. 5, the same feedline is used for qubit, storage mode, and readout drives. Devices are fabricated on 40 mm \times 4 mm chips and are inserted into the cylindrical waveguide tunnels. The chips are clamped within the tunnels on either end by beryllium-copper leaf-springs. The clamps on either end of the tunnel also serve

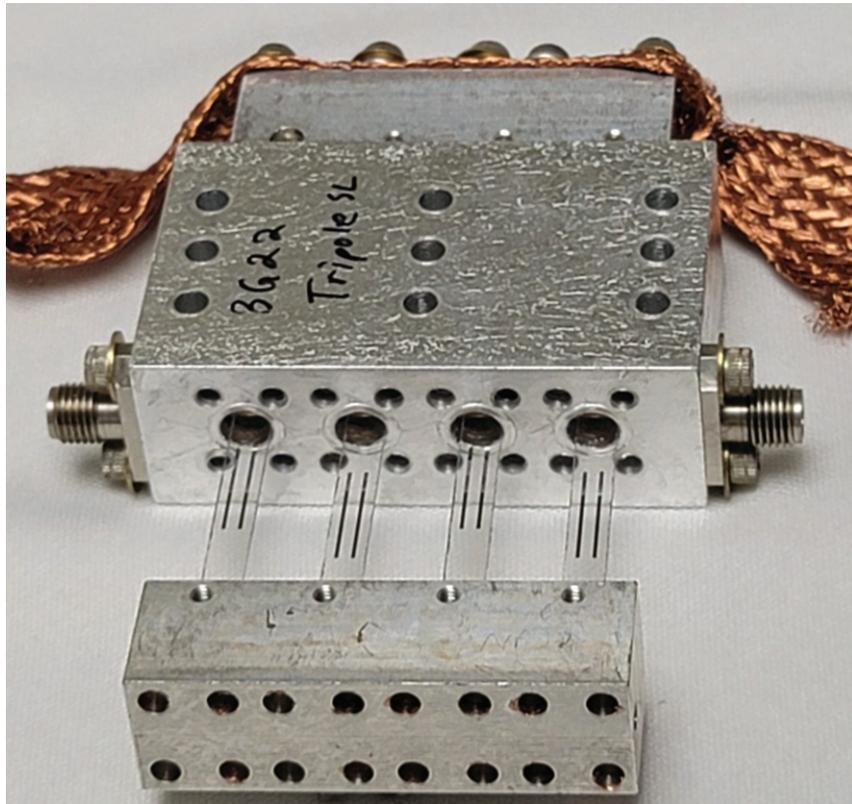


Figure 3.8: **Device Packaging.** Multiplexed coaxial tunnel package. Stripline resonators or transmon qubits are patterned on dielectric substrates that are then inserted into coaxial waveguide tunnels made of high-purity aluminum. Coupling is done using a transverse centerpin connected on either end by coaxial SMA flanges. Figure from Ganjam et al. [86].

as end-caps for the tunnels themselves, thereby defining the locations of the seams and completing the enclosure with well-defined package modes that are high (>18 GHz) in frequency.

3.5.2 Measurement Setup

A fridge wiring diagram can be found in Fig. 3.9. Device packages are mounted to the mixing chamber stage of a dilution refrigerator operating at 20 mK. The packages are enclosed in multiple layers of shielding. First, a light-tight gold-plated copper shield internally coated with Berkeley black acts as an IR photon absorber[98]. A superconducting shield made of 1/64" thick lead foil is wrapped around the copper shield. Finally, a Cryoperm can serves as the outermost shield to attenuate the ambient magnetic fields at the package. Input lines are attenuated at both the 4 K stage (20 dB) and mixing chamber stage (50-60 dB depending on the line; 20 dB of reflective attenuation is achieved through the use of a directional coupler) and are filtered at multiple locations using 12 GHz K&L low-pass filters and custom-made eccosorb CR-110 IR filters. Output lines are also low-pass filtered and isolated from the devices using circulators and isolators. A SNAIL parametric amplifier (SPA) is used on the qubit output line to provide quantum-limited amplification for qubit readout. HEMT amplifiers at the 4 K stage provide additional low-noise amplification for the output signals.

Resonators are measured in frequency domain using a vector network analyzer (Agilent E5071C). Qubits and quantum memories are measured in time domain using an FPGA-based quantum controller (Innovative Integration X6-1000M) which can output arbitrary waveforms in pairs of I and Q quadratures at ≈ 50 MHz that are then up-converted to GHz frequencies using an LO tone generated by an Agilent N5183A (Readout drive uses a Vaunix LMS-103 for the LO) and a Marki IQ-0307-LXP mixer. Qubit, readout, and storage mode drives are all generated the same way and are combined and amplified using a

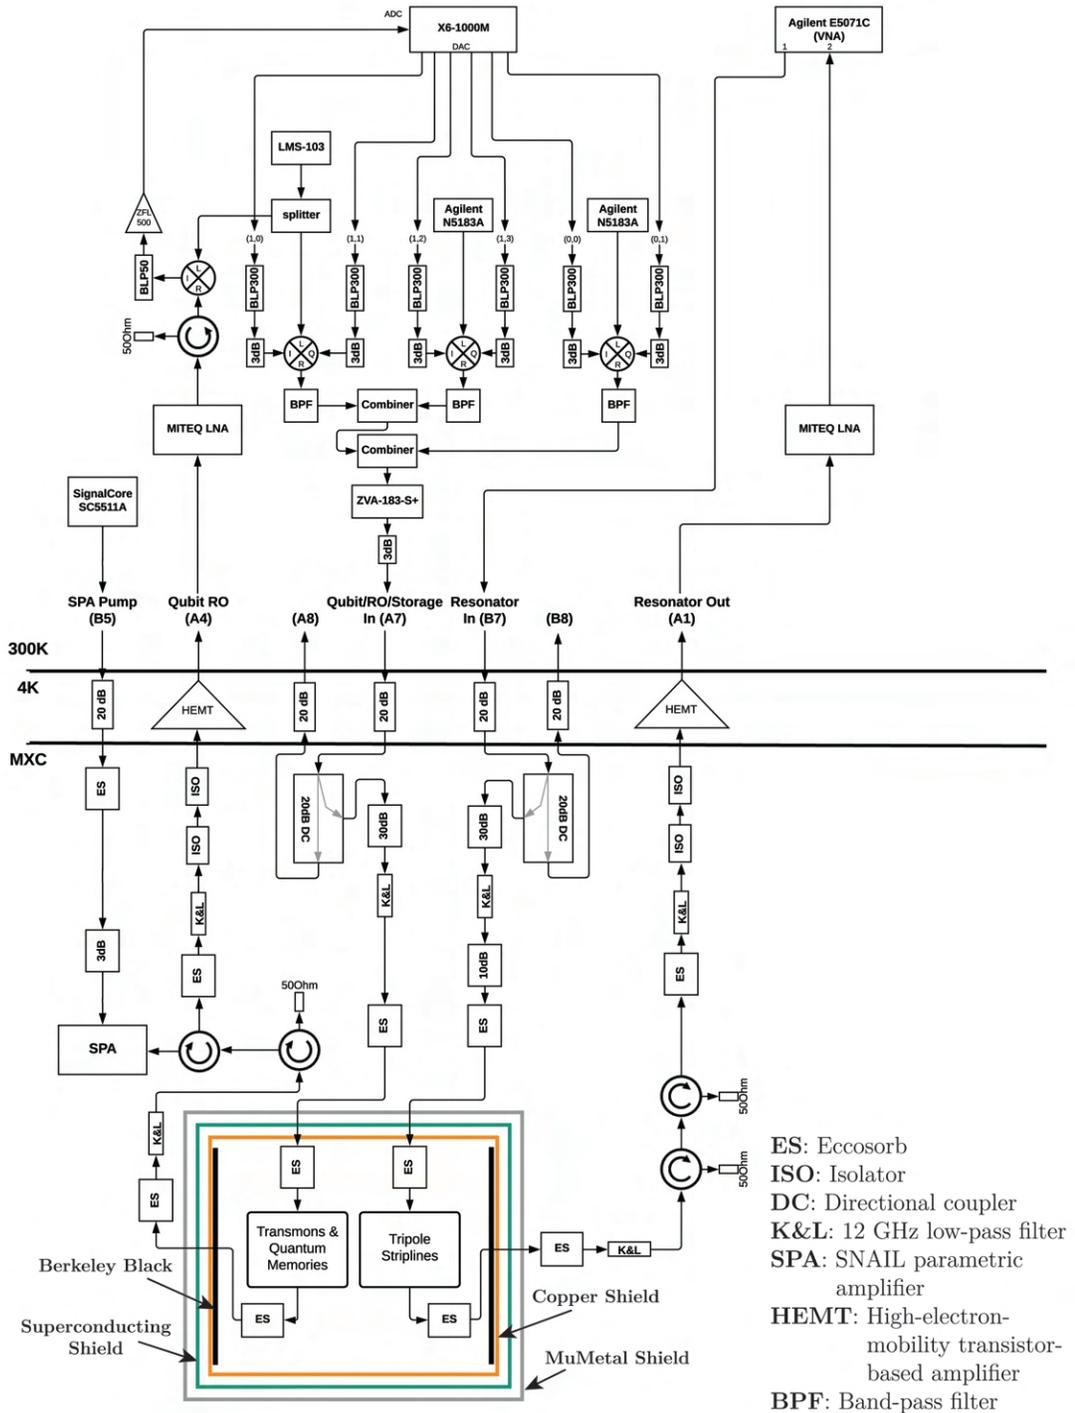


Figure 3.9: Fridge Wiring Diagram. Figure from Ganjam et al. [86].

Mini-Circuits ZVA-183-S+. The signals are finally attenuated by a room-temperature 3 dB attenuator to reduce the thermal noise temperature before being fed into the fridge. Readout responses from the fridge are amplified with a room-temperature amplifier (MITEQ LNA-40-04001200-15-10P) and isolated before being down-converted using a Marki IR-0618-LXP mixer (the same LO is used for both the up-conversion and down-conversion of the readout signals). Down-converted signals are then amplified using a Mini-Circuits ZFL 500 before being fed into the ADC of the FPGA. All signal generator sources and VNA are clocked to a 10MHz Rb frequency standard (SRS FS725).

We can use this wiring diagram to estimate the thermal photon occupation of the control line. Using the expressions from Sec. 3.4.2 and accounting for the attenuation of RF lines going from the 300 K stage to the mixing chamber stage, we arrive at a photon occupation of 3×10^{-5} which, from Eq. (3.27) gives a dephasing rate of $\Gamma_{\phi}^{th} = 1/537 (\mu\text{s})^{-1}$. However, this assumes that the attenuators are well thermalized to their respective stages; if, for example, the final-stage attenuator is not well-thermalized, the equilibrium temperature due to the noise of the room-temperature electronics could be far higher; this would change the photon occupation of the readout resonator and therefore increase the dephasing rate.

3.5.3 Resonator Quality Factor Measurements

Here, I will describe how resonator quality factors are measured. For this thesis, resonators were measured in frequency domain using a vector network analyzer (VNA). The VNA performs a frequency sweep around the resonance frequency and measures the scattering parameter. However, due to the long integration time at each frequency point, VNA measurements tend to be phase-sensitive. As a result, frequency fluctuations that occur on time scales faster than the sweep time manifest as a broadening of the Lorentzian linewidth. This is effectively a T_2 measurement; $Q = \omega_r T_2$. In most cases, resonators have no in-

trinsic dephasing. However, in the case of significant dephasing due to vibrational noise, the VNA cannot capture the energy relaxation rate. In such cases, it is preferred to do a time-domain, or ring-down measurement. A microwave pulse at the resonator frequency (the pulse width determines the bandwidth and height determines the input energy, which will determine the circulating photon number) is used to drive the resonator, populating it with photons. The time-dependent behavior is then observed; the ring-up dynamics are determined by Q_c whereas the ring-down dynamics are related to the coherence time. By independently measuring the I and Q quadratures of the signal separately, T_1 and T_2 can be distinguished based on how the quadratures are averaged; i.e. T_1 can be extracted from $\langle I^2 + Q^2 \rangle$ and T_2 can be extracted from $\langle I \rangle^2 + \langle Q \rangle^2$. A comprehensive description of the ring-down measurement can be found in Read et al. [106].

Vibrational noise is relatively easy to identify with a VNA. By observing the Lorentzian under high IF bandwidth settings, frequency jitter can be observed. Devices used in this thesis were not observed to dephase. This is likely due to the chip clamping approach detailed in Sec. 3.5.1; by clamping the chip at both ends, the “cantilever”-like vibrational mode can be suppressed. As a result, the VNA is sufficient to measure the energy relaxation of the resonators.

Measurement Configurations

Measuring a resonator depends on the measurement configuration. If we consider the resonator to be an arbitrary microwave network, there are various locations we can place the microwave port or ports in order to measure the resonator quality factor. There are three such configurations: transmission, reflection, and hanger.

The transmission configuration is a two-port measurement where the resonator is sandwiched between two transmission lines (Fig. 3.10a) and the transmission of the system is

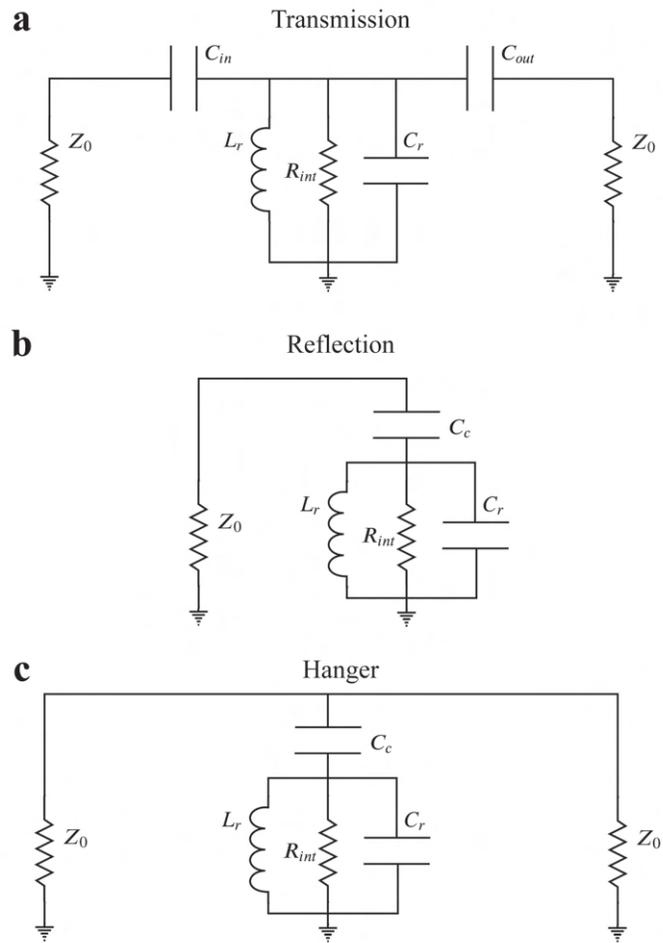


Figure 3.10: Resonator Measurement Configurations.

measured with the VNA. The scattering parameter $S_{21}(f)$ is given by[129]

$$S_{21}(f) = ae^{i\alpha}e^{-2\pi if\tau} \left[\frac{2(Q_l/|Q_c|)e^{i\phi}}{1 + i2Q_l(f/f_r - 1)} \right] \quad (3.30)$$

The various prefactors in from the bracketed expression describe the behavior of the input and output lines. Attenuation in the input line and amplification of the output line combine to produce the net attenuation a . The global phase shift α occurs due to the presence of various filters, directional couplers, and even attenuators that may produce slight phase shifts. The finite speed of light results in an electrical delay τ , which results in a frequency-dependent phase shift of the signal. f_r is the resonance frequency, and $Q_c = |Q_c|e^{-i\phi}$ is a complex coupling quality factor where ϕ describes an asymmetry in the response due to slight impedance mismatches between the ports[130]. The real-valued loaded quality factor Q_l is the total quality factor due to both internal and external (coupling) loss, $1/Q_l = 1/Q_{\text{int}} + 1/Q'_c$, where $1/Q'_c = \cos \phi/|Q_c|$. Additionally, in the case of a transmission configuration, there are two ports to couple to; as a result, there is an input and output coupling quality factor. The extracted Q_c is the geometric mean of the two coupling quality factors, $Q_c = \sqrt{Q_{\text{in}}Q_{\text{out}}}$.

Off-resonance, no signal is transmitted through the resonator, so $S_{21} = 0$. On resonance, in the case where $\alpha = \tau = \phi = 0$ and $a = 1$, transmission is maximized, $S_{21} = 2Q_l/Q_c$. Of interest is the critically-coupled regime, where $Q_{\text{int}} = Q_c$, and $S_{21} = 1$; the maximum available power is delivered to the resonator. The critically-coupled regime provides the highest signal-to-noise ratio for resonator measurement; it is therefore desired to critically couple the resonator for purposes of quality factor measurement.

The drawback of a transmission configuration is that it is uncalibrated. Nominally, the VNA's reference is set to its ports, and it measures the transmission of the device in addition to the several meters of coaxial cabling, attenuators, isolators, and amplifiers between

the ports. These lines and components in principle need to be calibrated away such that the new reference would be at the input and output of the device under test. This can be done at room temperature where the lines inside the fridge can be accessed; however, because the line attenuation changes when the fridge is cooled down, this calibration is not useful. A proper calibration requires some form of *in situ* cryogenic switch that can switch between measuring the device and a through-transmission line, which would define the unity transmission that the VNA can use to calibrate. Additionally, since there are separate coupling capacitances to each port, full characterization of both Q_{in} and Q_{out} requires measurement of S_{21} and one of S_{11} and S_{22} . In practice, this requires a specialized setup; it is often times far more convenient to measure the resonator in reflection or hanger configurations.

The reflection configuration is somewhat simpler (Fig. 3.10b), it is a single-port configuration where the signal-carrying transmission line is terminated with a load impedance that is the resonator. The measured reflected signal $S_{11}(f)$ is given by [129, 130]

$$S_{11}(f) = ae^{i\alpha}e^{-2\pi i f \tau} \left[1 - \frac{2(Q_l/|Q_c|)e^{i\phi}}{1 + i2Q_l(f/f_r - 1)} \right] \quad (3.31)$$

This expression looks remarkably similar to Eq. (3.30) with one major difference: the additional term of 1 in the bracketed expression. This effectively makes the measurement self-referencing; as a result, measurement of S_{11} fully characterizes the resonator. In the case where $\alpha = \tau = \phi = 0$ and $a = 1$, $S_{11} = 1$ off-resonance and $S_{11} = 1 - 2Q_l/Q_c$ on-resonance. The Lorentzian line-shape is inverted relative to the transmission configuration. In the critically-coupled regime, the depth of the Lorentzian peak is maximized; $S_{11} = 0$, an impedance-matching condition is satisfied where the maximum available power is delivered to the load. However, in the cases where $Q_c \gg Q_{\text{int}}$ (undercoupled regime) or $Q_c \ll Q_{\text{int}}$ (overcoupled regime), $S_{11} \approx 1$; the signal becomes harder to distinguish from the background as the SNR falls.

In practice, because the input line has significant attenuation, measuring the reflected signal would not be possible due to the significantly low SNR. Instead, a circulator or directional coupler is used to reroute the reflected signal towards the amplifier chain. As a result, the VNA technically measures S_{21} despite the fact that the reflected signal is what is measured. However, the difference in path traversed between the incoming and outgoing signals creates a Fano interference effect, resulting in asymmetry; hence, the presence of ϕ .

One drawback of reflection and transmission configurations is that only one resonator can be measured per measurement line. Devices in the coaxial tunnel package were therefore measured in hanger configuration (Fig. 3.10c). In this configuration, a transmission line is shunted midway by a resonator and a coupling capacitance. The input signal sees an effective load impedance that is the equivalent of the resonator in parallel with the terminating impedance Z_0 of the output line. As a result, some of the power will be shunted through the resonator, resulting in a reduction in transmission. S_{21} is given by [130, 131]

$$S_{21}(f) = ae^{i\alpha}e^{-2\pi if\tau} \left[1 - \frac{(Q_l/|Q_c|)e^{i\phi}}{1 + i2Q_l(f/f_r - 1)} \right] \quad (3.32)$$

Once again, the expression looks very similar to those of the transmission and reflection configuration; the only difference being the absence of a factor of 2 in the fraction in the brackets. In hanger configuration, the difference in incoming and outgoing paths for the signal coupled with slight impedance mismatches between the input and output ports results in Fano interference, necessitating the use of the asymmetry parameter ϕ . In the case where $\alpha = \tau = \phi = 0$ and $a = 1$, $S_{21} = 1$ off-resonance and $S_{21} = 1 - Q_l/Q_c$ on-resonance; hanger configuration is also self-referencing. In the critically-coupled regime, $S_{21} = 1/2$; power is split between the terminating impedance of the output line and the resonator. Similar to reflection, in the undercoupled regime $S_{21} \approx 1$. Unusually, $S_{21} =$

0 in the overcoupled regime. This implies that the highest SNR is achieved when the resonator is overcoupled; however, in this regime, the Q_l is entirely determined by Q_c ; extraction of Q_{int} would therefore propagate a tremendous amount of uncertainty. While the fitting algorithms are otherwise quite robust, it is therefore recommended that for the purposes of extracting Q_{int} , the resonator should be over- or undercoupled by more than a factor of 10 or 20.

Fitting

Fitting the expressions in Eqs. (3.30), (3.31), and (3.32) to extract f_r , Q_c , and Q_{int} requires two separate fits. Here, I will describe the fitting process for the hanger configuration; although the process for the other configurations are very similar. The first fit is called the circle fit, where S_{21} is plotted on the complex plane and fit to a circle. The electrical delay distorts the circle and must first be measured and calibrated out. To do this, the phase versus frequency response is fitted over a broad frequency span of ≈ 1 GHz. Once τ has been measured, it is used to subtract out the distortion in the complex plane. The circle is the result of a particular parameterization:

$$\theta(f) = \tan^{-1}[2Q_l(f/f_r - 1)] \implies \tan \theta = 2Q_l(f/f_r - 1) \quad (3.33)$$

This modifies Eq. (3.32) to (after cancelling the electrical delay by multiplying S_{21} by $e^{2\pi i f \tau}$):

$$S_{21}(f) = ae^{i\alpha} \left[1 + \frac{1}{2} \frac{Q_l}{|Q_c|} (e^{-i2\theta} - 1) e^{i\phi} \right] \quad (3.34)$$

This produces a transformed circle in the complex plane. The circle is parameterized by θ with radius $Q_l/2|Q_c|$ and center $(-\frac{Q_l}{2|Q_c|}, 0)$. The effect of ϕ is to rotate the circle about the origin, and the the 1 term translates the circle by $(1, 0)$. Finally, α rotates the translated

circle by an angle α , and a attenuates the circle's radius and distance from the origin. The circle fit yields the circle's radius r_c and center (x_c, y_c) :

$$\begin{aligned}
 r_c &= \frac{Q_l}{2|Q_c|}a \\
 x_c &= a \cos \alpha - r_c \cos(\phi + \alpha) \\
 y_c &= a \sin \alpha - r_c \sin(\phi + \alpha)
 \end{aligned} \tag{3.35}$$

This allows us to translate the circle to the origin and perform the second fit, called the phase fit:

$$\theta'(f) = \theta_0 - 2 \tan^{-1}[2Q_l(f/f_r - 1)] \tag{3.36}$$

where $\theta_0 = \phi + \alpha$. From the phase fit, we extract θ_0 , Q_l , and f_r , which allows us to solve for the rest of the parameters:

$$\begin{aligned}
 \alpha &= \tan^{-1} \left(\frac{y_c + r_c \sin \theta_0}{x_c + r_c \cos \theta_0} \right) \\
 a &= \sqrt{(x_c + r_c \cos \theta_0)^2 + (y_c + r_c \sin \theta_0)^2} \\
 \phi &= \theta_0 - \alpha \\
 Q_c &= \frac{Q_l}{2r_c} a e^{-i\phi} \\
 \frac{1}{Q_{\text{int}}} &= \frac{1}{Q_l} - \text{Re} \left[\frac{1}{Q_c} \right]
 \end{aligned} \tag{3.37}$$

Calculating \bar{n}

To characterize the power dependence of Q_{int} , it is necessary to determine the circulating photon number in the resonator as function of input power. This is done by expressing the

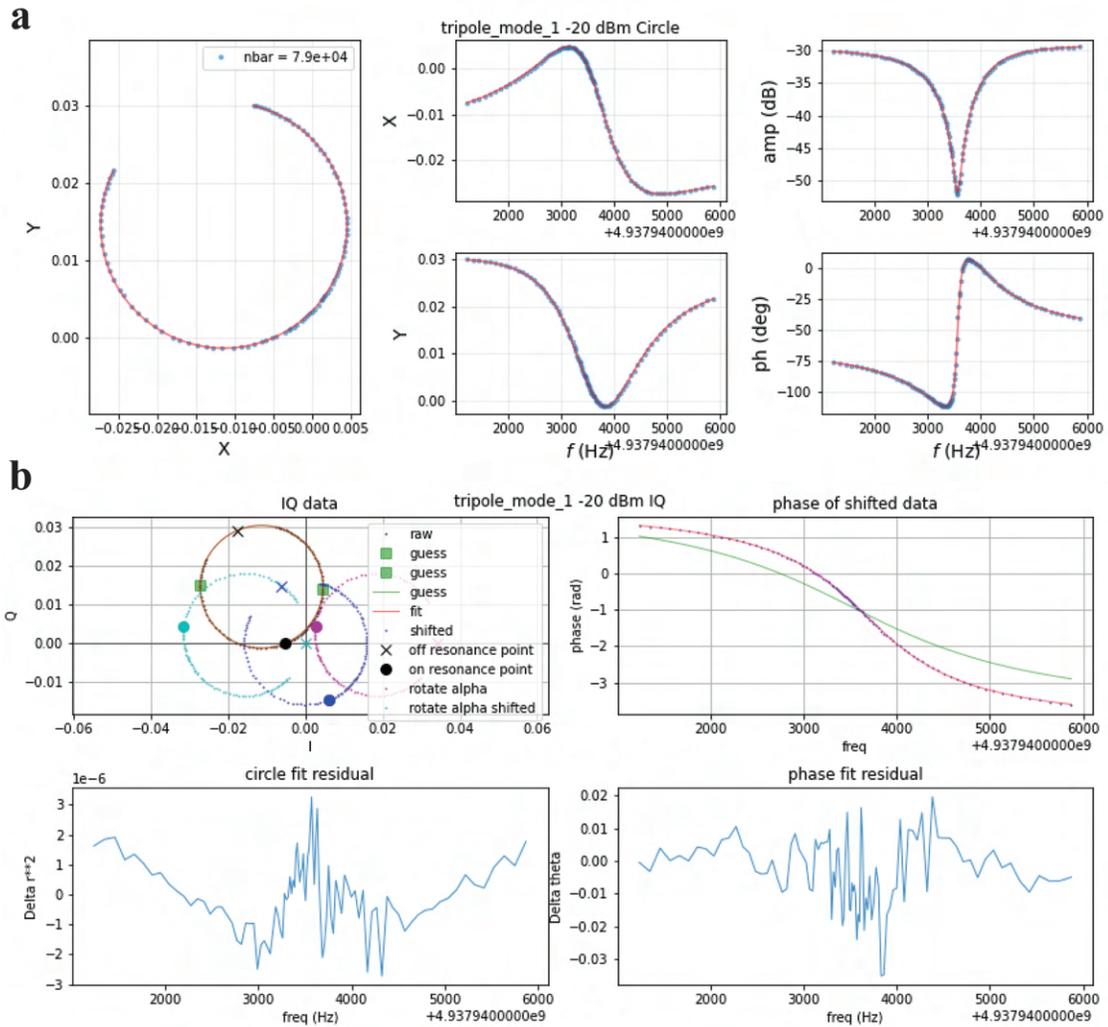


Figure 3.11: **Resonator Fitting.** **a** Circle fit in complex plane (left). The raw data (top right) showing the Lorentzian dip has a slight asymmetry to it. **b** Phase fit. The phase is fit versus frequency (top right) to extract f_r , θ_0 , and Q_l . The circle is transformed by each term in S_{21} (top left).

average photon number \bar{n} in terms of the dissipated power P_{diss} :

$$\bar{n} = \frac{P_{\text{diss}} Q_{\text{int}}}{\hbar \omega_r^2} \quad (3.38)$$

Calculating P_{diss} requires the analysis of the circuits for the various geometries. Another method to calculate \bar{n} is using input-output theory. We define the mode operator \hat{a} and input signal mode operator \hat{a}_{in} , where $P_{\text{in}} = \hbar \omega \langle \hat{a}_{\text{in}}^\dagger \hat{a}_{\text{in}} \rangle$. The time evolution of \hat{a} is then given by the quantum Langevin equation[86]:

$$\frac{d}{dt} \hat{a}(t) = \frac{i}{\hbar} [\hat{H}, \hat{a}(t)] - \frac{\kappa_l}{2} \hat{a}(t) + \sqrt{\kappa_c/2} \hat{a}_{\text{in}}(t) \quad (3.39)$$

where \hat{H} is the harmonic oscillator Hamiltonian, and $\kappa_{l,c} = \omega_r / Q_{l,c}$. Solving this equation and assuming resonant driving we arrive at for the hanger configuration:

$$\bar{n} = \frac{2\kappa_c}{\kappa_l^2} \frac{P_{\text{in}}}{\hbar \omega_r} = \frac{2}{\hbar \omega_r^2} \frac{Q_l^2}{Q_c} P_{\text{in}} \quad (3.40)$$

The average photon number can similarly be calculated for the other configurations. For transmission,

$$\bar{n} = \frac{2}{\hbar \omega_r^2} \frac{Q_l^2}{Q_{\text{int}}} P_{\text{in}} \quad (3.41)$$

and for reflection,

$$\bar{n} = \frac{4}{\hbar \omega_r^2} \frac{Q_l^2}{Q_c} P_{\text{in}}. \quad (3.42)$$

3.5.4 Transmon Coherence Measurements

Transmon coherence is measured in time domain using an FPGA-based controller that outputs I-Q pairs that are up-converted and sent to the device (see Sec. 3.5.2. Transmons

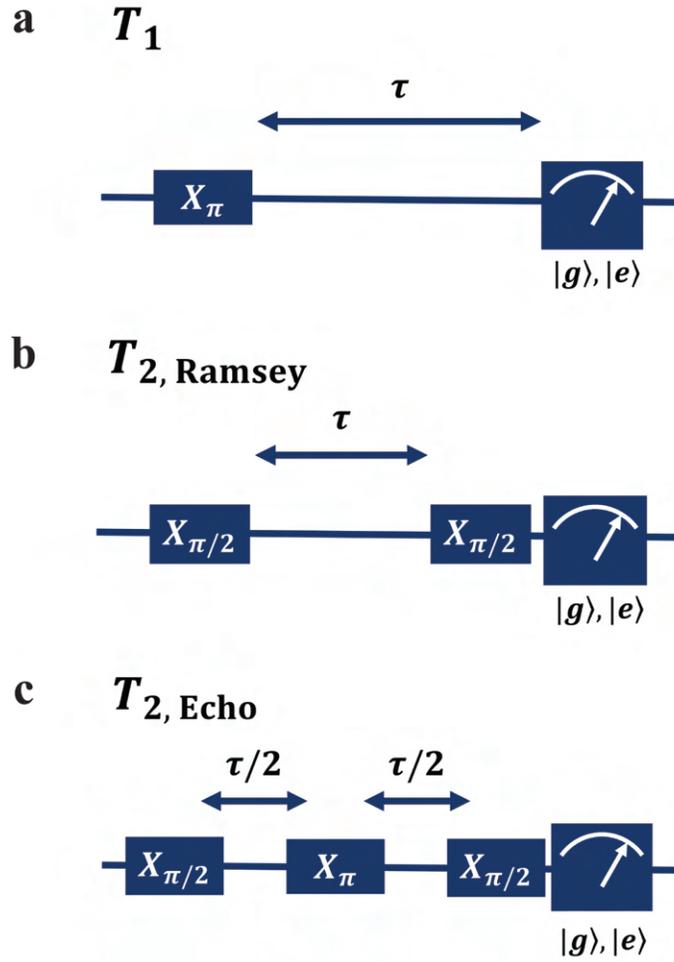


Figure 3.12: **Transmon coherence measurement pulse sequences.**

in this thesis were measured in multiplexed coaxial tunnel packages (see Sec. 3.5.1) in hanger configuration with the same line used for readout and transmon control. Here, I will describe the pulse sequences to measure transmon coherence.

Transmon Relaxation

Relaxation is measured by observing the exponential decay in the probability of finding the qubit in $|e\rangle$. An initial X_π pulse is used to prepare the transmon in $|e\rangle$. after a variable delay τ during which the qubit can relax, the qubit's state is measured by displacing the readout resonator (Fig. 3.12a). The experiment is repeated and the results are averaged to

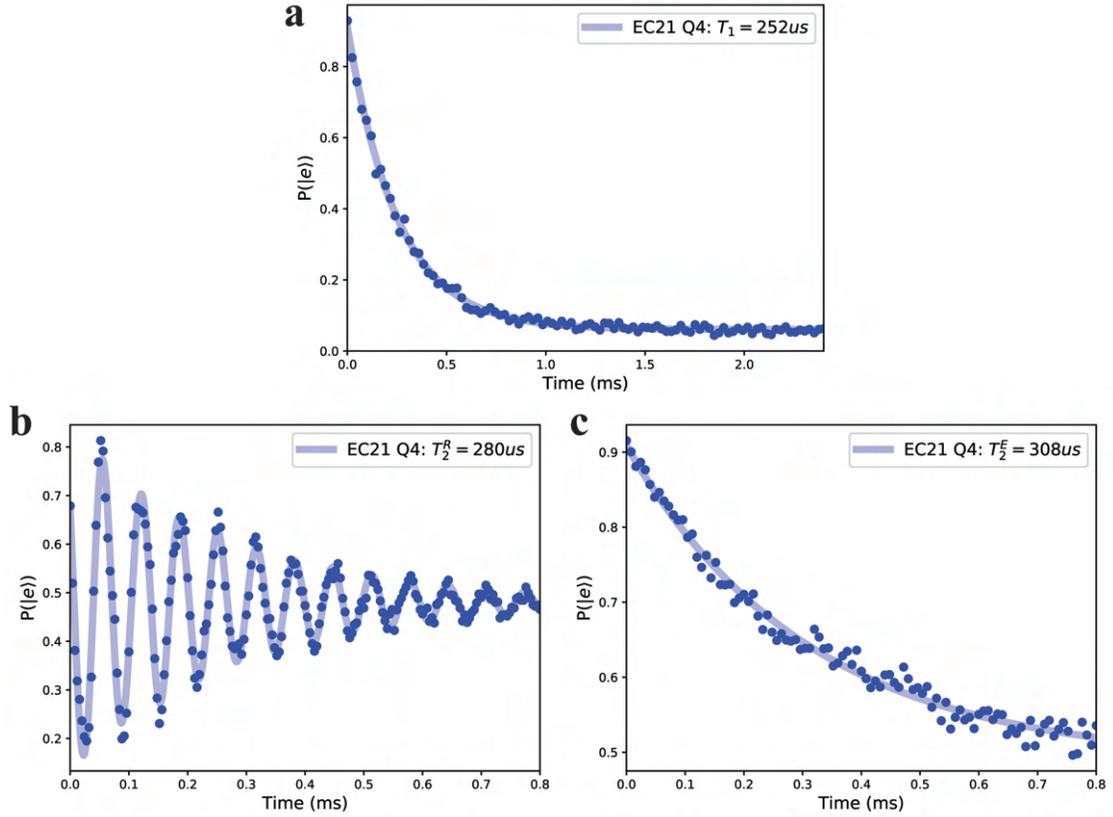


Figure 3.13: **Transmon coherence measurements.** **a** T_1 measurement of a tantalum-based transmon. **b** Transmon Ramsey T_2 coherence. **c** Transmon Hahn echo coherence. Since $T_2^E \approx T_2^R$, high-frequency noise must be dominating. Assuming T_2^E is dominated by thermal photon shot noise, this corresponds to $\bar{n} \approx 10^{-3}$ in the readout resonator, or $T \approx 62$ mK. this indicates that the control line is not well thermalized to the mixing chamber stage (20 mK) of the dilution refrigerator.

determine the probability of finding the qubit in $|e\rangle$ after a length of time. The data is then fit to an exponential; the characteristic time is T_1 (Fig. 3.13a).

Transmon Coherence

Transmon coherence T_2 is measured using Ramsey interferometry[42]. The qubit is initialized to the equator of the Bloch sphere using an $X_{\pi/2}$ pulse. After a variable delay during which decoherence processes scramble the phase relationship of the superposition, another $X_{\pi/2}$ is applied to project the qubit back to the z -axis, followed by readout (Fig.

3.12b). In practice, the qubit pulses are slightly detuned from ω_q , which causes the qubit to precess about the z -axis at a rate equal to the detuning of the drive from the qubit. This results in an approximately exponentially decaying sinusoidal behavior of $P(|e\rangle)$; fitting the oscillation frequency allows precise determination of the qubit frequency, and fitting the decay time yields the Ramsey decoherence time T_2^R (Fig. 3.13b).

The decoherence time is sensitive to how it is measured[42]. Low frequency noise, for example, can cause dephasing over time scales longer than the Ramsey experiment. As a result, fluctuations can vary from sequence to sequence. Ramsey interferometry is therefore highly sensitive to low-frequency noise. A Hahn echo experiment is much less sensitive to low-frequency noise, and can coarsely reveal aspects of the dephasing noise spectrum to which the qubit is subjected. In a Hahn echo experiment, the same Ramsey sequence is performed, but an additional X_π at $\tau/2$ after applying the initial $X_{\pi/2}$ pulse. This refocuses the noise-frequency fluctuations, resulting in a measurement that is insensitive to them (Fig. 3.12c). The echo sequence does not need to be driven off-resonantly; as a result, the behavior is a decaying exponential that decays to a mixed state; fitting the exponential gives the echo time T_2^E (Fig. 3.13c).

If $T_2^E > T_2^R$, then it is likely that low-frequency noise dominates (possibly due to Lorentzian noise from a single TLS, $1/f$ noise from an ensemble of TLSs, or $1/f$ flux noise, in the case of a flux-tunable transmon); whereas if $T_2^E \approx T_2^R$ high-frequency noise dominates (the most common culprit of which is thermal photon shot noise). Ramsey interferometry and Hahn echo are essentially noise spectroscopy measurements and are a subset of CPMG pulses which apply N X_π pulses between the two $X_{\pi/2}$ pulses, where $N \geq 0$. As N becomes larger, the measurement probes the noise spectrum at higher frequencies. CPMG pulses are described in Krantz et al. [42]. Another commonly used form of noise spectroscopy is spin-locking, where the qubit is locked into a rotating frame through a continuous Rabi drive, and the qubit's relaxation is measured in the rotating

frame of the spin-locking drive. The measurement then probes the noise at the Rabi frequency, which can then be swept. Spin-locking is described in Yan et al. [126].

Transmon Q_c

Measured transmon T_1 is proportional to the loaded quality factor of the mode, $(\omega T_1)^{-1} = Q_l^{-1} = Q_{\text{int}}^{-1} + Q_c^{-1}$. The coupling quality factor must be measured for transmons in order to determine Q_{int} which is limited by intrinsic loss mechanisms in the internal environment. While this can be done in finite element simulation, the true Q_c is dependent on the transmon chip's placement within the tunnel package and can vary by as much as 50% if the chip's position varies by as little as 0.5 mm from the nominal. We can instead determine the Q_c in situ by calibrating the qubit Rabi rate as a function of drive power; strongly coupled qubits need lower drive powers to achieve fast Rabi rates, and vice versa. We can recall Eq. (2.31) which relates the Rabi rate $\Omega_{\text{Rabi}} = \Omega_d$ to the displacement $\xi(t)$. Since the circulating photon number in the mode is just $\bar{n} = |\xi|^2$, We can write down the photon number in terms of the Rabi rate:

$$\bar{n} = \frac{\Omega_{\text{Rabi}}^2 Q_l^2}{\omega_q^2} \quad (3.43)$$

but from Eq. (3.40), we can relate the photon number to the drive power P_{in} , which gives us the relationship between Rabi rate and Q_c :

$$Q_c = \frac{2P_{\text{in}}}{\hbar\Omega_{\text{Rabi}}^2}. \quad (3.44)$$

3.5.5 Quantum Memory Coherence Measurements

Quantum memories are storage resonators that are dispersively coupled to the transmon, allowing for the generation and control of nonclassical states. Measurement of quantum

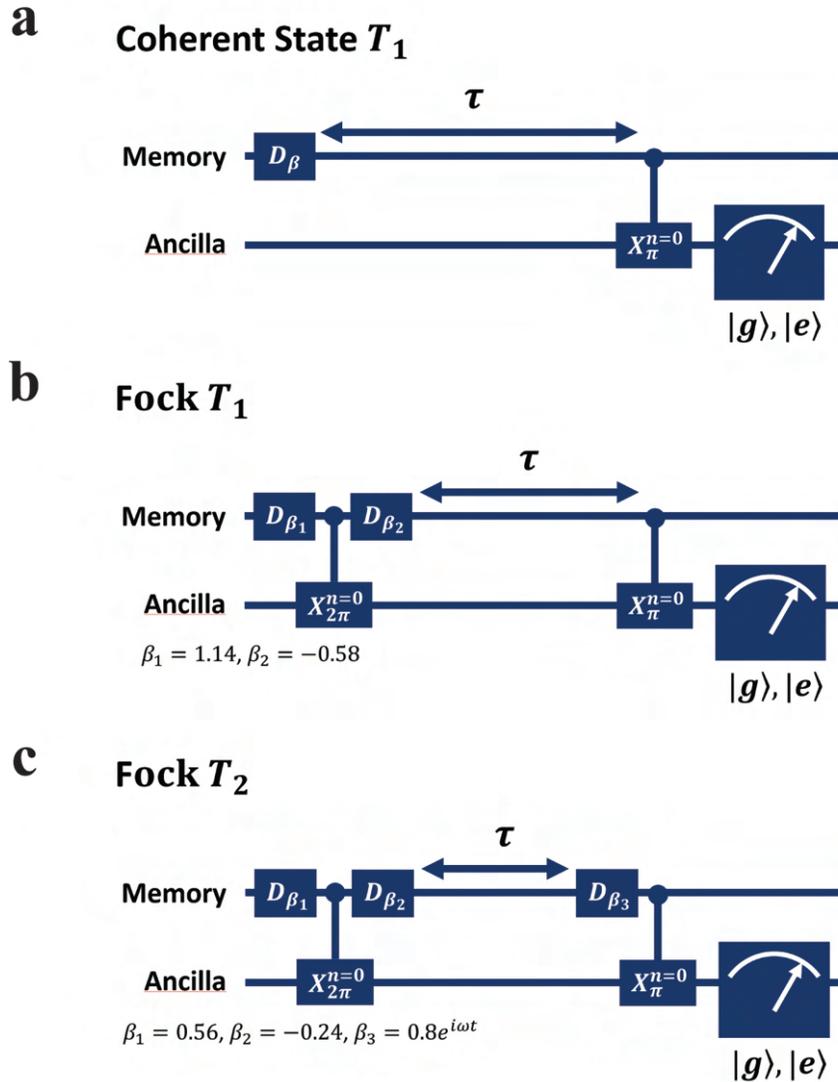


Figure 3.14: **Quantum memory coherence measurement pulse sequences.**

memory coherence requires the preparation of a state in the storage mode (see Ch. 2.2.1 and Ch. 2.3.3) and subsequent measurement by probing the state of the storage mode with the qubit. This measurement is done by checking if the storage mode is in *vacuum* or Fock $|0\rangle$ by applying a qubit X_π pulse conditioned on the storage mode being in Fock $|0\rangle$. Readout of the qubit is then performed; if the qubit is flipped successfully, then the storage mode must be in Fock $|0\rangle$. As a result, the probability of measuring the qubit in $|e\rangle$ as a function of time can be fit to determine the coherence.

Coherent State T_1

Coherent state decay is one of the basic coherence measurements one can do for a storage resonator. The resonator is first displaced by a small amount $\beta \approx 2$, followed by a measurement of the resonator state with the ancilla qubit after a variable delay τ (Fig. 3.14a). The probability of the qubit being in $|e\rangle$ is then measured as a function of delay time. The average photon number in the resonator decays as $\bar{n}(t) = |\beta|^2 e^{-t/T_1}$, but because we are measuring the population of the $n = 0$ peak, we must consider the probability of measuring 0 photons in the resonator for a coherent state of size $\beta = \sqrt{\bar{n}}$, which is $P_{n=0} = e^{-\bar{n}}$; therefore, the probability of measuring zero photons is given by

$$P_{n=0} = e^{-|\beta|^2 e^{-t/T_1}} \quad (3.45)$$

which is an exponential of an exponential. Fitting the probability of measuring the qubit in $|e\rangle$ gives the T_1 of a coherent state (Fig. 3.15a).

Fock T_1

Measurement of Fock state decay time T_1 is similar in principle to measuring coherent state decay; the only difference is in state preparation. Fock $|1\rangle$ is prepared using the methods outlined in Ch. 2.3.3. After a variable delay, the resonator state is measured by the qubit (Fig. 3.14b), and the probability of finding the qubit in $|e\rangle$ is fit to a decaying exponential to yield the T_1 time. Due to the linearity of the resonator, coherent state T_1 and Fock state T_1 times are expected to be the same (Fig. 3.15b).

Single-Photon Ramsey Coherence

Similar to the transmon T_2 measurement, a Ramsey experiment can be performed in the Fock $|0\rangle, |1\rangle$ manifold. For a Fock T_2 measurement, a superposition of Fock $|0\rangle$ and $|1\rangle$ is

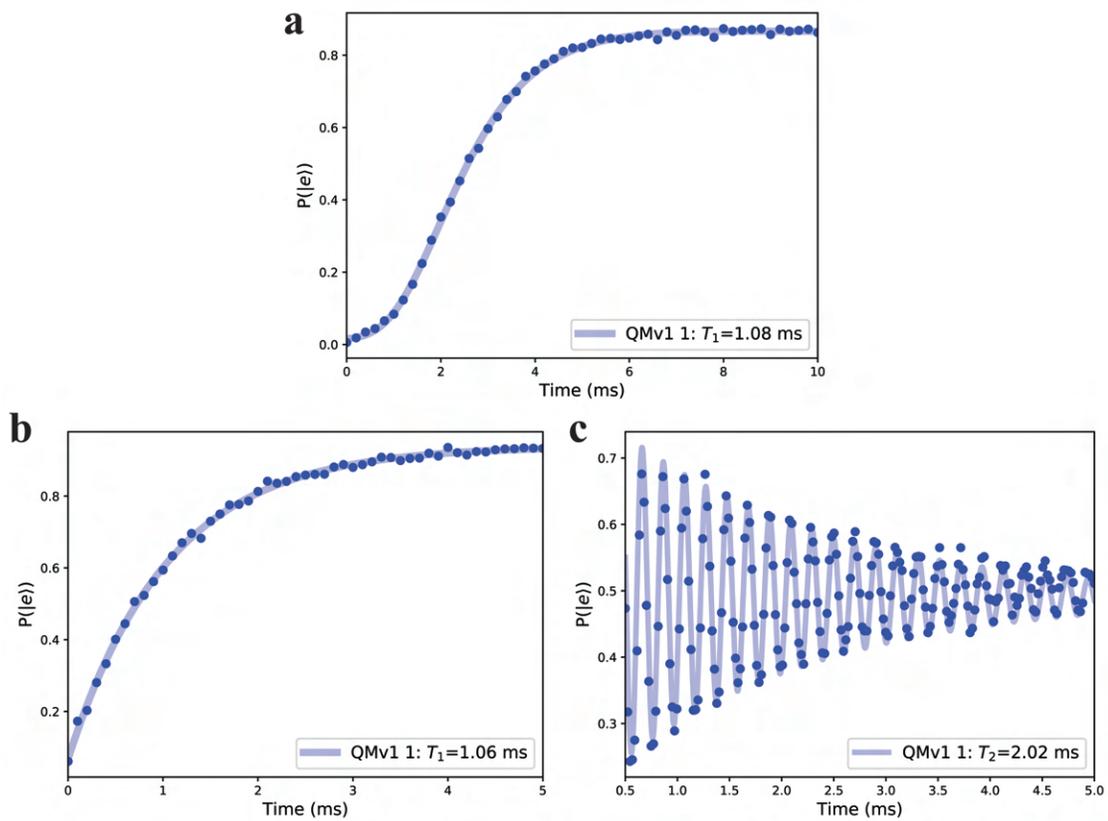


Figure 3.15: **Quantum memory coherence measurements.** **a** Coherent state decay. **b** Fock $|1\rangle$ decay. Fock T_1 and coherent state T_1 are approximately equal. **c** Fock T_2 coherence measurement.

prepared using the methods in Ch. 2.3.3 and is allowed to decohere over a variable delay time. Finally, in similarity to Ramsey interferometry, a small displacement $\beta_3 = 0.8e^{i\omega t}$ is applied to interfere with the resonator state, projecting it back towards the poles of the Bloch sphere defined by the Fock $|0\rangle$ and $|1\rangle$ computational space. The resonator state is then measured by the qubit (Fig. 3.14c). The displacement β_3 is also detuned by a small amount ω in order to produce oscillations in $P(|e\rangle)$ whose frequency is equal to the detuning between the drive and the resonator. A fit to an exponentially decaying sinusoid determines the Fock T_2 (Fig. 3.15c).

Chapter 4

Microwave Loss Characterization

As discussed in Ch. 1.3, improving the coherence of superconducting quantum circuits is crucial towards making better quantum processors. In the previous chapter, I discussed various sources of energy relaxation using the participation ratio model, which allows us to describe how resonator or transmon circuit geometry/design combines with intrinsic material/process-based loss channels to contribute to device quality factors. Many of these sources of loss have been explored substantially using microwave resonators and transmon qubits by measuring power dependence, temperature dependence, geometry dependence, and quasiparticle tunneling[90, 97, 98, 101–103, 106, 110, 114, 115, 117, 132–134]. Improvements in coherence have been made over the years by using materials thought to be lower in loss such as sapphire substrates over silicon[70, 135] and tantalum-thin-films over aluminum or niobium[39, 64]. Additionally, contamination-minimizing fabrication processes such as acid-based etching[90, 112] and substrate annealing[39, 110, 132] have also been explored. Finally, circuit geometry-based approaches have been shown to produce dramatic improvements by reducing energy participation in lossy regions, which had led to the rise of 3D transmon qubits[73] and cavity-based quantum memories with millisecond coherence times[53, 76, 80].

However, since the loss factors are unknown and the only parameter we can measure is

the quality factor, determination of loss factors is very difficult. A comprehensive understanding of relaxation mechanisms does not yet exist, which limits our ability to systematically explore new materials, processes, and circuit designs. Determining which sources of loss dominate in our devices and to what extent they do so is crucial to maximizing the performance of superconducting quantum circuits. It is therefore highly desired to develop a systematic approach to characterize loss factors in order to improve coherence in a predictable way.

Microwave resonators are the key to accomplishing this, and are a convenient loss characterization tools because they can be measured easily and quickly with high precision, and as discussed in Ch. 3, they are subjected to the same loss channels as transmon qubits. Additionally, resonators have a high dynamic range in their participation ratios; their geometries can be tuned easily to vary their sensitivities to particular source of loss. This feature has been heavily utilized to investigate various sources of loss in thin-film and cavity resonators[82, 105, 107–110, 112, 122]. In this chapter I will discuss the various strategies utilized in the field to understand losses in superconducting quantum circuits. I will begin by describing some simple methods used in the past to either set upper bounds on loss factors or determine if a particular source of loss is significant. I will then describe a more recently developed characterization approach that is capable of comprehensive characterization of loss factors. I call this method matrix loss characterization, and was first demonstrated in Calusine et al. [105]. I will show how we adapted this approaching multimode resonators to perform nearly full loss characterization with just one device[112]. Finally, I will show two experimental implementations of this approach, one to characterize losses in bulk superconductors and the other to characterize losses in thin-films.

4.1 Simple Loss Characterization Methods

A simple strategy to estimate the loss factors uses information that can be obtained through simulation; namely, participation ratios. The general approach revolves around designing a resonator that is sensitive to a specific source of loss but insensitive to others. The loss sensitivity of the i th source of loss is defined as

$$s_i = \frac{p_i \Gamma_i}{\sum_k p_k \Gamma_k} \leq 1 \quad (4.1)$$

Without knowledge of Γ_i , making the participation p_i as large as possible achieves the desired regime. If s_i can be made close to 1, we can assume that other sources of loss are negligible:

$$\frac{1}{Q_{\text{int}}} = p_i \Gamma_i + \sum_{k \neq i} p_k \Gamma_k \approx p_i \Gamma_i \quad (4.2)$$

Since p_i is calculated in simulation, measuring Q_{int} approximately determines Γ_i . However, the lack of prior knowledge about the loss factors makes it difficult to determine when the above condition is satisfied. In some special cases, however, $p_k = 0$ for all k . One example of this is shown in Reagor et al. [90], where a cylindrical cavity made of high purity aluminum was used to measure conductor loss $q_{\text{cond}} = X_s/R_s$. The TE₀₁₁ mode of this cavity has no electric energy stored on the surface and no currents flowing through the seam, rendering it insensitive to surface dielectric loss and seam loss; $p_{\text{MA}} = y_{\text{seam}} = 0$ (in practice, a shape perturbation was used to eliminate a mode degeneracy that resulted in a small but nonzero p_{MA}). This satisfies the condition required for Eq. (4.2) to be applicable; therefore, measurement of Q_{int} provides a good approximation for q_{cond} for bulk superconductors.

In most cases, however, it is impossible to design a resonator or resonance mode where

$p_k = 0$ for all k . As a result, measurement of Q_{int} can be used to place an upper bound on Γ_i :

$$\frac{1}{Q_{\text{int}}} = p_i \Gamma_i + \sum_{k \neq i} p_k \Gamma_k > p_i \Gamma_i \implies \Gamma_i < \frac{1}{p_i Q_{\text{int}}} \quad (4.3)$$

Loss factor upper bounds are useful to calculate with resonators of particularly large Q_{int} and p_i . An example of this is in Axline et al. [84], where the quality factors of various thin-film and cavity resonators were used along with their participations to determine how large a loss factor can be. This bound can then be used to prove a conservative estimate for the loss sensitivity of a particular resonance mode of interest. If $s_i \ll 1$, then the mode must be insensitive to the i th source of loss. Note that this only applicable if the resonator of interest was made using the same materials and fabrication processes as the resonator that was used to calculate the upper bound, in order to keep Γ_i fixed.

A significant drawback of the method of bounding loss factors is that it is often unclear whether the resonator being used to probe that source of loss is actually limited by it. A slightly more advanced approach is the method of participation scaling. This approach relies on designing multiple resonators that are made with the same materials and processes but differ in their geometry; specifically, the goal is to vary the participation of one source of loss while keeping the others fixed:

$$\frac{1}{Q_{\text{int}}^j} = p_i^j \Gamma_i + \sum_{k \neq i} p_k \Gamma_k \quad (4.4)$$

Each resonator, indexed by j , has a different value of p_i , but the same value of $p_{k \neq i}$. Therefore, by measuring how Q_{int} scales with p_i , we can determine Γ_i . At low values of p_i , Q_{int} would not change due to its insensitivity to the i th source of loss. As p_i is increased, at some point the Q_{int} would begin to decrease; the resonator is now unambiguously limited by the i th source of loss. This approach was demonstrated to study seam loss in 3D

cavities[122], indium bump-bond loss in thin-film devices[82], and surface dielectric loss in both thin film resonators and transmons[74, 104, 110, 136].

The above approaches to loss characterization have led to many significant insights over the years; however, they are all limited in that they can only study one type of loss at a time. While the participation scaling method can also be used to determine a background loss $\sum_{k \neq i} p_k \Gamma_k$ that is independent on the scaling of p_i , the background loss term is a combination of losses from many other sources of loss which cannot be distinguished from each other. Additionally, for many sources of loss it is very difficult if not impossible to independently scale them while leaving the others unchanged. In the next section I will describe a modern approach to loss characterization which is similar to the participation scaling method but allows all of the participations to vary; this allows us to extract multiple loss factors using the same set of resonators or resonance modes.

4.2 Matrix Loss Characterization

The matrix approach is based on the fact that the participation ratio model is simply a linear equation. As a result, if the loss of a resonator is determined by n sources of loss, then m appropriately resonators can be used to set up a system of equations:

$$\begin{pmatrix} Q_1^{-1} \\ Q_2^{-1} \\ \vdots \\ Q_m^{-1} \end{pmatrix} = \begin{pmatrix} p_{11} & p_{12} & \cdots & p_{1n} \\ p_{21} & p_{22} & \cdots & p_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ p_{m1} & p_{m1} & \cdots & p_{mn} \end{pmatrix} \begin{pmatrix} \Gamma_1 \\ \Gamma_2 \\ \vdots \\ \Gamma_n \end{pmatrix} \implies L_j = \sum_i^n P_{ji} \Gamma_i \implies \vec{L} = \mathbf{P} \vec{\Gamma} \quad (4.5)$$

where \vec{L} is a column vector of measured internal losses $L_j = 1/Q_j$, \mathbf{P} is an $m \times n$ participation matrix with elements P_{ji} , and $\vec{\Gamma}$ is a column vector that contains the loss

factors Γ_i . If $m \geq n$ (in other words, if the rank of \mathbf{P} is greater than or equal to the number of loss channels), then there exists a matrix \mathbf{P} that gives a nontrivial solution for $\vec{\Gamma}$. Therefore, by measuring the internal quality factors of the resonators and simulating their participation ratios, the loss factors can be determined.

Importantly, the measured L_j have some uncertainty associated with them. When solving this equation, the errors σ_{L_j} must be propagated onto the extracted Γ_i . This can be done using least-square minimization[137, 138]. We can write down the least-squares sum as

$$S = \sum_j \left(\sum_i \tilde{P}_{ji} \Gamma_i - \tilde{L}_j \right)^2 \quad (4.6)$$

where $\tilde{P}_{ji} = P_{ji}/\sigma_{L_j}$ and $\tilde{L}_j = L_j/\sigma_{L_j}$ are the measurement-error-weighted participation matrix and internal loss, respectively. In matrix form this can be expressed as

$$S = (\tilde{\mathbf{P}}\vec{\Gamma} - \vec{\tilde{L}})^T (\tilde{\mathbf{P}}\vec{\Gamma} - \vec{\tilde{L}}) \quad (4.7)$$

S is minimized when $\partial S/\partial \vec{\Gamma} = 0$:

$$\frac{\partial S}{\partial \vec{\Gamma}} = -2\vec{\tilde{L}}^T \tilde{\mathbf{P}} + 2\vec{\Gamma}^T \tilde{\mathbf{P}}^T \tilde{\mathbf{P}} = 0 \quad (4.8)$$

The solution is then

$$\vec{\Gamma} = \mathbf{C} \tilde{\mathbf{P}}^T \vec{\tilde{L}} \quad (4.9)$$

where $\mathbf{C} = (\tilde{\mathbf{P}}^T \tilde{\mathbf{P}})^{-1}$ is defined as the covariance matrix and is symmetric. We can then calculate the propagated error $\vec{\sigma}_{\vec{\Gamma}}$ as

$$\vec{\sigma}_{\vec{\Gamma}}^2 = \langle \delta \vec{\Gamma} \delta \vec{\Gamma}^T \rangle = \mathbf{C} \tilde{\mathbf{P}}^T \langle \delta \vec{\tilde{L}} \delta \vec{\tilde{L}}^T \rangle \tilde{\mathbf{P}} \mathbf{C}^T \quad (4.10)$$

where $\delta\Gamma_i = \sum_j \frac{\delta\Gamma_i}{\delta L_j} \delta L_j = \sum_j \sum_k C_{ik} \tilde{P}_{kj} \delta \tilde{L}_j$. Since $\langle \delta \tilde{L}_i \delta \tilde{L}_j \rangle = \langle \frac{1}{\sigma_i} \frac{1}{\sigma_j} \delta L_i \delta L_j \rangle = \delta_{ij}$, we arrive at the final expression for $\vec{\sigma}_{\tilde{\Gamma}}$:

$$\vec{\sigma}_{\tilde{\Gamma}}^2 = \mathbf{C} \tilde{\mathbf{P}}^T \tilde{\mathbf{P}} \mathbf{C}^T = \mathbf{C} \mathbf{C}^{-1} \mathbf{C}^T = \mathbf{C} \quad (4.11)$$

Therefore, the propagated error on the extracted loss factors are given by $\sigma_{\Gamma_i} = \sqrt{C_{ii}}$.

The least-squares minimization approach is elegant because an analytical solution exists. However, it makes no assumptions about the loss factors; they can be negative, which is unphysical for a linear resonator. Negative loss factors can be extracted in cases where the loss channels under consideration are over-constraining the loss model; i.e. the system of resonators as designed is insensitive to a specific source of loss. An alternate approach to solving the matrix problem is to instead utilize a numerical non-negative least-squares algorithm, which adds the constraint that $\Gamma_i \geq 0$. In order to estimate the error propagation, a Monte Carlo analysis is used[105, 107, 108]. Operationally, this proceeds as follow: first, a normal distribution of internal quality factors for each resonator is generated using the measured L_j as the mean and the measurement uncertainties σ_{L_j} as the standard deviation. Then, the non-negative least-squares algorithm is applied by sampling a set of L_j from the distributions for each resonator and extract a set of loss factors for each sample. This generates a distribution of loss factors whose mean and standard deviation reflect the solution to the matrix equation and propagated error. In the case of an over-constrained system, the Monte Carlo method still will not be able to extract all the loss factors, but it will be able to effectively provide an upper bound. In the case of a well-conditioned system, however, the Monte-Carlo method reproduces the results of the linear least-squared algorithm.

In order to fully bound the extracted loss factor, it is required that

$$\frac{\sigma_{\Gamma_i}}{\Gamma_i} = \frac{\sqrt{C_{ii}(\mathbf{P}, \vec{\sigma}_{\tilde{\mathbf{L}}})}}{\left[\mathbf{C} \tilde{\mathbf{P}}^T \tilde{\mathbf{L}} \right]_i} < 1 \quad (4.12)$$

The measurement sensitivity $\frac{\sigma_{\Gamma_i}}{\Gamma_i}$ is clearly dependent on both the geometric design of the loss characterization system (defined as the set of resonators that makes up the participation matrix) and the loss factors themselves. As a result, there is a range of loss factors for which the participation matrix is well-conditioned to solving the matrix equation with minimal propagated error. It is therefore necessary to ensure that this range of loss factors for which the system is solvable encompasses the loss factors we want to extract. This is accomplished by designing a set of resonators such that the rows of participation matrix are linearly independent; there must exist a subset of n resonators where each one is sensitive to a different source of loss. In practice, this can be difficult to do, as some sources of loss tend to scale together with regards to their participations. For example, in thin-film resonators the surface participations p_{SA} , p_{MS} , and p_{MA} all scale in the same way with geometry[74]. This is due to the fact that the charge density tends to accumulate at the edges of the thin-films; this region is where the three interfaces are in very close proximity to each other and also where the majority of the participation resides. Since this field behavior occurs on very small length scales ($< 1 \mu\text{m}$), it is insensitive to far-field boundary conditions and changes in resonator geometries that occur on much larger length-scales ($\sim 10 - 100 \mu\text{m}$). In such cases, attempting to distinguish between the interfacial loss factors results in a large amount of propagated error onto the loss factors, often times to the point where the fractional error on the loss factors exceed unity, which would yield an unbounded solution.

Strategies to develop well-conditioned matrices have been demonstrated in earlier works. TLS loss in the various surface and bulk dielectric regions in thin-film CPW res-

onators was studied in Calusine et al. [105], Woods et al. [107], and Melville et al. [108]. These works utilize the fact that CPW resonators have internal quality factors that change by over an order of magnitude as a function of photon number, which implies that they are dominated by TLS loss. As a result, the quality factors of these devices must be limited only by TLSs in the three surface dielectric interfaces and in the bulk of the substrate. This creates a loss model where only four sources of loss need to be considered; therefore, a minimum of four resonators, each with different geometries, were needed to extract the loss factors. These works utilized substrate etching to alter the scaling between p_{SA} , p_{MS} , and p_{MA} , which was possible due to their choice of using a silicon substrate, and were able to design a participation matrix that was well-conditioned to extracting the loss factors with acceptable amounts of propagated error. Another study was done here at Yale to study dielectric loss associated with a substrate. This technique, called the “dielectric dipper”, utilized a high-Q coaxial post cavity with a dielectric substrate inserted into it [106]. The substrate was attached to a piezoelectric positioner, allowing its insertion in to the cavity (and therefore participation) to be varied in situ. Each quality factor measurement at a different insertion is effectively a new resonance mode with a different set of participations. This defines a participation matrix as defined in Eq. (4.5) with m rows, where m is the number of insertions points for which the quality factors were measured. Because the substrate participation is significantly dependent on insertion, this scheme defines a participation matrix that is well-conditioned to extract dielectric loss with high sensitivity.

4.3 The Multimode Approach to Loss Characterization

The participation ratio model is a macroscopic model that assumes that Γ_i are globally spatially uniform. In reality, local material defects and localized anomalies as a result of processing during fabrication can result in inhomogeneous loss factors. This results in

device-to-device variation of resonator Q of nominally identical devices. In a multimode approach to loss characterization, a single device can manifest multiple resonance modes, each of which have different spatial distributions of electromagnetic field therefore different participation ratios[112]. This allows for the use of a single device to study multiple sources of loss, eliminating systematic errors due to device-to-device or run-to-run variation. Furthermore, by measuring multiple multimode devices, the variation in loss factors can be determined. The ability to reproduce the statistical distribution of loss factors allows for the evaluation of the consistency of a particular fabrication process and for the prediction of the range of quality factors to be expected for a device of given geometry[86].

In the following sections, I will review two implementations of the multimode approach: the forky whispering-gallery-mode resonator (FWGMR) described in Lei et al. [112] and the tripole stripline (TSL) describe in Ganjam et al. [86].

4.4 Forky Whispering-Gallery-Mode Resonator

The FWGMR is a multimode resonator designed to characterize losses in bulk superconducting materials. Specifically, the work presented in Lei et al. [112] focused on measuring losses associated with 6061 and high-purity (5N5) aluminum, which are commonly used in 3D cavity-based cQED experiments. Additionally, various surface treatments such as acid-based etching and diamond turning were explored. This work also demonstrated loss characterization in another multimode device called the ellipsoidal cavity, which was subjected to additional treatments such as thin-film coating. The ellipsoidal cavity is similar to the FWMGR and investigates the same sources of loss albeit with slightly different measurement sensitivity. I will not discuss the ellipsoidal cavity here in the interest of brevity and refer the reader to the published work of Lei et al. [112] for a more detailed description.

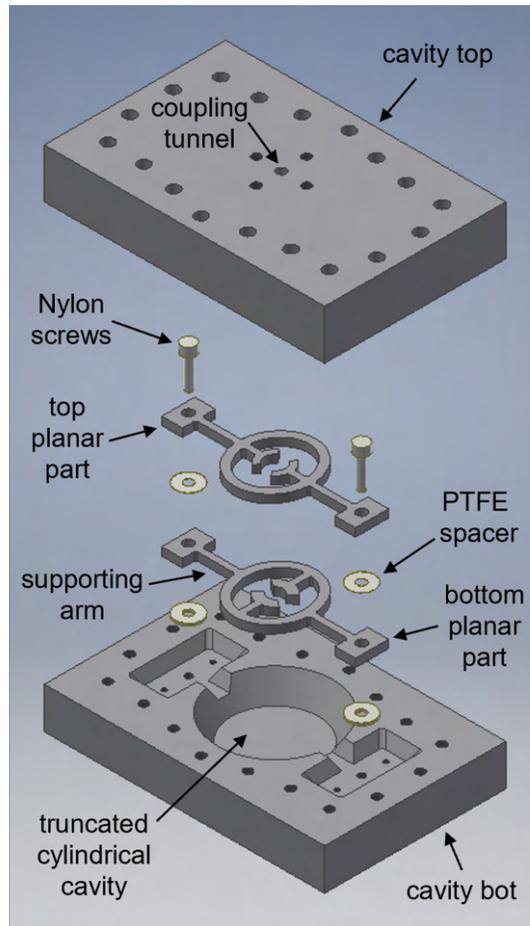


Figure 4.1: **Forky whispering-gallery-mode resonator (FWGMR)**. Exploded view: two planar parts stacked on top of each other and separated by Teflon spacers define the modes. Figure obtained from Lei et al. [112]

4.4.1 FWGMR device design

The FWGMR is a cavity resonator comprised of two planar components that are separated and galvanically isolated by Teflon spacers and assembled into a truncated cylindrical cavity enclosure (Fig. 4.1). Both the planar components and the enclosure are made of the bulk superconductor under study, machined using conventional CNC techniques. The shape of the planar parts define the modes of this device and their participations. The elliptical ring supports a whispering-gallery resonance mode. Since there are two elliptical rings stacked on top of each other, two whispering-gallery modes (WGMs) result: a common mode (CWGM) where the currents on both rings flow in the same direction, resulting in spatially diffuse electromagnetic fields that terminate at the walls of the enclosure (Fig. 4.2d). This mode is sensitive to the seam loss associated with closing the two halves of the enclosure together. The other WGM is the differential mode (DWGM), where the currents in the rings flow in opposite directions (Fig. 4.2c). This confines the field in the small spacing between the two spacers (the thickness of the spacers can vary due to part variation and thermal contraction and must be inferred through the frequency measurement of the mode; spacer thickness variation was 50 – 150 μm). The thickness of the spacers determine the strength and spatial confinement of the fields. Since both electric and magnetic fields are confined, this mode has high electrical energy density on the MA interface and high current density on the surface the metal. This mode is therefore sensitive to both the MA surface dielectric loss and conductor loss. Finally, the forks in the middle of the ring are arranged such that the forks from each planar part overlap each other, resulting in a parallel-plate capacitance. The currents the flow between the forks form an inductive element, giving rise to the differential fork mode (DFM) (Fig. 4.2e). Because the electrical energy is confined between the forks and the magnetic energy is less confined, this mode becomes sensitive only to MA surface dielectric loss.

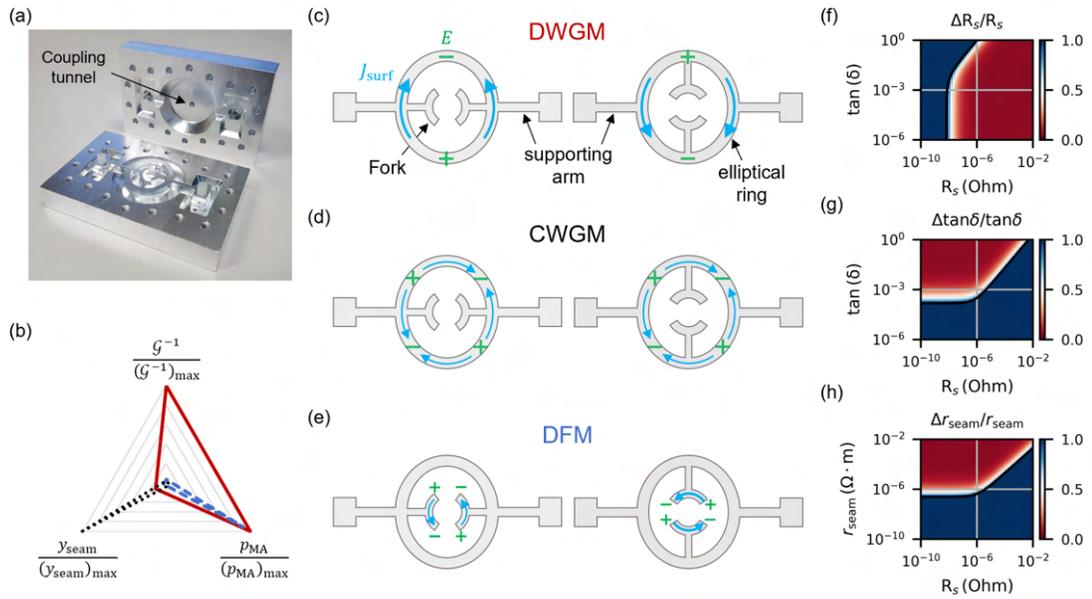


Figure 4.2: **FWGMR modes and measurement sensitivity.** **a** Assembled FWGMR made with 5N5 aluminum after chemically etching. **b** Radar plot showing relative participation ratios between the modes. **c-e** Diagrams of the two planar parts, showing current flows and polarities for the three modes. Green $+/-$ symbols represent the relative polarities of the electric fields; blue arrows show the relative direction of current flow. **f-h** Measurement sensitivity plots of the loss factors that define the range of loss factors for which the loss extraction can be done without propagating a fractional error of greater than unity. Figure obtained from Lei et al. [112]

4.4.2 Participations & measurement sensitivity

The participation matrix for this system is given in Table 4.1, assuming 100 μm spacer thickness. The first column is the conductor participation G^{-1} , the second column is p_{MA} (assuming a region thickness of 3 nm), and the third column is y_{seam} . The participations of the DWGM is given by the first row, the DFM is given by the second row, and the CWGM is given by the third row. The radar plot in Fig. 4.2b provides a visual representation of the relative participations of the three types of modes. The CWGM has relatively high y_{seam} while having low $G^{-1} \propto p_{\text{cond}}$ and p_{MA} . The DFM similarly has high p_{MA} while having low y_{seam} and G^{-1} . The DWGM is the only mode with high G^{-1} .

Table 4.1: FWGMR Participation Matrix, assuming 100 μm spacer thickness

Mode	Participations		
	G^{-1}	p_{MA}	y_{seam}
DWGM	$2.8 \times 10^{-1} \Omega^{-1}$	3.8×10^{-6}	$2.7 \times 10^{-4} (\Omega\text{m})^{-1}$
DFM	$8.9 \times 10^{-3} \Omega^{-1}$	3.5×10^{-6}	$7.1 \times 10^{-5} (\Omega\text{m})^{-1}$
CWGM	$5.5 \times 10^{-3} \Omega^{-1}$	1.5×10^{-7}	$2.1 \times 10^{-3} (\Omega\text{m})^{-1}$

However, the participations alone do not determine how well-conditioned the matrix is. To determine the range of loss factors over which this system propagates minimal error, measurement sensitivity plots are made, as in Figs. 4.2f-h. To make these plots, two of the loss factors are varied while the third is fixed (technically, since there are three sources of loss the full sensitivity plot is a color plot on a three-dimensional surface; for illustrative purposes, $r_{\text{seam}} = 1/g_{\text{seam}} = 10^2 \mu\Omega\text{m}$ for Figs. 4.2f-g, and $\Gamma_{\text{MA}} = \tan \delta = 5 \times 10^{-2}$ for Fig. 4.2h). For each point, the matrix equation is solved and the error is propagated assuming a measurement error of 5%. From these plots, we can identify a measurement sensitivity for the system where $\sigma_{\Gamma_i}/\Gamma_i < 1$; the region in red encompasses a range of loss factors where they are resolvable. The contour that divides the red (resolvable) and the blue (unresolvable) regions is determined by the influence of the other sources of loss. In Fig.

4.2f, for example, R_s as low as $10^{-6} \Omega$ can be resolved with fractional error < 1 ; however, lower values of R_s also require lower values for $\tan \delta$ in order to keep the DWGM from being dominated by MA loss. Once both $\tan \delta$ and R_s are sufficiently low, the residual seam loss dominates the loss of the DWGM, and the system loses measurement sensitivity for conductor loss. Prior studies of conductor loss in etched high-purity aluminum have measured $R_s \sim 10^{-7} \Omega$ [90], to which the system is sensitive as long as $\tan \delta < 10^{-1}$.

It should be noted that the FWGMR is a complex structure that consist of other modes that could also be included in the participation matrix. For example, a hybridized mode exists called the differential forky whispering-gallery mode (DFWGM) where the electric field is stored both between the forks and the elliptical rings. Additionally, there are cavity modes that are highly sensitive to seam loss. Finally, all of these modes each have higher order harmonics that can be measured to potentially increase the measurement sensitivity or probe the spatial variation of the loss factors.

4.4.3 Device measurement & loss extraction

Table 4.2 shows the measured single-photon quality factors and extracted loss factors for a FWGMR device made with 6061 aluminum. Measuring the quality factor of these devices was difficult because of the device's sensitivity to vibrational noise, which caused the spacing between the two planar parts to fluctuate, thereby causing the resonance frequencies of the DFM and DWGM modes to fluctuate. This noise was suppressed if not eliminated entirely by turning off the pulse-tube cooler of the dilution refrigerator during measurement, allowing the use of the VNA to measure the quality factor due to energy loss.

For this device, loss factor extraction was performed using the second-order DFM mode (DFM-2) and the fundamental DWGM and CWGM modes. However, in principle, other modes of the device can be used for the loss extraction. One would expect that using

Table 4.2: Extracted loss factors at single-photon powers of 6061 Al measured by a FWGMR resonator with spacer thickness 100 μm . The mode name followed by a “1” or “2” corresponds to the fundamental and second order harmonics, respectively. Measurement error was fixed at 5%.

Device	Material	Mode	Freq (GHz)	Q_{int} ($\times 10^6$)	R_s ($\mu\Omega$) $\tan \delta$ r_{seam} ($\mu\Omega\text{m}$)
FWGMR-4 (100 μm)	6061 Al	DWGM-1	5.858	0.45	6.48 ± 0.43
		DFM-2	9.199	2.2	0.11 ± 0.01
		CWGM-1	10.863	7.4	39.1 ± 3.5

other modes would still yield the same loss factors because the geometric dependence of loss is appropriated to the participation matrix. For this device, however, the extracted loss factors varied when different modes were used for the loss analysis. There are several possible explanations for this. For one, imperfections in the machining and assembly of the device could result in geometrical defects in the resonator that are not captured by the simulations. This could result in an inaccurate participation matrix. Additionally, spatial inhomogeneity in machining quality or contaminant density could result in inhomogeneous loss factors. Since different modes have different spatial distributions of their electromagnetic fields, the different modes would effectively be sensitive to different loss factors. Finally, the loss factors are also assumed to be frequency independent or very weakly so. Since the different modes have different frequencies, the extracted loss factors may vary due to the frequencies of the modes that are used to extract them.

4.4.4 Comparison of loss factors of different materials & processes

We can compare the loss factors of 6061 aluminum with those of both 6061 and 5N5 aluminum subjected to different surface treatments. In Fig. 4.3, we can see how the combination of materials and fabrication processes affect the intrinsic material loss. Additionally, we can see device-to-device variation due to loss factor inhomogeneity. There seems to

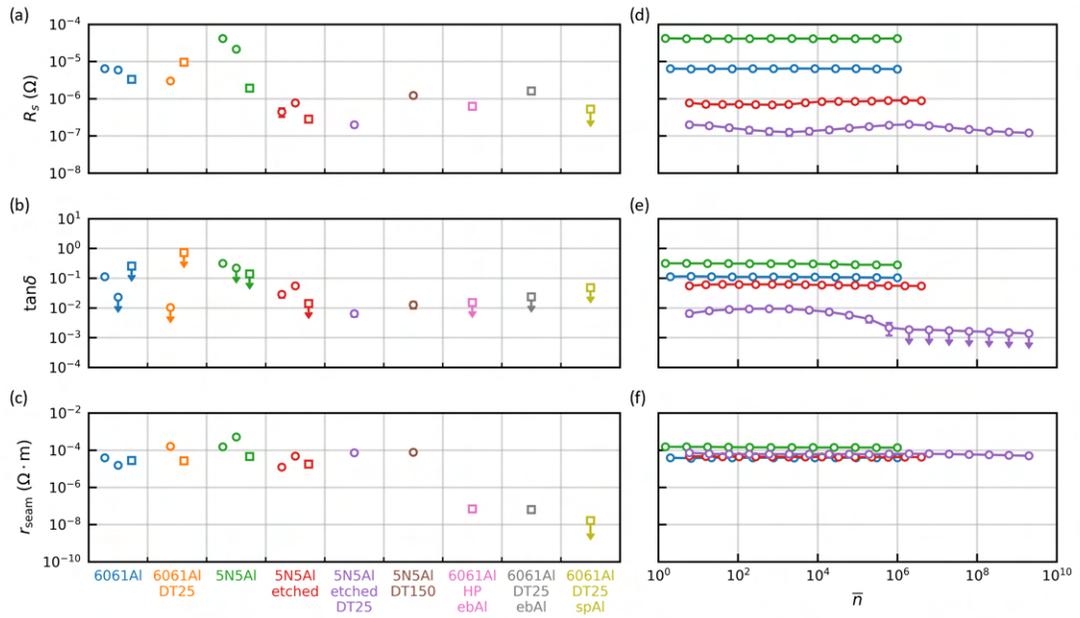


Figure 4.3: **Extracted single-photon loss factors for bulk superconductors.** **a-c** Extracted surface resistance **(a)**, MA loss tangent **(b)**, and seam resistance per unit length **(c)** of bulk-machined aluminum of different alloys (6061 and 5N5) subjected to different surface treatments (diamond turning, chemical etching, and thin-film coating). These loss factors were obtained from FGMR and ellipsoidal cavity measurements. **d-f** Photon number dependence of the loss factors. Figure obtained from Lei et al. [112]

be little to no difference in the loss factors between untreated 6061 Al (Figs. 4.3a-c, blue) and 5N5 Al (Figs. 4.3a-c, green); however, there is significant sample-to-sample variation within 5N5 devices, as the surface resistance varies by over an order of magnitude. The device-to-device consistency improved dramatically (Figs. 4.3a-c, red) after applying a chemical etching treatment (Transene Aluminum Etchant Type A heated to 50 °C for two hours[90, 112]) to the 5N5 Al device that removed $\approx 100 \mu\text{m}$ of material from the surface of the metal[90]. Additionally, etching 5N5 Al reduces the surface resistance by a factor of around 40, and the MA loss tangent by a factor of around 5. These loss factor improvements allow 3D post cavities to serve as highly coherent quantum memories, as without them the cavities would not have been able to reach millisecond coherence times[53, 90].

To understand the source of these improvements, 5N5 Al samples before and after etching were analyzed using transmission electron microscopy (TEM) and electron energy loss spectroscopy (EELS) (Fig. 4.4). The unetched samples have substantial damage around 100 nm below the surface. This damage was likely generated during the machining processes and contributes to the increased surface resistance. The etching treatment removes the damaged material from the surface, exposing the previously undisturbed material underneath, thereby yielding lower surface resistance. Additionally, the surface oxide that grows on unetched 5N5 Al has thickness ranging from 3 nm to 16 nm (Fig. 4.4a), which rescales the MA loss tangent to an average of $\approx 1.1 \times 10^{-1}$ nm. After etching, the oxide is more uniform and thinner (Fig. 4.4b) with thickness ≈ 3 nm (it should be noted that larger scale variations in oxide thickness have not yet been explored). Comparing the two, we see that the factor of 5 reduction in MA loss tangent originates from both a reduction in oxide thickness and a factor of 2 reduction in intrinsic loss.

The ability to understand how materials and processes affect intrinsic loss motivates the exploration of other surface treatments. In particular, diamond-turning, a precision machining processes to uniformly remove material from the surface to create an extremely

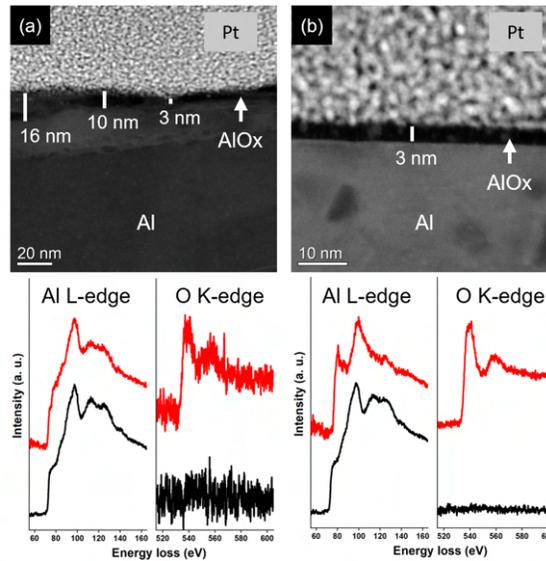


Figure 4.4: **TEM of unetched and etched bulk Al.** **a** Cross-sectional HAADF-STEM image of unetched 5N5 aluminum with EELS spectra showing that the region labelled as “AlO_x” is an oxide of aluminum. **b** Etched 5N5 Al. The oxide seems to be more uniform in thickness; however, larger scale variation has not yet been explored. TEM done by Kim Kisslinger and Sooyeon Hwang at Brookhaven National Laboratory. Figure obtained from Lei et al. [112]

smooth finish, could potentially be an alternative to chemical etching. Although no improvement was seen in 6061 Al after a 25 μm diamond-turning (Figs. 4.3a-c, orange), 5N5 aluminum saw significant reductions in the loss factors after being subjected to a 150 μm diamond-turning (Figs. 4.3a-c, brown) and in fact produces similar results as chemical etching. This indicates that diamond-turning may be a chemical-free alternative to processing 5N5 Al for high-Q cavities, as they both accomplish the same task of removing material damaged and contaminated by conventional machining processes. Moreover, an additional diamond-turning on an already etched 5N5 Al surface can further reduce both the MA loss tangent and the surface resistance (Figs. 4.3a-c, purple), indicating that the etching process may not have removed enough damaged or contaminated material from the surface.

Another alternative to etching or diamond turning was shown in this work to be thin-

film coating. Instead of removing the damaged and contaminated material, thin-film aluminum is deposited using both electron-beam evaporation and sputtering. The currents in the mode then flow through the higher-quality thin film instead of the more lossier bulk Al surface. Again, this surface treatment seems to yield similar surface resistances and MA loss tangents as diamond turning and chemical etching. However, surprisingly, thin-film coatings seem to produce much lower seam resistances than the other surface treatments. This may be due to the relatively higher ductility of thin-film aluminum over the bulk, allowing it to form more points of direct metal-to-metal contact within the seam. This also indicates that using thin-films might be the key to achieving ultra-low seam or joint losses, as was demonstrated in Lei et al. [82], and that utilizing such techniques would be extremely useful towards reducing seam losses in cavity architectures.

By measuring the quality factors at higher photon numbers, we can perform the loss extraction and describe the intrinsic power dependence of the loss factors themselves (Figs. 4.3d-f, purple). Here, we see that for most of the materials/process combinations, the extracted loss factors do not seem to have a much power dependence. However, for 5N5 Al that was etched and then diamond-turned, the MA loss does seem to have power dependence; it seems as if the MA losses from the other samples were limited by a power-independent loss mechanism. In the 5N5 Al etched and diamond-turned sample, the processes have removed or suppressed the power-independent source of loss, revealing the TLS-like behavior.

4.4.5 Conductor loss in bulk Al

The extraction of R_s allows us to make some statements about the mechanisms of conductor loss. In Ch. 3.3.1, I showed that for aluminum resonators operated at 20 mK, thermal quasiparticle loss is negligible. However, assuming the surface resistance is due to nonequilibrium quasiparticles, we can estimate the residual quasiparticle fraction in

etched 5N5 Al using Eq. (3.18). From Reagor et al. [90], we obtain a penetration depth $\lambda \approx 50$ nm, and to calculate the normal-state resistivity ρ_n just before superconducting transition, we use the following relation[93]:

$$\rho_n l = \frac{m^* v_F}{n_0 e^2} = 4.31 \times 10^{-16} \Omega \text{m}^2 \quad (4.13)$$

The right-hand side is computed for aluminum[95], where m^* is the effective mass of charge carriers, v_F is the Fermi velocity, and n is the number density of charge carriers. We can calculate the mean free path l from the measured penetration depth and how it relates to the London length λ_0 and coherence length ξ_0 :

$$\frac{1}{\lambda^2} = \frac{1}{a \lambda_0^2} \left[\frac{\pi}{2} - \frac{\cosh^{-1}(a)}{\sqrt{a^2 - 1}} \right] \quad (4.14)$$

where, $a = \pi \xi_0 / 2l > 1$ in the dirty limit of superconductivity. The mean free path reflects the characteristic length over which charge carriers can flow before being scattered. The coherence length and London length are intrinsic properties of an ideal superconductor in the limit of large mean free path and are given by $\xi_0 = \hbar v_F / \pi \Delta = 1.6 \mu\text{m}$ and $\lambda_0 = \sqrt{m^* / \mu_0 n_0 e^2} = 16$ nm for aluminum. For etched 5N5 Al, the above expressions give $l \approx 180$ nm, which is much lower than ξ_0 and confirms that the superconductor is in the dirty limit. With this value of l , we obtain $\rho_n^{20 \text{ mK}} = 2.4 \times 10^{-9} \Omega\text{m}$, which is around 11 times lower than the room-temperature resistivity of $\rho_n^{300 \text{ K}} = 2.65 \times 10^{-8} \Omega\text{m}$. Now we can calculate x_{qp} from Eq. (3.18). Assuming an average $R_s = 5 \times 10^{-7} \Omega$ and $\omega / 2\pi = 5$ GHz over the devices measured in this work, we calculate $x_{qp} \approx 3 \times 10^{-4}$. This is orders of magnitude larger than what has been measured in Al-based transmon qubits (made with thin-films) through quasiparticle tunneling experiments ($x_{qp} = 10^{-8} - 10^{-10}$)[97, 98, 134]. The reason for this is still unclear, but may not be intrinsic to bulk superconductors. It has been shown in bulk niobium TESLA cavities that residual resistances of $R_s < 5$ n Ω can

be achieved (assuming conductor loss is power independent)[77], which corresponds to $x_{qp} \approx 7 \times 10^{-8}$. The SRF community over the years have developed highly specialized surface treatments for niobium to achieve this; perhaps there is more to explore for high-purity aluminum.

Instead of assuming that bulk aluminum is limited by excess quasiparticles, we can assume they are limited by vortices, and use R_s to calculate what the ambient cooling magnetic field would need to be to explain it through Eq. (3.19). The critical magnetic field B_c is dependent on the effective penetration depth λ and the effective coherence length ξ . This is different from ξ_0 which only applies in the limit where $l \gg \xi_0$. ξ can be calculated from x_0 and l using the following relation:

$$\frac{1}{\xi} = \frac{1}{\xi_0} + \frac{1}{l} \quad (4.15)$$

Using $l = 180$ nm, we find that $\xi \approx 162$ nm, which gives $B_c = 0.026$ T. This therefore implies that if the residual resistance $R_s = 5 \times 10^{-7} \Omega$ of etched 5N5 Al is due to vortex loss, then the device must have been subjected to a cooling field of $B_0 \approx 2 \times 10^{-6}$ T. However, this is inconsistent with measured magnetic field strengths inside our Cryo-erm magnetic shields of $B_0 \approx 10^{-7}$ T, which would imply a vortex flow resistance of $R_{\text{vortices}} \approx 25$ n Ω , assuming that the cooling field remains constant from room temperature to 20 mK. Moreover, if R_s was limited by vortex loss, more effective magnetic shielding should be able to dramatically reduce it. This experiment was conducted and detailed in Lev Krayzman's thesis[83], where a degaussed magnetic shield with residual magnetic field strengths of $B_0 \approx 10^{-8}$ T at room temperature was used while it was nested in an outer protective magnetic shield. The quality factors of the FWGMRs and ellipsoidal cavities did not appear to change when measured inside this shielding setup, indicating that reducing the cooling field by a factor of 10 did not have an affect on R_s . As a result, it is

highly likely that the residual resistance we see in surface-treated 5N5 aluminum is due to excess quasiparticles of unknown origin.

4.4.6 Predicting losses in 3D cavity resonators

The usefulness of the extracted loss factors lies in their ability to predict the loss of a resonator of a given geometry. In Fig. 4.5a, Q_{int} was predicted for a variety of resonator geometries including the FWGMR and ellipsoidal cavity. The errors on the loss factors propagate to error on the predicted Q_{int} , which implies that the origin of device-to-device variation lies in the intrinsic variations in the loss factors. The predicted quality factors of the FWGMR and ellipsoidal cavity should be self-consistent, as these modes themselves are used to extract the loss factors. The predicted quality factors of the rectangular cavity in the “V2” geometry (see Brecht et al. [122]) and the coaxial post cavity (“ $\lambda/4$ ”) are between 50 – 100 million, consistent with other experimental observations of identical cavity geometries[53, 75, 106, 122]. This demonstrates that the loss model can predict with reasonable accuracy the losses of resonators of dramatically varying geometries.

In Fig. 4.5b, a loss budget is made by calculating the fractional contribution of each source of loss to the total loss in the resonator, $p_i \Gamma_i / \sum_j p_j \Gamma_j$. Here, we can confirm the loss sensitivity of the FWGMR modes, as the DWGM and DFWGM modes are both almost equally limited by conductor and MA dielectric loss, whereas the DFM and CWGM modes are dominated by MA and seam loss, respectively. For the post cavity in particular, MA dielectric loss and conductor loss seem to limit the quality factor almost equally, indicating that future coherence improvements in the 3D cavity architecture necessitate improvements to both sources of loss.

The demonstration of the FWGMR validates the loss model and general matrix approach to loss characterization of bulk superconductors. By extracting the loss factors, we have developed a comprehensive understanding of what limits the Q_{int} of 3D cavities and

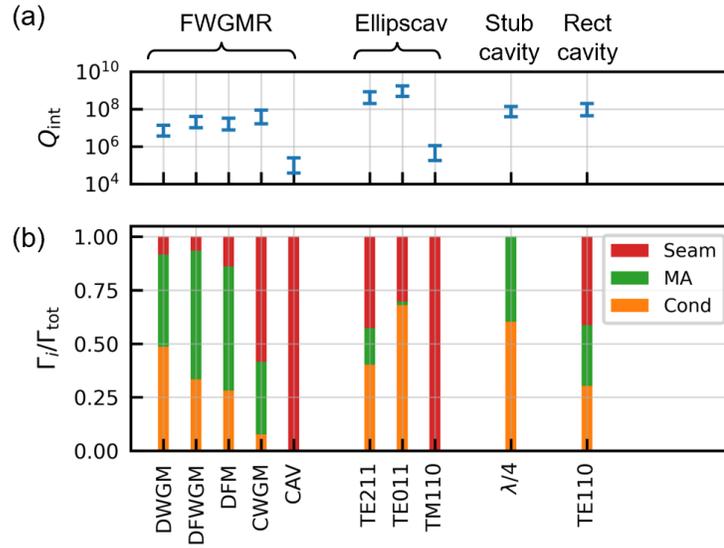


Figure 4.5: **Predicted quality factor and loss budget of 3D cavity resonators.** **a** Predicted single-photon quality factor based on participations and extracted loss factors for etched 5N5 Al. The uncertainties of the loss factors propagate onto the predicted Q , giving a range quality factors to expect in a device of that geometry. **b** Single-photon loss budget for the modes of the FWGMR, ellipsoidal cavity, $\lambda/4$ coaxial post cavity, and a rectangular cavity resonator. The fractional loss contribution is given by $p_i\Gamma_i/\sum_j p_j\Gamma_j$. Figure obtained from Lei et al. [112]

to what extent they limit a resonator of given geometry. This provides a path towards improving coherence by developing new processes or exploring new geometries whose loss sensitivities may be more favorable. The developments described in this section motivated a parallel study to investigate and characterize losses in thin-film circuits, which I will describe in the next section.

4.5 Tripole Stripline

The tripole stripline (TSL) is a multimode stripline resonator designed to characterize losses in thin-film superconductors[86]. The TSL is measured in a coaxial tunnel waveguide package made of machined 5N5 Al that has been chemically etched (see Fig. 3.8). The modes of this device can distinguish between surface losses, bulk dielectric loss, and package losses in on-chip circuits. This loss characterization device was used to measure and compare the loss factors associated with different thin films and fabrication processes. Additionally, sapphire substrates that differed by their growth method and preparation were compared by measuring their bulk dielectric loss factors. The ability to distinguish between these sources of loss allowed us to obtain a comprehensive understanding of the roles of materials, fabrication processes, and device geometry towards determining the relaxation times of thin-film superconducting quantum circuits. This knowledge can then be used to predict the losses of transmon qubits, as well as optimize devices to maximize coherence, which will be discussed in Ch. 5. This subsection follows the work in Ganjam et al. [86].

4.5.1 Multimode stripline device design

The tripole stripline consists of three striplines patterned adjacently to each other on a sapphire substrate (Fig. 4.6b). The size of the three striplines and the spacings between

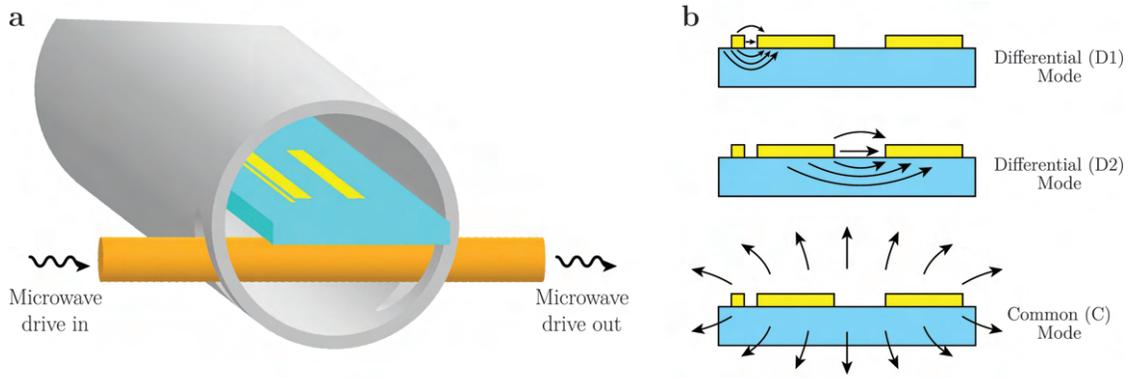


Figure 4.6: **Tripole striplines in the coaxial tunnel architecture.** **a** Thin-film striplines are patterned on a substrate and loaded into a cylindrical tunnel made of high-purity aluminum. A transverse coupling pin capacitively drives the resonators in hanger configuration. **b** Cross-sectional view of the tripole stripline, showing the arrangement of the strips and electric field behaviors for each mode. While the electric field of the D1 mode is confined mostly on the surface, the electric field of the D2 mode penetrates far deeper into the bulk, rendering it sensitive to losses over a significant portion of the bulk of the substrate. Figure obtained from Ganjam et al. [86]

them determine the spatial distributions of the electromagnetic fields of the modes, which in turn determine their sensitivities to particular sources of loss. There is one narrow stripline and two wide striplines. The narrow stripline is placed very close to one of the wider striplines; this gives rise to the D1 differential mode whose fields are highly localized within the small spacing between them. This leads to a large energy density localized in the surface, giving this modes very high surface dielectric participation. The large spacing between the two wide striplines supports a D2 differential mode whose fields are more dilute than that of the D1 mode, but still has a large bulk dielectric participation; rendering it sensitive to bulk loss. Finally, a common mode exists where the electric field terminates at the walls of the tunnel package; this results in relatively high participation in the package's MA interface, the conductor, and the seams generated by closing the tunnel.

The design of the TSL evolved over the course of this project. Originally, a device consisting of two striplines instead of three; this device is called the adjacent stripline (ASL). The ASL contains two striplines that are spaced by a distance $d = 10 \mu\text{m}$ apart.

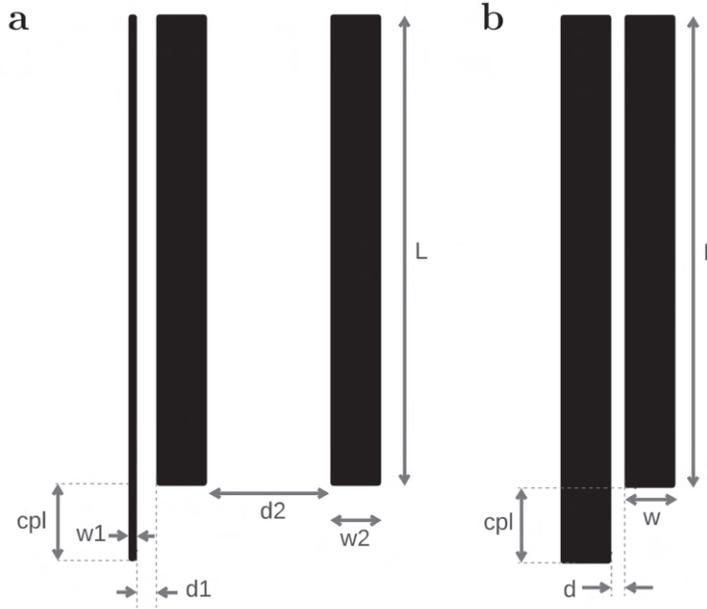


Figure 4.7: **TSL & ASL device design.** **a** Tripole stripline (TSL). **b** Adjacent stripline (ASL). Figure obtained from Ganjam et al. [86]

This device manifests two modes; a surface-sensitive differential (D) mode and a bulk and package-sensitive common (C) mode. At the time, it was thought that the package losses were insignificant and that the C mode would be dominated by bulk dielectric loss; the measurement and loss extraction with this device revealed bulk loss factors that were inconsistent with other works[106] and other measured devices on sapphire. The TSL was developed as a result to resolve this issue, and became the standard loss characterization device for this project. Moreover, ASL devices did not go to waste; by measuring TSLs in the same package as the ASLs, the package losses could be distinguished and subtracted from the ASL measurements, allowing the ASL to be used to extract bulk dielectric loss.

The basic layout of the TSL and ASL are shown in Fig. 4.7. The TSLs had three different designs that differed slightly from each other in terms of stripline length (L), wide conductor width (w_2), and narrow conductor spacing (d_1). The Dimensions the TSL and ASL are given in Table 4.3. Between TSLv1 and v2, the length of the stripline was decreased in order to increase the mode frequency to benefit from larger HEMT gain. The

Table 4.3: Tripole & adjacent stripline dimensions. ¹ ASLv1 was meandered. Table obtained from Ganjam et al. [86]

Dimension	Device Type			
	TSLv1	TSLv2	TSLv3	ASLv1
cpl	0.5 mm	0.5 mm	0.5 mm	0.3 mm
L	14 mm	12 mm	12 mm	14 mm ¹
$w1$	10 μm	10 μm	10 μm	-
$w2$	400 μm	400 μm	100 μm	-
$d1$	10 μm	10 μm	20 μm	-
$d2$	1200 μm	1200 μm	1200 μm	-
d	-	-	-	10 μm
w	-	-	-	150 μm

TSLv3 design was used for the aluminum TSLs, which were patterned with electron-beam lithography, and suffered from a lack of consistency and reproducibility when writing the TSLv1/2 patterns. It was observed that these caused by local resist heating effects and delamination; as a result, $d1$ was increased and $w2$ was decreased in order to increase device yield. Additionally, in all devices, the narrow stripline was longer than the other two striplines by length cpl , which introduced a field perturbation in the D1 mode that increases its coupling to the drive line, allowing the mode to be measured in the critical coupling regime.

The ASL device is simpler than the TSL; it contains two striplines that are spaced by a distance $d = 10 \mu\text{m}$ apart. The ASL's D mode is analogous to the TSL's D1 mode; the C mode is analogous to the TSL's C mode. There is no analogous D2 mode. The ASL also has one of its striplines slightly longer than the other in order to facilitate coupling to the D mode. Finally, it should be noted that the ASL is a meandered stripline, which confines its field slightly in order to reduce sensitivity to seam loss.

Table 4.4: TSL and ASL participation matrices, calculated using the methods in Appendix C. Table obtained from Ganjam et al. [86]

Device	Type	Substrate Thickness (μm)	Mode	Freq (GHz)	Participation Matrix					$y_{\text{seam}}^{\text{seam}} (\Omega\text{m})^{-1}$
					p_{surf}	p_{bulk}	$p_{\text{pkg,cond}}$	$p_{\text{pkg,MA}}$		
AM22, DZ22	TSLv1	430	D1	4.52	1.2×10^{-3}	0.90	4.3×10^{-8}	5.9×10^{-10}	7.9×10^{-9}	
			D2	5.34	3.5×10^{-5}	0.73	3.4×10^{-6}	2.9×10^{-8}	1.3×10^{-8}	
			C	6.64	2.2×10^{-5}	0.38	1.2×10^{-5}	1.2×10^{-7}	3.1×10^{-6}	
A23AI	TSLv3	650	D1	5.24	9.0×10^{-4}	0.90	1.5×10^{-7}	1.6×10^{-9}	8.6×10^{-6}	
			D2	5.51	8.9×10^{-5}	0.85	1.7×10^{-6}	1.4×10^{-8}	3.4×10^{-8}	
			C	6.68	5.2×10^{-5}	0.54	8.7×10^{-6}	9.7×10^{-8}	9.3×10^{-6}	
EF21	TSLv1	530	D1	4.52	1.2×10^{-3}	0.90	4.9×10^{-8}	6.1×10^{-10}	1.5×10^{-8}	
			D2	5.14	3.4×10^{-5}	0.76	3.6×10^{-6}	2.9×10^{-8}	2.0×10^{-8}	
			C	6.45	2.1×10^{-5}	0.41	1.2×10^{-5}	1.2×10^{-7}	5.8×10^{-6}	
EC21	ASLv1	530	D	3.68	6.7×10^{-4}	0.90	6.5×10^{-8}	4.8×10^{-10}	3.2×10^{-8}	
			C	6.03	2.8×10^{-5}	0.55	4.4×10^{-6}	7.5×10^{-8}	1.3×10^{-5}	
R22	TSLv1	650	D1	4.52	1.2×10^{-3}	0.90	6.0×10^{-8}	6.5×10^{-10}	2.5×10^{-8}	
			D2	4.97	3.4×10^{-5}	0.79	3.9×10^{-6}	2.9×10^{-8}	2.6×10^{-8}	
			C	6.27	2.1×10^{-5}	0.43	1.8×10^{-5}	1.3×10^{-7}	9.4×10^{-6}	
BF22	TSLv2	650	D1	5.25	1.2×10^{-3}	0.90	8.1×10^{-8}	8.5×10^{-10}	4.0×10^{-8}	
			D2	5.74	3.5×10^{-5}	0.80	3.8×10^{-6}	2.7×10^{-8}	3.6×10^{-8}	
			C	7.13	2.2×10^{-5}	0.45	1.3×10^{-5}	1.3×10^{-7}	1.4×10^{-5}	

4.5.2 TSL/ASL participations & measurement sensitivity

The participation matrices for each ASL and TSL device measured in this work are given in Table 4.4. Devices were fabricated on substrates of different thicknesses; this can slightly change the bulk dielectric and package participations, but leave the other participations relatively unchanged. Regardless, the striplines are susceptible to many different sources of loss. The C modes of the TSL and ASL have high $p_{\text{pkg}_{\text{cond}}}$, $p_{\text{pkg}_{\text{MA}}}$, and y_{seam} , while the D1 (TSL) and D (ASL) modes have high $p_{\text{surf}} = p_{\text{SA}} + p_{\text{MS}} + p_{\text{MA}}$. The D2 mode of the TSL, on the other hand, has high bulk dielectric participation p_{bulk} while have low surface dielectric and package participations.

The TSL has three modes; however, there are seven loss channels (p_{SA} , p_{MS} , p_{MA} , p_{bulk} , y_{seam} , $p_{\text{pkg}_{\text{cond}}}$, and $p_{\text{pkg}_{\text{MA}}}$), and there are not enough modes to distinguish between all of them. Two strategies are used to effectively shrink the participation matrices into square matrices with linearly independent rows. The first us to define a total surface loss $1/Q_{\text{surf}} = p_{\text{surf}}\Gamma_{\text{surf}}$, where $p_{\text{surf}} = p_{\text{SA}} + p_{\text{MS}} + p_{\text{MA}}$ is the sum of surface dielectric participations from the three interfaces associated with thin-film superconductors patterned on a substrate. These surface participations are calculated in the same way as described in Ch. 3.3.2 using the same assumed $t_{\text{surf}} = 3$ nm and $\epsilon_r = 10$. $\Gamma_{\text{surf}} = \sum_{k=\text{SA,MS,MA}} \frac{p_k}{p_{\text{surf}}} \tan \delta_k$ is then the corresponding surface loss factor, which differs from how surface loss is described in other studies[105, 107, 108] that attempt to independently characterize the surface loss factors $\tan \delta_{\text{SA,MS,MA}}$. In our case, Γ_{surf} is the most relevant descriptor of intrinsic surface loss because the three surface loss factors are heavily interdependent. Since they are all heavily influenced by substrate preparation, metal deposition, and circuit patterning processes, they must be described collectively; therefore, Γ_{surf} can be used to accurately predict the total surface loss in a thin-film device. Importantly, however, Γ_{surf} now has an implicit geometric term $\frac{p_k}{p_{\text{surf}}}$ for $k = \text{SA, MS, MA}$. Fortunately, since the surface partici-

pation scales similarly with geometry for planar thin-film devices, the ratios $\frac{p_k}{p_{\text{surf}}}$ are either independent or weakly-dependent on resonator geometry. As a result, Γ_{surf} is effectively still a geometry-independent loss factor.

The second strategy to shrink the participation matrix is to define a total package loss $1/Q_{\text{pkg}} = p_{\text{pkg}_{\text{cond}}}\Gamma_{\text{pkg}_{\text{cond}}} + p_{\text{pkg}_{\text{MA}}}\Gamma_{\text{pkg}_{\text{MA}}} + p_{\text{seam}}\Gamma_{\text{seam}}$ as a combination of conductor loss due to current flow induced on the walls of the tunnel package, MA surface dielectric loss due to the oxide and other dielectric contaminants on walls' surfaces, and seam loss due to the end-caps that close the tunnel package (see Fig. 3.8). If we attempt to apply the same approach as we did for the surface dielectric losses, we would find that the package participations do not scale the same way with geometry; as a result, a unifying package loss factor is not a geometry-independent descriptor for package loss. To resolve this, we instead recognize that the FWGMR is a sensitive probe for these losses that are associated with bulk superconductors; since these tunnel packages are made with the same material and undergo the same processes (5N5 etched Al), we can use the loss factors extracted by the FWGMR to characterize and subtract the package losses from the TSLs and ASLs. We use the averaged loss factors from FWGMR devices (F1(e) and F2(e)) (see Lei et al. [112], Table I) to obtain $R_s = (0.61 \pm 0.28) \mu\Omega$ and $\tan \delta_{\text{pkg}_{\text{MA}}} = (4.1 \pm 1.8) \times 10^{-2}$. The FWGMR measurement shows that these loss factors are power-independent; we assume this to be the case in the TSL and ASL as well. We then calculate $\Gamma_{\text{pkg}_{\text{cond}}} = R_s/(\mu_0\omega\lambda)$ and $\Gamma_{\text{pkg}_{\text{MA}}} = \tan \delta_{\text{pkg}_{\text{MA}}}$, where $\lambda = 50 \text{ nm}$ and ω is the mode frequency. With this, we subtract $p_{\text{pkg}_{\text{cond}}}\Gamma_{\text{pkg}_{\text{cond}}} + p_{\text{pkg}_{\text{MA}}}\Gamma_{\text{pkg}_{\text{MA}}}$ from the internal quality factors of each mode. Doing so reveals that these losses are a small but not insignificant contributor to the total internal loss. For the TSL in particular, the package conductor and MA losses for each mode are $1/Q_{\text{D1}} = (1/17) \times 10^9$, $1/Q_{\text{D2}} = (1/4.7) \times 10^8$, and $1/Q_{\text{C}} = (1/1.3) \times 10^8$ for the D1, D2, and C modes, respectively. Since Q_{int} is expected to be at most $\approx 2.5 \times 10^6$ for the D1 mode, $\approx 40 \times 10^6$ for the D2 mode, and $\approx 20 \times 10^6$ for the C mode, these

sources of loss contribute negligibly in the D1 and D2 modes and no more than 10-15% of the total loss in the C mode. Importantly, we do not use the FWGMR to subtract seam loss, because the seam in the coaxial tunnel package is prepared differently from the seam in the FWMGR due to differences in applied clamping force. By utilizing the two aforementioned strategies, we can thus contract the participation matrix of the TSL to a 3x3 matrix, where the D1 mode is sensitive to surface loss, the D2 mode is sensitive to bulk dielectric loss, and the C mode is sensitive to seam loss (after subtracting out the package conductor and MA losses). Precise extraction of g_{seam} along with the loss factors from the FWGMR therefore allow us to fully subtract the package loss contribution of the ASL C mode, which allows it to be used to distinguish between surface and bulk dielectric loss.

Similar to the FWGMR, measurement sensitivity plots were calculated for TSLv1/3 and ASL and are shown in Fig. 4.8. TSLv2 sensitivity is very similar to TSLv1 and is not shown. Here, we distinguish between regions where the fractional error $\sigma_i/\Gamma_i < 1$ (blue) and region where the loss factors are not resolvable (red). We additionally fix the package losses for these plots using the FWGMR loss factors and seam loss $r_{\text{seam}} = 1/g_{\text{seam}} = 4.8 \times 10^{-3} \Omega$. From these plots, the TSL can resolve bulk loss factors as low as 2×10^{-9} and surface loss factors as low as 1×10^{-6} . The ASL has an order of magnitude reduced sensitivity due to the C mode being partially limited by package losses. As will be demonstrated further in this subsection, these sensitivities are still much higher than what is actually measured.

4.5.3 Device measurement & loss extraction

I will first demonstrate loss extraction in thin-films using tripole striplines made of tantalum. TSLs were fabricated on a HEMEX-grade sapphire substrate. The substrate was annealed at 1200 °C in oxygen before tantalum was deposited via DC magnetron sputter-

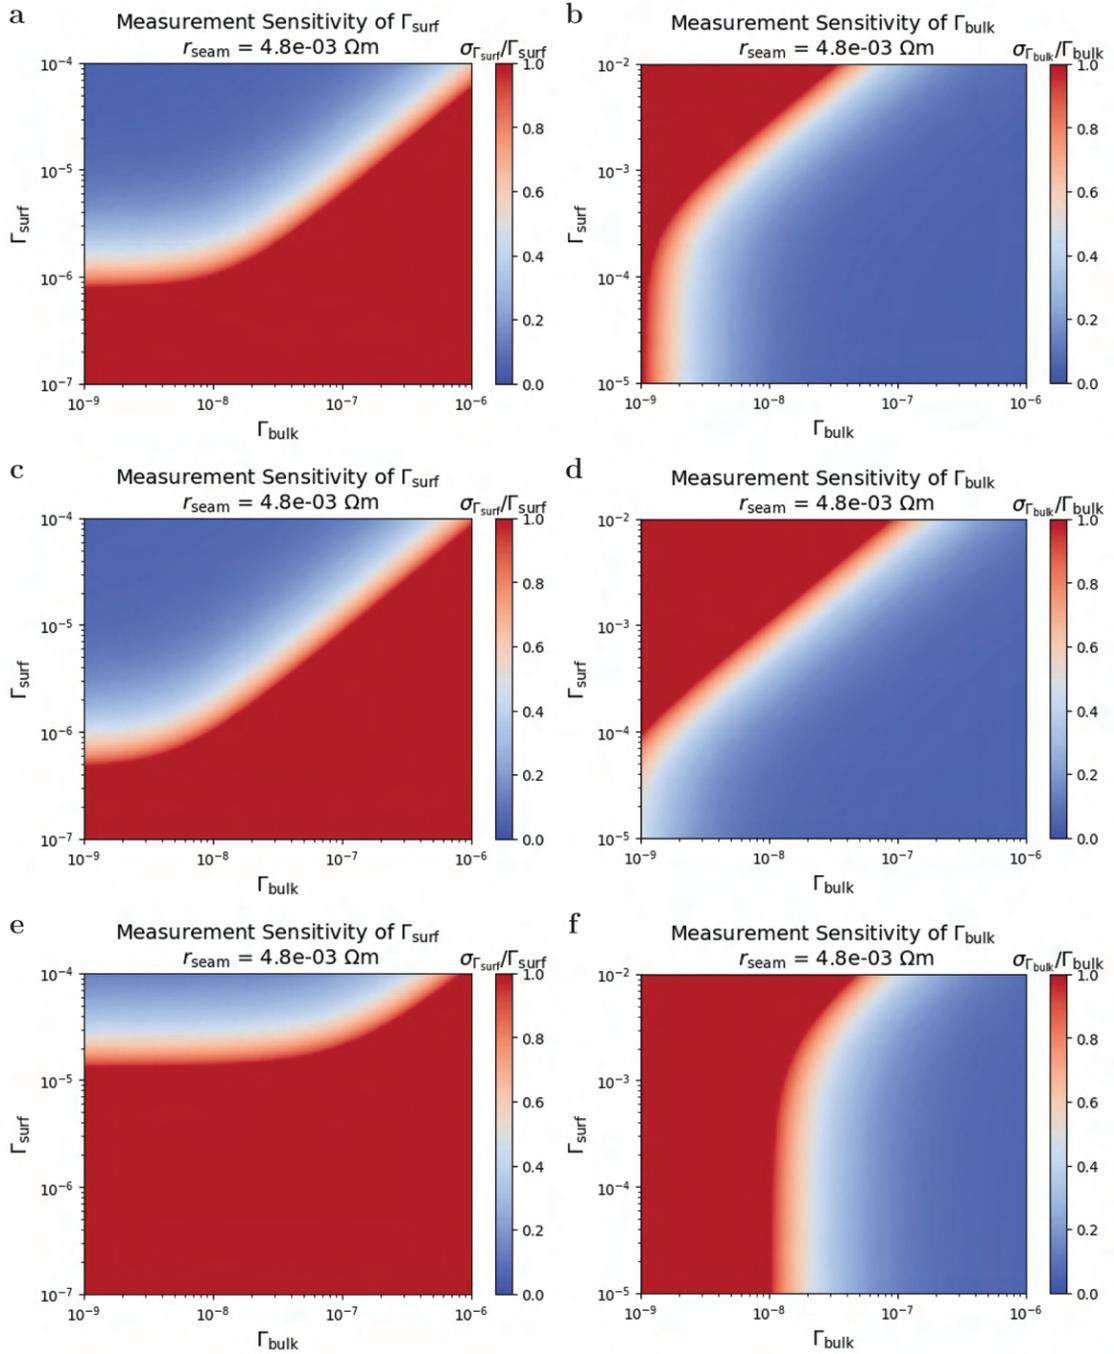


Figure 4.8: **TSL & ASL measurement sensitivity.** Surface and bulk loss measurement sensitivity for TSLv1 (a, b), TSLv3 (c, d), and ASLv1 (e, f). Here, $r_{\text{seam}} = 1/g_{\text{seam}}$. Figure obtained from Ganjam et al. [86]

ing at 800 °C. The striplines were subtractively patterned using optical lithography and flourine-based reactive ion etching (RIE), and loaded in multiplexed coaxial tunnel packages made of etched 5N5 aluminum. The full fabrication details are given in Appendix A. The package was mounted in a fridge setup that is described in Ch. 3.5.2. The resonance modes are measured in hanger configuration using a VNA. Q_{int} and \bar{n} were extracted using the methods in Ch. 3.5.3.

Fig. 4.9a shows the power dependence of the internal quality factors for the three modes. Here, we see that the D1 mode has the largest power dependence, with Q_{int} changing by over an order of magnitude from $\bar{n} = 10^0 - 10^6$. As discussed in Ch. 3.3.2, this can be attributed to TLSs that couple to the electric field of the mode. On the other hand, the D2 and C modes are much less power dependent, changing by less than a factor of two over the same range of \bar{n} . This behavior is consistent with the relative surface participations of the three modes; the D2 and C modes have nearly two orders of magnitude less surface participation than the D1 mode. This implies that the D2 and C modes have much weaker coupling to TLSs, which allow them to obtain single-photon quality factors that are over an order of magnitude larger than that of the D1 mode. The D2 mode, which is insensitive to both surface and package losses, can achieve single-photon internal quality factors of around 3×10^7 , far higher than the $Q_{\text{int}} \sim 10^6$ that is typically observed in planar stripline resonators.

To extract the intrinsic loss factors for this multimode device, we first fit the power dependence of each mode to the TLS model given by Eq. (3.24). Assuming a constant operating temperature of 20 mK, $\tanh\left(\frac{\hbar\omega}{2k_B T}\right) \approx 1$ for $\omega/2\pi > 3$ GHz, so the temperature-dependent term in the TLS model can be ignored. We use the TLS fit as an interpolating function to obtain a value for Q_{int} for all values of \bar{n} between 10^0 and 10^6 ; the error on these interpolated values is propagated from the measurement error and is therefore used to calculate the propagated error of the extracted loss factors. These errors were small, less

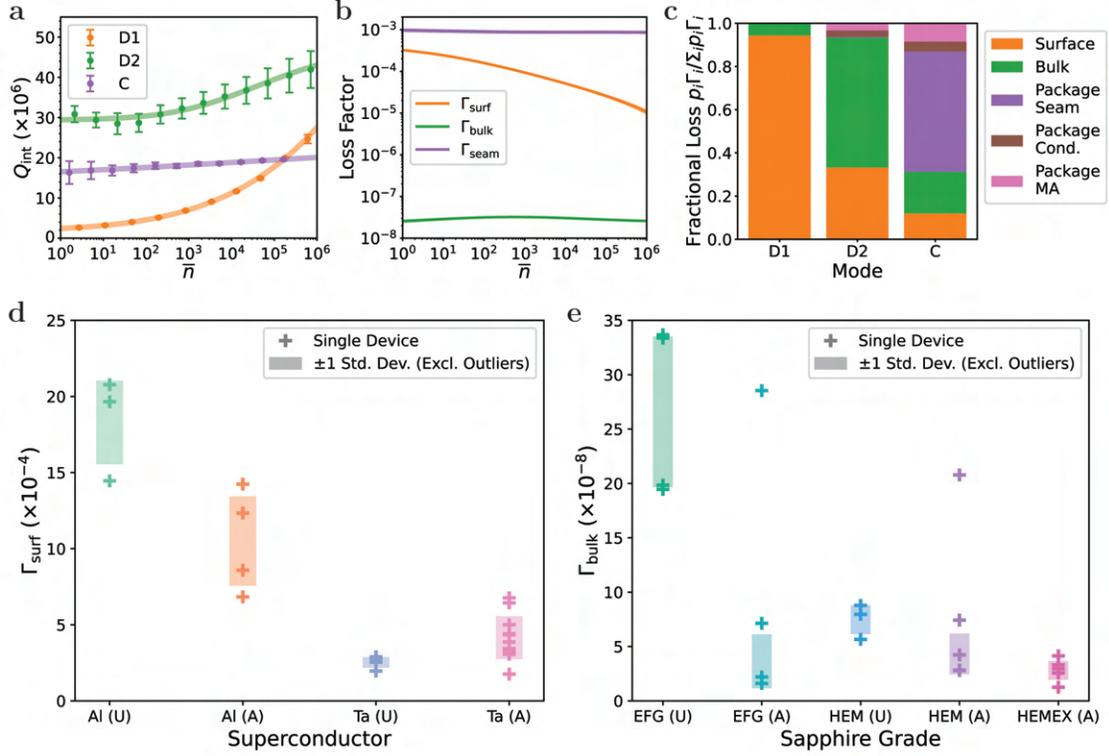


Figure 4.9: **TSL loss extraction.** **a** Power dependence of internal quality factor of the modes of a particular tripole stripline device, made using tantalum patterned on an annealed HEMEX sapphire substrate. Circles are measured Q_{int} ; lines are TLS fits using Eq. (3.24). Error bars represent the propagated fit error on Q_{int} obtained from least-squares minimization of Eq. (4.9) and for some points are small enough to not be visible. **b** Power dependence of extracted loss factors (solid lines). Propagated errors (shaded regions) are small ($\sim 10\%$) and are hidden within the width of the solid lines. Seam loss here has been normalized to be dimensionless, $\Gamma_{\text{seam}} = \omega\epsilon_0/g_{\text{seam}}$. **c** Single-photon loss budget for the modes of the tripole stripline. While the D1 mode is clearly dominated by surface loss, the D2 mode is dominated by bulk dielectric loss, and the C mode is dominated by seam loss. **d, e** Comparison of surface (**d**) and bulk (**e**) loss factors from multiple tripole stripline devices made using either aluminum- or tantalum-based fabrication processes on annealed (A) or unannealed (U) sapphire substrates. Figure obtained from Ganjam et al. [86]

than 5% in most cases. With these interpolated values of Q_{int} , we can subtract the package conductor and MA losses using the FWGMR measurements and subsequently extract the seam, bulk, and surface loss factors as a function of photon number.

The power dependence of the extracted loss factors allows us to observe general trends in the different sources of loss. The relatively much larger power dependence of the D1 mode coupled with its high surface participation implies that the surface loss factor is power-dependent while the other loss factors are not. This is confirmed in Fig. 4.9b, where the surface loss factor varies by over an order of magnitude, while seam loss and bulk dielectric loss do not. This also implies that the small amount of power dependence in the D2 and C modes stem from their small but nonzero surface participation. This observation directly confirms that the dominant source of TLSs that couple to superconducting circuits are localized in surface regions[115]. Additionally, the distinction between surface loss and bulk dielectric loss is made clear in the several orders of magnitude difference between their loss factors. At single-photon powers, the extracted bulk loss factor is $(2.6 \pm 0.2) \times 10^{-8}$, almost 4 orders of magnitude smaller than the surface loss factor of $(3.4 \pm 0.3) \times 10^{-4}$. This is qualitatively similar to what is observed in other studies[106, 110, 136], and highlights the difference in processing seen by the two regions as well as the relative degree to which the surface may be contaminant or defect-rich.

With the extracted loss factors, we can again quantify how much each source of loss contributes to the total loss of each mode. Similar to Fig. 4.5b, we can plot the fractional loss contribution $p_i \Gamma_i / \sum_i p_i \Gamma_i$ of each source of loss for each mode in Fig. 4.9c. This shows that the D1 mode is dominated by surface loss, the D2 mode is dominated by bulk dielectric loss, and the C mode is dominated by seam loss. This fulfills the conditions for an ideal loss characterization platform discussed in Sec. 4.2: each mode is sensitive to a different source of loss.

4.5.4 Comparison of loss factors of different materials & processes

The TSL device measured is one made using a particular set of materials and fabrication processes. We can quantify the device-to-device variation of loss factors by measuring several nominally devices made with a nominally identical process. This allows us to capture any inhomogeneity or lack of consistency in the extracted loss factors without relying on a statistical measurement uncertainty. Additionally, we can measure how using different materials such as the grade of sapphire wafer, the wafer annealing treatment, or the superconductor affects the measured surface and bulk loss factors (since we keep the package the same, we expect the seam loss to be uncorrelated with the choice on on-chip materials or processes). The results of this are shown in Figs. 4.9d,e. Here, we see that for most materials and process combinations, the data is well-clustered; however, some outliers do exist; however, the outliers are not representative of the expected loss factor an average device would be subjected to. As a result, averages and deviations for the loss factors were calculated by excluding any outlier with median relative deviation $|X_i - \tilde{X}|/\tilde{X} > 3$, where \tilde{X} is the median.

Fig. 4.9d clearly shows that surface loss factors can vary greatly based on the type of substrate preparation, superconductor, and lithography process. Aluminum-based fabrication processes (show evaporation through a resist mask followed by liftoff; see Appendix A) on unannealed sapphire substrates yield the largest surface loss factors. Annealing the substrate can improve this by a factor two; however, this improvement seems to be specific to aluminum. For Tantalum-based process, there seems to be little difference in surface quality if the substrate is annealed or not. Regardless, we see that the tantalum-based process yields surface loss factors that are over two times lower than those of an aluminum-based process. To gain additional understanding as to why this is the case, we can look to physical materials characterization to search for physical signatures of loss

between Al and Ta (see Appendices B.2 and B.3). Cross-sectional transmission electron microscopy (TEM) reveals significant differences between the MA and MS interfaces of Al and Ta on sapphire (Figs. B.2 and B.3). For one, the MA interface of aluminum has a ≈ 5 nm-thick oxide layer (Fig. B.2a), while the MA interface of tantalum only has a ≈ 3 nm-thick oxide (Fig. B.3a). This rescales the aluminum MA loss tangent $\tan \delta_{MA}^{Al}$, thereby affecting Γ_{surf}^{Al} . In other words, the difference in surface loss factors may be due to the increased oxide thickness. TEM also reveals significant differences in the MS interfaces; here, the aluminum-sapphire interface has a thin, ≈ 2 nm-thick amorphous region between the metal and the substrate (Fig. B.2b), which may be amorphous AlO_x or Al, while the tantalum-sapphire interface is incredibly clean, with nearly epitaxial film growth and no observable sign of an amorphous region (Fig. B.3b). It should be emphasized here that this interfacial behavior may not be intrinsic to the material itself; what we see is a convolution of the materials used and the fabrication processes employed. Aluminum-based devices were deposited using electron-beam evaporation and patterned using a liftoff process while the tantalum-based devices were deposited via high-temperature sputtering and patterned subtractively. These fabrication processes can significantly influence how the films grow and interact with the substrate and the air, as well as the level of contamination introduced to the device[139, 140].

From Fig. 4.9e we can see that bulk dielectric loss factors can also vary based on the choice of sapphire grade and annealing treatment. Here, it seems that EFG sapphire has the highest bulk dielectric loss, over three times higher than HEM. by annealing either, however, the loss factor reduces by a factor of approximately eight for EFG and two for HEM. Additionally, annealing HEMEX-grade sapphire yields the lowest bulk loss factor with the smallest device-to-device variation as measured over six devices. Importantly, annealing the sapphire regardless of grade reduces the bulk loss factor to approximately the same level. This improvement might originate the degree of crystalline order or defect density

present in the various sapphire grades. The EFG sapphire growth method may produce a higher defect or dislocation density, whereas the HEM growth method, which is specially optimized for optical uniformity and graded based on it (HEMEX being the highest grade), may have greater crystalline order and uniformity[106, 141]. Annealing the wafer may repair such defects or disorder, resulting in reduced loss. While bulk crystalline order was not probed in this work, surface-level morphological changes were observed after annealing through atomic force microscopy (AFM) of the sapphire surface (see Appendix B.1). After annealing, the sapphire surface was atomically flat with monatomic step-heights that correspond to the wafer miscut angle. It should be noted that the differences in EFG and HEM loss factor are in qualitative agreement with other studies[103, 106]; however, the absolute loss tangents vary significantly. We attribute these effects as being due to the fabrication process of the device, which may have introduced additional damage to the crystal structure of the substrate. In particular, the devices used in this work were patterned on chips that were diced with a diamond-embedded saw, which is a violent process that causes chipping of the sapphire at the edges of the chip that may have caused deeper subsurface damage. This highlights the fact that microwave loss characterization is principally about characterizing the convolution of materials used and fabrication processing employed.

The device-by-device loss factors at single-photon powers are shown in Table 4.5. The devices were first grouped according to superconductor used and annealing process employed and the surface loss factors were averaged for each group while excluding outliers with median relative deviation $MRD > 3$. The devices were then grouped according to the sapphire grade used and annealing process employed to determine an average bulk loss factor (again, excluding outliers). These loss factors are given in Table 4.7, and capture the intrinsic variation we expect to see for our various fabrication processes. As a result, they are the useful loss factors that will be used to predict the losses of transmons and on-chip quantum memories in Ch. 5. Loss factors at high powers are also shown in Ta-

Table 4.5: TSL & ASL loss factors at $\bar{n} = 1$. ¹ ASL differential (D) mode Q. Table obtained from Ganjam et al. [86]

Device ID	Sapphire Growth Method	Anneal	Thin-Film Superconductor	Mode $Q_{\text{int}}(\bar{n} = 1)$ ($\times 10^6$)			Loss Factors ($\bar{n} = 1$)		
				D1	D2	C	Γ_{surf} ($\times 10^{-4}$)	Γ_{bulk} ($\times 10^{-8}$)	$1/g_{\text{seam}}$ (Ωm) ($\times 10^{-3}$)
AM22 TSL1				0.55	8.95	9.00	14.5 ± 1.0	7.96 ± 0.5	13.4 ± 0.9
AM22 TSL2	HEM	No	Al	0.41	8.79	4.30	19.7 ± 3.2	5.65 ± 1.5	51.6 ± 2.4
AM22 TSL3				0.39	7.13	6.41	20.8 ± 6.3	8.78 ± 3.0	22.4 ± 1.2
DZ22 TSL3	HEM	Yes	Al	0.56	9.25	3.13	14.2 ± 3.5	7.42 ± 1.7	81.2 ± 1.2
DZ22 TSL4				0.82	5.41	6.81	8.58 ± 0.5	20.8 ± 0.4	13.1 ± 1.1
A23Al TSL2	HEMEX	Yes	Al	1.53	10.33	12.21	6.83 ± 0.4	4.14 ± 1.8	1.95 ± 1.6
A23Al TSL3				0.88	7.56	10.17	12.3 ± 0.5	2.57 ± 0.7	1.56 ± 0.3
EF21 TSL1				2.41	6.24	10.21	1.96 ± 0.2	19.8 ± 0.7	0.87 ± 0.7
EF21 TSL2				1.91	6.22	9.93	2.89 ± 0.1	19.5 ± 1.2	1.26 ± 1.2
EF21 TSL3	EFG	No	Ta	1.58	3.66	0.81	2.72 ± 0.2	33.7 ± 0.2	186 ± 1.0
EF21 TSL4				1.63	3.65	0.48	2.54 ± 0.3	33.4 ± 2.0	331 ± 1.5
EC21 ASL1				2.67 ¹	-	5.04	1.96 ± 0.2	19.8 ± 0.7	-
EC21 ASL2				3.41 ¹	-	10.99	2.89 ± 0.1	19.5 ± 1.2	-
EC21 ASL3	EFG	Yes	Ta	3.19 ¹	-	10.16	2.72 ± 0.2	33.7 ± 0.2	-
EC21 ASL4				2.14 ¹	-	15.38	2.54 ± 0.3	33.4 ± 2.0	-
R22 TSL1				1.59	24.04	24.20	5.00 ± 1.0	2.78 ± 0.5	1.08 ± 0.4
R22 TSL3	HEM	Yes	Ta	2.04	26.01	7.06	3.87 ± 0.2	2.83 ± 0.2	11.9 ± 1.4
R22 TSL4				1.78	19.66	13.69	4.35 ± 0.6	4.23 ± 0.3	3.92 ± 0.3
BF22 TSL1				1.24	19.57	14.25	6.44 ± 0.7	3.30 ± 0.5	2.38 ± 0.2
BF22 TSL2				2.53	27.40	14.96	3.07 ± 0.1	2.95 ± 0.1	2.81 ± 0.2
BF22 TSL3	HEMEX	Yes	Ta	2.38	41.23	10.80	3.40 ± 0.3	1.25 ± 0.2	5.18 ± 0.2
BF22 TSL4				2.44	29.47	16.54	3.22 ± 0.1	2.56 ± 0.2	2.45 ± 0.3

Table 4.6: TSL & ASL loss factors at high power. Loss factors for some devices are not resolvable due to the dominance of bulk or package loss at high powers. ¹ ASL differential (D) mode Q.

Device ID	Sapphire Growth Method	Anneal	Thin-Film Superconductor	Measured Photon Number	Mode Q_{int} ($\times 10^6$)				Loss Factors				$1/g_{\text{scam}}$ (Ωm) ($\times 10^{-3}$)
					D1	D2	C	Γ_{surf} ($\times 10^{-4}$)	Γ_{bulk} ($\times 10^{-8}$)	Γ_{surf} ($\times 10^{-4}$)	Γ_{bulk} ($\times 10^{-8}$)		
AM22 TSL1				10^6	4.91	18.42	15.08	1.20 ± 0.1	6.51 ± 0.1	1.20 ± 0.1	6.51 ± 0.1	9.99 ± 0.8	
AM22 TSL2	HEM	No	Al	10^7	2.96	16.98	5.95	2.0 ± 0.6	5.73 ± 0.3	2.0 ± 0.6	5.73 ± 0.3	42.2 ± 0.9	
AM22 TSL3				10^7	1.74	11.87	8.22	3.8 ± 0.2	8.37 ± 0.2	3.8 ± 0.2	8.37 ± 0.2	22.5 ± 0.8	
DZ22 TSL3	HEM	Yes	Al	10^7	12.96	16.34	3.73	-	-	-	-	-	
DZ22 TSL4				10^6	33.43	6.95	8.62	-	-	-	-	-	
A23Al TSL2				10^5	63.01	87.49	59.13	0.11 ± 0.1	1.60 ± 0.6	0.11 ± 0.1	1.60 ± 0.6	0.66 ± 0.4	
A23Al TSL3	HEMEX	Yes	Al	10^5	45.05	57.56	27.49	0.14 ± 0.1	2.20 ± 0.2	0.14 ± 0.1	2.20 ± 0.2	2.18 ± 0.2	
EF21 TSL1				10^6	23.10	7.70	16.83	-	-	-	-	-	
EF21 TSL2				10^6	17.18	10.07	14.31	-	-	-	-	-	
EF21 TSL3	EFG	No	Ta	10^5	16.13	4.56	0.83	-	-	-	-	-	
EF21 TSL4				10^5	15.16	5.57	0.48	-	-	-	-	-	
EC21 ASL1				10^5	16.88 ¹	-	5.45	-	-	-	-	-	
EC21 ASL2				10^5	12.56 ¹	-	12.57	0.42 ± 0.1	6.81 ± 0.7	0.42 ± 0.1	6.81 ± 0.7	-	
EC21 ASL3	EFG	Yes	Ta	10^5	13.97 ¹	-	12.36	1.18 ± 0.2	0.83 ± 0.6	1.18 ± 0.2	0.83 ± 0.6	-	
EC21 ASL4				10^5	11.14 ¹	-	21.01	1.20 ± 0.1	1.65 ± 0.8	1.20 ± 0.1	1.65 ± 0.8	-	
R22 TSL1				10^5	13.78	62.55	36.82	0.52 ± 0.1	1.82 ± 0.1	0.52 ± 0.1	1.82 ± 0.1	1.27 ± 0.3	
R22 TSL3				10^5	16.31	48.38	7.64	0.38 ± 0.02	2.39 ± 0.1	0.38 ± 0.02	2.39 ± 0.1	12.0 ± 1.1	
R22 TSL4	HEM	Yes	Ta	10^5	18.32	41.02	15.62	0.22 ± 0.1	3.04 ± 0.1	0.22 ± 0.1	3.04 ± 0.1	4.66 ± 0.3	
BF22 TSL1				10^5	2.36	27.48	17.33	3.47 ± 0.1	3.01 ± 0.1	3.47 ± 0.1	3.01 ± 0.1	2.17 ± 0.2	
BF22 TSL2				10^7	22.95	45.21	18.54	0.08 ± 0.02	2.25 ± 0.1	0.08 ± 0.02	2.25 ± 0.1	2.49 ± 0.2	
BF22 TSL3	HEMEX	Yes	Ta	10^6	17.27	66.03	11.81	0.37 ± 0.1	1.44 ± 0.1	0.37 ± 0.1	1.44 ± 0.1	5.04 ± 0.2	
BF22 TSL4				10^6	25.54	43.05	20.09	0.11 ± 0.02	2.59 ± 0.1	0.11 ± 0.02	2.59 ± 0.1	2.17 ± 0.2	

Table 4.7: Average single-photon surface & bulk loss factors for thin-film device, calculated by excluding outliers with median relative deviation $\text{MRD} > 3$. Table obtained from Ganjam et al. [86]

Material/Process System	$\Gamma_{\text{surf}} (\times 10^{-4})$
Al (Unannealed)	18.3 ± 2.7
Al (Annealed)	10.5 ± 2.9
Ta (Unannealed)	2.53 ± 0.4
Ta (Annealed)	4.15 ± 1.4

Material/Process System	$\Gamma_{\text{bulk}} (\times 10^{-8})$
EFG (Unannealed)	26.6 ± 6.9
EFG (Annealed)	3.64 ± 2.5
HEM (Unannealed)	7.46 ± 1.3
HEM (Annealed)	4.31 ± 1.9
HEMEX (Annealed)	2.80 ± 0.9

ble 4.6. Here, the saturation of TLS loss results in much lower surface loss factors. For some devices, the bulk and seam loss factors dominate the loss of even the D1 mode; the TSL loses sensitivity to surface loss, and therefore is unable to resolve the loss factors at high powers. In such cases, we could opt to perform the Monte Carlo method in order to estimate the upper bounds on the loss factors; this was not done for this work.

Finally, we observe a large amount of variation in the extracted seam loss for these devices, as can be seen in Table 4.5 ($1/g_{\text{seam}}$). We attribute this to device-to-device variation in interface quality due to residual contamination, surface roughness, and clamping force. This highlights the significance of package losses in the coaxial architecture and in the 3D cavity architecture as a potential source of large device-to-device variation. In 3D cavity architectures, this was circumvented by cleverly designing cavities that are insensitive to seam loss, as is the case for the post cavity. For the coaxial architecture, this can also be done, and can be seen in the high-Q D2 modes of the tripole stripline. We will see in Ch. 5 that using this strategy is what enables us to realize highly coherent on-chip quantum memories. Nevertheless, for loss characterization and predictive purposes we can calculate

an average $g_{\text{seam}} = (2.1 \pm 2.0) \times 10^2 (\Omega\text{m})^{-1}$ by excluding outliers with $\text{MRD} > 3$. The large variation in the average g_{seam} reflects the variation one should expect when preparing the seam in this way. However, better methods of seam preparation may exist that could enable far higher and more consistent seam qualities, such as surface polishing, thin-film coating, which were both explored with the FWGMR in Sec. 4.4. Additionally, a greater clamping force may also improve the seam.

4.5.5 Conductor loss in thin-films

So far, I have neglected the inclusion of conductor loss. While conductor loss has been characterized in bulk superconductors using the FWGMR, the tripole stripline is not capable on uniquely extracting the residual resistance of thin-film superconductors. However, this does not break the loss model, as conductor participation scales in the same way as $p_{\text{SA,MS,MA}}$. In fact, we can simply modify our definition of Γ_{surf} to include conductor loss:

$$\Gamma_{\text{surf}} = \frac{p_{\text{SA}}}{p_{\text{surf}}} \tan \delta_{\text{SA}} + \frac{p_{\text{MS}}}{p_{\text{surf}}} \tan \delta_{\text{MS}} + \frac{p_{\text{MA}}}{p_{\text{surf}}} \tan \delta_{\text{MA}} + \frac{p_{\text{cond}}}{p_{\text{surf}}} \Gamma_{\text{cond}} \quad (4.16)$$

where $\Gamma_{\text{cond}} = R_s/\mu_0\omega\lambda$. For this work, the quantity $\frac{p_{\text{cond}}}{p_{\text{surf}}} \Gamma_{\text{cond}}$ is additionally assumed to be small, and contributes negligibly to loss in the single-photon regime. The conductor participations for the modes of the ASL and TSL are given in Table 4.8. It is important to note here that conductor participation is much larger in thin-film resonators than in 3D cavities. Here, the D1 mode has conductor participation of order $\sim 10^{-2}$; the DWGM mode of the FWGMR, which is especially sensitive to conductor loss, only has conductor participation of order $\sim 10^{-4}$. This is due to the high current densities found along the edges of the thin film; surface dielectric participations are high in thin-films for the same reason.

Table 4.8: TSL & ASL conductor participation, calculated using the methods in Appendix C. A $\lambda = 50$ nm was assumed for these calculations.

Mode	p_{cond}			
	TSLv1	TSLv2	TSLv3	ASLv1
D1 (D)	6.0×10^{-2}	6.1×10^{-2}	4.2×10^{-2}	2.6×10^{-2}
D2	8.3×10^{-4}	8.0×10^{-4}	2.7×10^{-3}	-
C	5.5×10^{-4}	5.5×10^{-4}	1.9×10^{-3}	1.3×10^{-3}

Unlike bulk superconductors such as etched 5N5 Al which seems to have $x_{qp} \sim 10^{-4}$, thin-film aluminum seems to have much lower residual quasiparticle fraction. Quasiparticle tunneling experiments using offset charge-sensitive transmons have shown that as long as the proper radiation shielding and RF filtering are used, x_{qp} can be as low as $x_{qp} = 10^{-8} - 10^{-10}$ [97, 98, 134]. We can therefore use Eq. (3.18) to calculate the expected R_s for Al and Ta under the conservative assumption that $x_{qp} = 10^{-8}$ in both superconductors.

For Ta, we use the DC transport measurements described in Appendix B.3 that give $T_c = 4.3$ K and $\rho_n = 2.55 \times 10^{-9}$ Ωm , calculated by dividing the room-temperature resistivity by the residual resistivity ratio (RRR). Using values for the Fermi velocity and effective electron mass from Halloran et al. [142] and Bobrov et al. [143], we can calculate $\lambda_0 = 30.7$ nm and $\xi_0 = 96.4$ nm. From Eqs. (4.13), (4.14), and (4.15), we calculate $l = 145$ nm, $\lambda = 41.1$ nm, and $\xi = 57.8$ nm. The thickness of the deposited film is nominally 150 nm; this implies that the mean free path is limited by the film's thickness. Remarkably, the long mean free path puts this tantalum film in the clean limit of superconductivity, i.e. $l > \xi_0$. In this limit, scattering impurities are negligible and quasiparticles travel in ballistic trajectories within the superconductor. However, Eq. (3.18) implies that $R_s \propto \lambda^3/\rho_n \propto 1/l^{1/2}$ when $l \ll \xi_0$; but as $l \rightarrow \infty$, λ becomes independent of l and $R_s \propto l$, resulting in an increase of surface resistance as the superconductor becomes cleaner. This is inconsistent with numerical simulations of R_s [93] and is resolved by using an effective

resistivity $\rho_{eff} \approx \rho_n l / \lambda$ for the calculation of R_s where l is large[93]. This gives $R_s = 1.48 \times 10^{-11} \Omega$, which is four orders of magnitude smaller than what was measured in etched 5N5 Al, and corresponds to $\Gamma_{cond} \approx 10^{-8}$ at 5 GHz. Therefore, conductor loss in tantalum thin films should limit the D1 mode of the TSL to $Q_{cond} > 10^9$.

A similar calculation can be done for Al; unfortunately, DC transport measurement were not done for the films grown in this work. Moreover, the Al process mimics the Al-based transmon fabrication process in which a trilayer of 20 nmAl/2 nm AlO_x/30 nmAl is grown. However, if we assume that the mean free path of the aluminum film is similarly limited by its thickness, we can assume $l \approx 25$ nm (average thickness of the two Al layers). From this, we can use Eq. (4.13) to get $\rho_n = 1.7 \times 10^{-8} \Omega\text{m}$ (RRR ≈ 1.5). For thin Al films, the T_c is higher than that of bulk Al; for 20 and 30 nm-thick films, we assume an average $T_c = 1.35$ K (measured in Kyle Serniak’s thesis[35]), which gives $\xi_0 \approx 1.4 \mu\text{m}$ [33, 95]. For this Al film, we are well in the dirty limit, $l \ll \xi_0$, and Eq. (3.18) is applicable. With these parameters and $\lambda_0 = 16$ nm[33], we calculate $\xi = 25$ nm and $\lambda = 121$ nm. Finally, we calculate $R_s = 3.48 \times 10^{-11} \Omega$ due to excess quasiparticles, which is quite similar to that of the Ta thin-film and likewise contributes negligibly to the total loss.

To estimate vortex flow resistance, we begin with the assumption that the cooling field $B_0 = 10^{-7}$ T, which yields $R_{vortices} \approx 24$ n Ω for aluminum and $R_{vortices} \approx 7$ n Ω for tantalum, or $\Gamma_{cond} \approx 5 \times 10^{-6}$ at 5 GHz for both. This would limit the D1 mode of the TSL to $\approx 3.7 \times 10^6$. However, this is inconsistent with high-power quality factor measurements of the D1 mode, which can be as high as 6×10^7 as in the case of device “A23Al TSL2”. A better estimation of R_s is provided by the high-power surface loss factor. Since $\Gamma_{cond} < p_{surf} \Gamma_{surf} / p_{cond}$, we can assume that conductor loss is power-independent and use the high-power surface loss factor to set an upper bound on Γ_{cond} . The lowest high-power surface loss factor for Ta is $\Gamma_{surf} \approx 8 \times 10^{-6}$ (from device “BF22 TSL2”; see Table 4.6), which

implies that $\Gamma_{\text{cond}} < 1.6 \times 10^{-7}$ or $R_s < 0.26 \text{ n}\Omega$ assuming $\omega/2\pi = 5 \text{ GHz}$. For Al, the lowest high-power surface loss factor is $\Gamma_{\text{surf}} \approx 1 \times 10^{-5}$ (from device “A23Al TSL2”; see Table 4.6), which implies that $\Gamma_{\text{cond}} < 2.4 \times 10^{-7}$ or $R_s < 1.2 \text{ n}\Omega$ assuming $\omega/2\pi = 5 \text{ GHz}$. These loss factors would limit the quality factor the D1 mode to $Q_{\text{cond}} > 10^8$; as a result, conductor loss contributes negligibly to the internal loss at single-photon powers. Interestingly, this bound on the residual resistance allows us to bound the ambient cooling field in the device’s environment. Using Eq. (3.19), we can set an upper bound for the cooling field to be $B_0 < 4 \times 10^{-9} \text{ T}$, which implies that the ambient field attenuates by almost 2 orders of magnitude as the Cryoperm shield cools down, either due to its changing magnetic permeability or the magnetic field expulsion from the lead foil that is wrapped around the copper can that houses the device (see Ch. 3.5.2).

The loss characterization of thin films and bulk superconductors leads to some very surprising conclusions that were long suspected but not directly verified. For one, it appears as if intrinsic loss factors associated with thin-film devices are far lower than those of bulk superconducting Al. With regards to conductor losses, we are able to place bounds on the residual resistance of thin-film Al and Ta that are several orders of magnitude lower than that of etched 5N5 bulk Al. Additionally, bulk Al seem to be limited by a very high excess quasiparticle population, over four orders of magnitude higher than that of thin films (even with conservative estimates). It should be noted, however, that the FWGMR was measured in a different fridge which did not have the same level of radiation shielding and RF filtering as the setup described in Ch. 3.5.2, which may have contributed to an increased x_{qp} . However, anecdotal evidence from my time at Yale indicates that even when more shielding filtering is added, there has not been a noticeable improvement in cavity quality factors. This may be due to the fact that Al-based cavity resonators tend to already have large device-to-device variation in quality factor, mostly due to inconsistent machining and etching processes. Tremendous device-to-device variation has been observed in

the coaxial architecture as well[84]; With the tripole striplines, we can pinpoint the origin of this variation once again to seam losses associated with the bulk superconductors. In contrast, aside from a few outliers, “on-chip” (i.e. surface and bulk) losses seem fairly consistent, and as we will see in Ch. 5, the consistency is what will allow us to make accurate predictions of device performance.

Chapter 5

Understanding and Optimizing Losses in Planar Transmon Qubits and Quantum Memories

With the tripole stripline, we gain a comprehensive understanding of how materials and fabrication processes limit the coherence of superconducting quantum circuits. In the previous chapter I have shown how using high-purity aluminum with a chemical etching process yields dramatic improvements to surface resistance and MA dielectric loss when compared with 6061 or unetched 5N5 Al. Additionally, by using tantalum-based fabrication and sapphire annealing processes for thin-film circuits, we can dramatically improve surface and bulk loss factors, and make stripline resonators with quality factors as high as $3 - 4 \times 10^7$ in the single-photon regime. In this chapter, I will show how we can use the knowledge gained from loss characterization to predictively model the relaxation times of aluminum- and tantalum-based transmon qubits. I will then demonstrate the accuracy of these predictions by comparing them to measured transmon T_1 . Finally, I will describe how we can utilize our newfound understanding of intrinsic loss to optimize materials, processes, and device geometry to design a stripline-based quantum memory with coherence times that exceed one millisecond.

5.1 Validating the Loss Model with Qubit Measurements

In this section, I will use the loss factors extracted by the FWGMR (Ch. 4.4) and the tripole striplines (Ch. 4.5) to predict the internal quality factors of a transmon qubit. These predictions will then be verified by comparing them with measured transmon coherence. In particular, I will describe how the standard aluminum-based transmons differ from tantalum-based transmons fabricated on various grades of sapphire.

5.1.1 Transmon Device Design

A standard 3D transmon device[73] was used for this work. This design is shown in Fig. 5.1. These devices were co-fabricated with the tripole striplines to ensure that the validation devices (transmon) were subjected to the same processing as the loss-characterization devices (TSL). Each transmon device is patterned on a chip that is inserted into the same cylindrical tunnel packages that were used to measure the TSLs, with the same drive line used for readout and qubit driving. On the chip we pattern a readout resonator that is coupled to the transmon to enable state measurement and a Purcell filter that is weakly coupled to the readout resonator to obtain a large $\kappa_r \approx 1$ MHz while preserving a low external coupling of the drive line to the transmon. The transmon and readout parameters are given in Table 5.1.

Table 5.1: Typical transmon (subscript “t”) & readout (subscript “r”) parameters. Table obtained from Ganjam et al. [86]

$\omega_t/2\pi$ (GHz)	4.5-5.1
$\omega_r/2\pi$ (GHz)	9.0-9.3
$\chi_{tt}/2\pi$ (MHz)	170-180
$\chi_{tr}/2\pi$ (MHz)	0.5-1.1
$\kappa_r/2\pi$ (MHz)	0.5-1.1

The Al and Ta-based transmons have very similar geometries but they differ in small

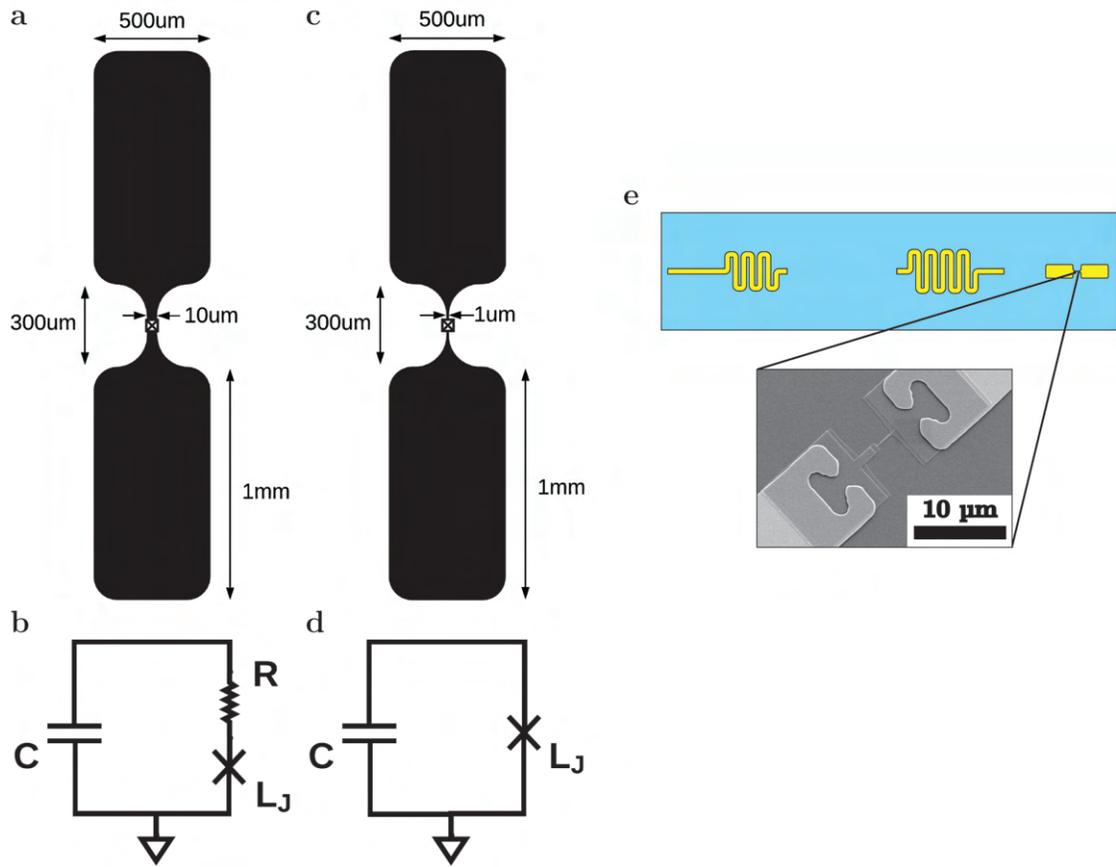


Figure 5.1: **3D Transmon qubit device design.** **a, b** Ta-based transmon and circuit model. The resistor is used to represent loss from the Ta/Al contact region. **c, d** Al-based transmon and circuit model. Junction leads are thinner to improve reliability and reproducibility of the electron-beam lithography process. **e** Transmon on chip with readout resonator and Purcell filter. Inset: SEM of Josephson junction and near-junction region on a tantalum-based transmon. The overlap region between the aluminum and tantalum is a source of contact loss. Figure adapted from Ganjam et al. [86]

ways. Both transmons have the same capacitor design. The Ta-based transmon has wider leads ($10\ \mu\text{m}$) to the junction, as can be seen in Fig. 5.1a. The wider leads have the effect of distributing the local electric field over a larger area, thereby slightly reducing participation. The Al-based transmon, on the other hand, has much skinnier leads at $1\ \mu\text{m}$, as seen in fig. 5.1c, d. Since the Al-based transmon is patterned entirely with electron-beam lithography, proximity effects must be considered. Patterning large features (such as a $10\ \mu\text{m}$ -wide lead) close to the submicron-sized junction results in significant proximity dosing which severely impacted device yield. By making the lead skinnier, the problem was averted. The skinnier lead has the side-effect of slightly increasing surface participation.

In a tantalum-based transmon, the capacitor pads are fabricated using subtractively-patterned tantalum, and the Dolan-bridge style Al/AlO_x/Al Josephson junctions are patterned additively using double-angle shadow evaporation through a resist mask defined with electron-beam lithography followed by liftoff (see Appendix A for fabrication details). Notably, this is a two-step fabrication process; since the tantalum is deposited on a clean substrate and is patterned subtractively, we would expect a much cleaner MS interface. The Al-based transmon is patterned in a single step where electron-beam lithography is used to pattern the junction and the capacitor pads into the resist mask. Then Al/AlO_x/Al is deposited using double-angle shadow evaporation, followed by liftoff. This has the advantage of only requiring one lithography step to fabricate the transmon; the drawback is a potentially more contaminated MS interface. The effects of the two fabrication processes can be seen in cross-sectional TEM (see Figs. B.2b and B.3b).

5.1.2 Transmon Participations

To describe the losses in transmon qubits, we can use the participation ratio model described in Ch. 3.1. Since these devices are measured in the same coaxial tunnel package

and are made with the same materials and fabrication processes as the tripole striplines, they are limited by the same sources of loss. In particular, Al-based transmons are limited by surface losses associated with the aluminum fabrication process, bulk dielectric loss, and package loss. Ta-based transmons, due to their hybrid nature of having tantalum pads and aluminum junctions, are limited by surface losses associated with both the tantalum and aluminum processes, as well as bulk dielectric loss and package loss. Additionally, since the Ta-based transmon uses two different metals deposited in different deposition steps, the contact between the Al and Ta may be resistive and therefore lossy. We can model this with a resistor in series with the junction, as in Fig. 5.1b. The source of this resistance may be contaminants or surface adsorbates in the region as well as native tantalum oxide. Two contact regions between the Ta and Al exist and are shown in Fig. 5.1e. Since this loss is due to the contact between two metals, it is convenient to describe it using a seam loss model[122]:

$$\frac{1}{Q_{\text{seamTa/Al}}} = \frac{2R_{\text{Ta/Al}}}{Z_0} = \frac{y_{\text{seamTa/Al}}}{g_{\text{seamTa/Al}}} \quad (5.1)$$

where $R_{\text{Ta/Al}}$ is the resistance of a single Ta/Al contact region, $Z_0/\sqrt{L_J/C_q}$ is the characteristic impedance of the transmon (transmons as designed had $L_J = 9$ nH and $C_q = 122$ fF), $y_{\text{seamTa/Al}} = 2/(Z_0 w)$, $g_{\text{seamTa/Al}} = 1/(R_{\text{Ta/Al}} w)$, and $w = 10$ μm is the width of the seam. To alleviate the Ta/Al contact loss, we employ an argon ion beam cleaning step prior to aluminum deposition. This process is intended to remove the tantalum oxide and any surface contamination such that after aluminum deposition, a good metal-to-metal contact forms. In order to fully characterize the loss of a Ta-based transmon, this ion beam cleaning process must be characterized and the contact resistance must be measured at microwave frequencies and single-photon powers. This will be discussed further in Sec. 5.1.3.

Table 5.2: Transmon participations, calculated using the methods in Appendix C. y_{seam} for both package and contact loss are given in units of $(\Omega m)^{-1}$. Table obtained from Ganjam et al. [86]

	Al transmon	Ta transmon
p_{surfTa}	-	8.1×10^{-5}
p_{surfAl}	1.5×10^{-4}	5.5×10^{-5}
p_{bulk}	0.84	0.84
$y_{\text{seamTa/Al}}$	-	7.4×10^2
$p_{\text{pkg}_{\text{cond}}}$	9.3×10^{-8}	9.3×10^{-8}
$p_{\text{pkg}_{\text{MA}}}$	5.1×10^{-9}	5.1×10^{-9}
$y_{\text{pkg}_{\text{seam}}}$	3.0×10^{-9}	3.0×10^{-9}

Ta/Al seam participation along with the other participations for both the Al- and Ta-based transmons are given in Table 5.2 (see Appendix C for simulation details). It should be noted here that some care should be taken when calculating participations of transmons. Since a significant ($> 90\%$) of the total magnetic energy of the transmon mode is stored in the kinetic inductance of the Josephson junction, the total magnetic energy U_{tot} must be calculated to include the energy stored in the junction as well as the energy stored in the magnetic field; $U_{H_{\text{tot}}} = \int_{\text{all}} \mu_0 |\vec{H}|^2 dv + \frac{1}{2} L_J I_J^2$. This can significantly affect the calculation of participation ratios that rely on computing magnetic energy or current density, as is the case for conductor loss and seam loss. Additionally, calculating participations of near-junction regions with size of order $\sim 1 \mu\text{m}$ is too computationally intensive to do using global electromagnetic simulations. In this region, the small capacitance between the junction electrodes give rise to a highly localized electric field with high energy density. As a result, this region is simulated using a local 3D electrostatic simulation where a voltage is applied across the electrodes and the electric field energies in the various interfacial regions are calculated. Participations from the region within 100 nm of the junction are excluded from the total surface participation due to the reasoning provided by Wang et al. [74]. Given the small size of the junction region and low TLS densities $\sim 1 \mu\text{m}^{-2} \text{GHz}^{-1}$ measured in Al/AlO_x/Al junctions[114, 144–146], The region is expected to include zero

TLs and therefore appear virtually lossless, assuming that surface dielectric losses are dominated by TLs. Indeed, earlier studies have bounded the loss tangent of the junction oxide to below 4×10^{-8} [38], which supports this argument and implies that dielectric loss of the junction region would limit the transmon $Q_{\text{int}} > 10^9$. Additionally, recent quasiparticle tunneling experiments have shown that charge-parity switching lifetimes can be as high as several hundreds of milliseconds if the appropriate radiation shielding and microwave filtering are used[98] (see Ch. 3.5.2).

5.1.3 Characterizing Ta/Al Contact Loss

Contact resistance between the Ta and Al films may contribute significantly to the total loss of the transmon. To alleviate this loss, we use an in situ argon ion beam cleaning process that removes the tantalum oxide and any surface contaminants before depositing aluminum without breaking vacuum. With cross-sectional TEM, we can observe the effects of this process, as shown in Fig. 5.2b. Here, the Ta/Al interface seems to be quite rough; presumably, the ion beam has milled through the oxide and into the tantalum itself, damaging the surface of the tantalum. The unprocessed tantalum surface, on the other hand, appears free from damage and is therefore much more smooth, as can be seen in Fig. B.3a. Additionally, an element map formed by energy-dispersive X-ray spectroscopy (Fig. 5.2c) shows no detectable traces of oxide or other contaminants between the Ta and Al, which indicates that a good metal-to-metal contact is present.

To characterize and quantify this loss channel, we use a device with a unique design: the segmented stripline. This device is analogous to the bump-bonded stripline, also known as the interrupted stripline, discussed in Lei et al. [82] and in Lev Krayzman's thesis[83]. In this device we pattern around 625 alternating segments of Ta and Al that contact each other, resulting in 624 contact points, all of which are assumed to be resistive. Each segment is $10 \mu\text{m}$ wide and $\approx 20 \mu\text{m}$ long. The segments are patterned in

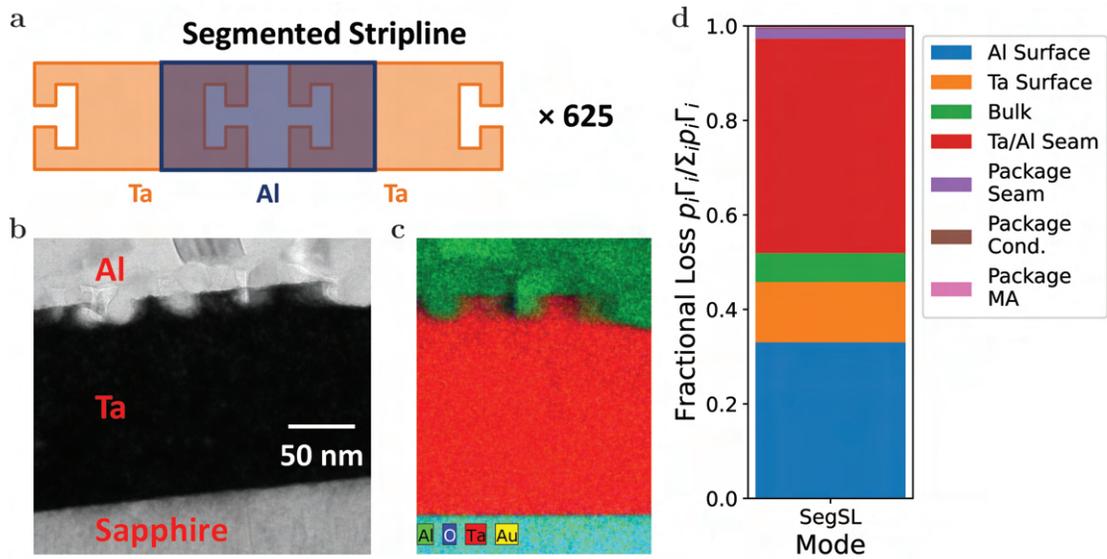


Figure 5.2: **Segmented stripline to extract Ta/Al contact loss.** **a** Segmented stripline design; alternating segmented of Ta and Al are repeated to make the full length of the stripline. **b** Brightfield TEM of Ta/Al interface, showing damage to the Ta film caused by the ion beam cleaning prior to Al deposition. **c** Elemental map formed by energy-dispersive X-ray spectroscopy (EDS) of the TEM sample shows no oxide in the Ta/Al interface, indicating that good metal-to-metal contact is present. **d** Single-photon loss budget for a representative segmented stripline, showing that 50% of the device's total loss is due to Ta/Al contact loss. TEM done by Kim Kisslinger at Brookhaven National Laboratory. Figure obtained from Ganjam et al. [86]

the same way as a Ta-based transmon: Ta segments are patterned subtractively, and Al segments are patterned additively. The result is a stripline that is 12.5 mm and is over 100 times more sensitive to Ta/Al contact loss than a transmon. A diagram of the segmented stripline is shown in Fig. 5.2a.

Because the segmented stripline is a thin-film resonator patterned on a substrate and inserted into a coaxial tunnel package, it is sensitive to the same sources of loss as the tripole striplines, with the addition of Ta/Al contact loss. Surface, bulk, and package participations are calculated using the same methods. For the contact loss, we use the same seam loss model as we did for the transmon and the package seam. The segmented stripline defines a $\lambda/2$ resonance mode, and because the stripline is very narrow (10 μm), we can assume that the current mostly flows along the propagation axis of the stripline. As a result, we can use an analytical model to calculate $y_{\text{seamTa/Al}}$ that assumes a sinusoidal current distribution throughout the stripline. The $y_{\text{seamTa/Al}}$ is therefore given by:

$$y_{\text{seamTa/Al}} = \frac{2 \sum_i \sin^2(\pi z_i/l)}{\pi w Z_0} \quad (5.2)$$

where $l = 12.5$ mm is the total length of the stripline, z_i is the position of the i th seam, $w = 10$ μm is the width of the stripline, and $Z_0 \approx 212 \Omega$ is the characteristic impedance of the stripline mode. The characteristic impedance can be calculated quite simply in a cross-sectional simulation using Ansys Q3D Extractor. The participations for the segmented stripline are given in Table 5.3.

We can use the loss factors extracted from the tripole striplines to account for the surface, bulk, and package losses in the segmented stripline. This allows for the measurement of the Ta/Al contact resistance by just measuring the quality factors. Two nominally identical segmented striplines were measured at single-photon powers. The quality factors and extracted contact resistances $R_{\text{Ta/Al}} = (w \cdot g_{\text{seamTa/Al}})^{-1}$ of the two devices are given in

Table 5.3: Segmented stripline participations, calculated using the methods in Appendix C. Table obtained from Ganjam et al. [86]

Freq (GHz)	5.74
p_{surfTa}	1.6×10^{-4}
p_{surfAl}	1.6×10^{-4}
p_{bulk}	0.72
$y_{\text{seamTa/Al}}$	9.4×10^4
$p_{\text{pkg}_{\text{cond}}}$	3.3×10^{-6}
$p_{\text{pkg}_{\text{MA}}}$	4.8×10^{-8}
$y_{\text{pkg}_{\text{seam}}}$	2.3×10^{-6}

Table 5.4: Segmented stripline loss. Table obtained from Ganjam et al. [86]

Device	$Q_{\text{int}}(\bar{n} = 1)$	$R_{\text{Ta/Al}} (\text{n}\Omega)$
SegSL1	1.97×10^6	246 ± 59
SegSL2	1.88×10^6	272 ± 73

Table 5.4. The fact that the quality factors differ by less than 10% indicates that the quality of the contact can be reliably achieved. From these two measurements, we calculate an average contact resistance of $260 \pm 47 \text{ n}\Omega$. We can also calculate a loss budget for the segmented stripline which shows that around 50% of the total internal loss is from the Ta/Al contact. This verifies that the segmented stripline has high sensitivity to this loss channel. Comparing this contact resistance with the bounds on surface resistance calculated in Ch. 4.5.5, we see that this contact resistance is over 100 times larger, indicating that the Ta/Al interface is resistive far beyond what would be expected from a continuous superconducting thin film.

It is interesting to compare the extracted contact resistance to conductor loss. In Ch. 4.5.5, I calculated some limits on residual surface resistance in thin-film Ta and Al to be $< 1 \text{ n}\Omega$. In a perfect metal-to-metal contact, the contact is equivalent to a continuous thin-film with resistance equal to R_s estimated in the previous chapter. From this, it is apparent that the Ta/Al contact is not perfect, and that there is some resistive current flow. With a measured resistance of $R_{\text{Ta/Al}} = 260 \pm 47 \text{ n}\Omega$ per Ta/Al contact, we can use Eq. (5.1)

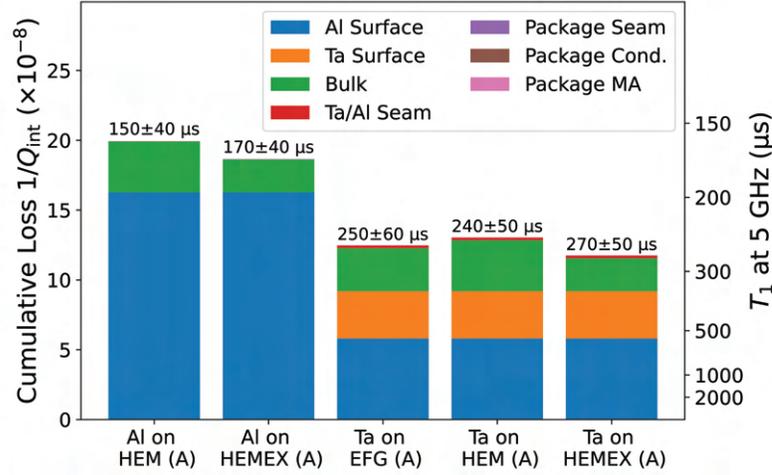


Figure 5.3: **Predicting transmon loss.** Predicted loss and expected T_1 for transmons made using different materials and processes (Al vs. Ta capacitor pads). Loss budget is also computed, showing the dominant sources of loss in Al- and Ta-based transmons. Figure obtained from Ganjam et al. [86]

estimate the loss to limit the transmon quality factor to $Q_{\text{int}} \approx 5 \times 10^8$. Therefore, while the Ta/Al contact is not perfect, it currently does not seem to limit the loss of a transmon.

5.1.4 Predicting Transmon Loss

Now that we have fully characterized the sources of loss that limit transmon relaxation times, we can use the transmon's participations from Table 5.2 to calculate the expected quality factors $Q_{\text{int}} = \omega T_1$ and their ranges for Al- and Ta-based transmon qubits. These predictions are shown in Fig. 5.3, where I have made predictions for Al- and Ta-based transmons on annealed sapphire substrates. The errors on the loss factors are propagated into the predicted T_1 's, yielding a range of values that reflect the inhomogeneity or natural variation of the loss factors due to the consistency (or lack thereof) of the fabrication processes employed to make the device.

From these predictions, we see that Al-based transmons are expected to achieve relaxation times of around 150 – 170 μs at 5 GHz, limited primarily by surface losses due to the

aluminum-based process (blue). Now if we replace the capacitor pads with tantalum using the respective process, the reduced surface loss (orange) is expected to yield dramatically improves T_1 's that exceed 240 μ s. This improvement seems to be relatively independent on sapphire grade (green), since annealed EFG, HEM and HEMEX sapphire all have approximately the same loss factor, as shown in Fig. 4.9e. From these predicted loss budget, the difference between Al- and Ta-based transmons is clear: surface loss is dramatically reduced. However, almost half of the Ta-based transmon's loss is from surface losses attributable to the aluminum region near the junction. While this is an extremely small area, the capacitance across the junction electrodes induces large electric fields that are densely localized, leading to high surface participations in this region. This, coupled with the large surface loss factor associated with the aluminum process, results in a much higher loss contribution. Additionally, bulk loss accounts for 15-20% of the Ta-based transmon's loss, which reflects the fact that as new materials and processes are utilized to reduce surface loss, the other sources of loss begin to play a larger role. As a result, bulk losses must be considered for transmons with T_1 's in the hundreds of microseconds range. Finally, as discussed in Sec. 5.1.3, the Ta/Al contact loss is negligible (red), and due to the compact electromagnetic field profile of the transmon, package losses are also very low.

5.1.5 Verifying Transmon Coherence Predictions

The transmon predictions allow us to see how a Ta- or Al-based fabrication process can affect qubit coherence. In particular, as Fig. 5.3 shows, Ta-based transmons are expected to have significantly higher T_1 than Al-based transmons. To verify this, Al- and Ta-based transmons were fabricated and their internal quality factors were measured and compared with the ranges predicted in Fig. 5.3. Transmon T_1 and T_2 were measured using the techniques described in Ch. 3.5.4. Transmons were fabricated and measured using an Al-based process and a Ta-based process, and using different grades of annealed sapphire.

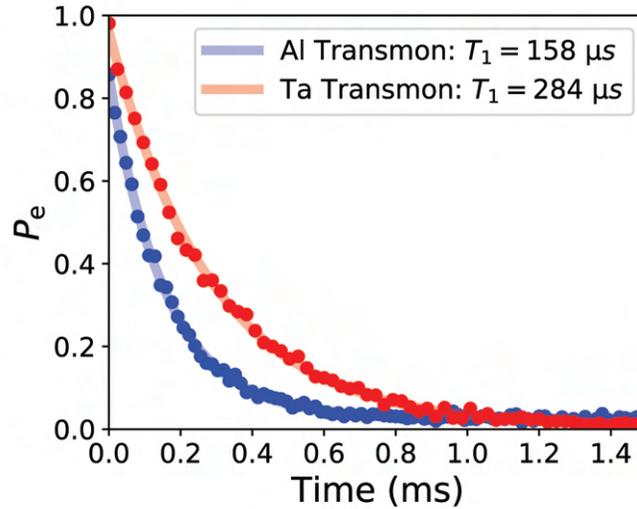


Figure 5.4: **Al- and Ta-based transmon T_1 .** Representative Al- and Ta-based transmon T_1 curves showing an almost factor of 2 improvement by adopting a tantalum-based process. Figure adapted from Ganjam et al. [86]

Fig. 5.4 shows representative T_1 measurements for and Al-based transmon (blue) and a Ta-based transmon (red). Here, it is clear that the Ta-based transmon has significantly higher relaxation times than the Al-based transmon, consistent with the predictions made in Sec. 5.1.4.

The coherence of each transmon is measured over at least a ten-hour period to capture temporal fluctuations, with some devices being measured over two days. Significant fluctuations in both T_1 and T_2 over long timescales, with T_1 fluctuating by as much as $\pm 30\%$ about the mean (Fig. 5.5a). These fluctuations can be attributed to fluctuating TLSs near or inside the Josephson junction, where the electric field densities are very high and coupling to TLSs can be quite large. Despite the statistical expectation to find zero TLSs in the region of the junction as discussed in Sec. 5.1.2, the variance around this expected value is high. Moreover, the region’s small area and high energy density renders the transmon highly sensitive to even small deviations from that expectation due to stochastically fluctuating TLSs both in space and frequency[100, 121]. This leads to large fluctuations in

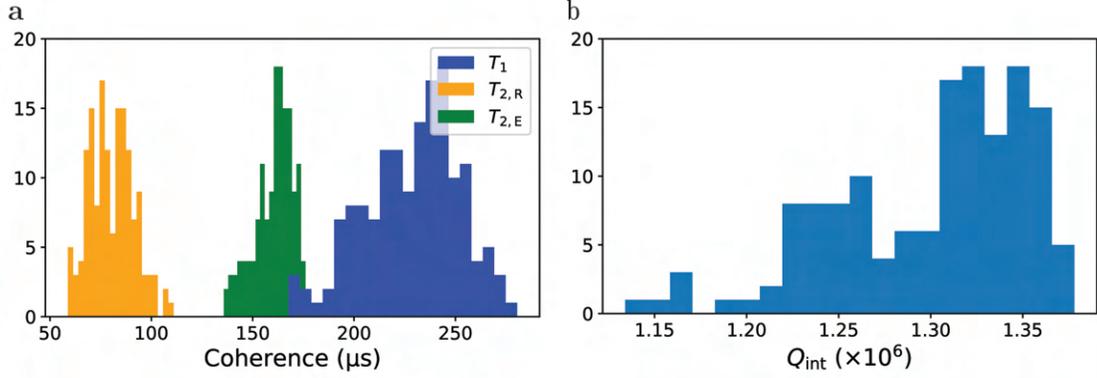


Figure 5.5: Temporal fluctuations in transmon coherence & resonator quality factor. **a** Representative histogram of temporal fluctuations for coherence in a Ta-based transmon device measured over 48 hours. **b** Representative histogram of temporal fluctuations for Q_{int} in a tripole stripline D1 mode measured over 35 hours. Data was taken at low power in the TLS-dominated regime, with $\bar{n} \sim 100$. Figure adapted from Ganjam et al. [86]

coherence over long periods of time due to TLSs drifting in frequency and becoming resonant with the transmon. This behavior is not captured by resonators and therefore cannot be accounted for by the loss model. Resonator internal quality factors tend to fluctuate by only $\pm 10\%$ about their mean, as shown in Fig. 5.5b. This is due to the resonator's much larger area and more dilute electric field. Resonators weakly couple to a large number of TLSs, whereas a transmon can couple very strongly to a small number of TLSs. Consequently, a single TLS fluctuation can dramatically affect the Transmon's coherence while negligibly affecting the resonator's.

The presence of temporal coherence fluctuations implies that loss factors extracted from resonator measurements do not completely explain transmon coherence, as TLS fluctuations near the junctions are not captured. Nevertheless, the fact that resonator quality factors fluctuate by $\pm 10\%$ and are insensitive to single TLS fluctuations indicates that the upper 80th percentile of transmon T_1 's are fluctuations that are captured by resonators. The remaining fluctuations are due to single or small numbers of TLSs fluctuating in frequency near the junction. As a result, we can conclude that resonators can be used to predict the upper 90th percentile of T_1 's achievable by a transmon as its coherence fluctuates over

long timescales.

To compare the upper 90th percentile of transmon T_1 's with the predicted values, we must first determine the transmons' internal quality factor. Since $Q^{-1} = Q_{\text{int}}^{-1} + Q_c^{-1}$, Q_{int} can be calculated by determining the external coupling of the transmon Q_c . To do this we use Eq. (3.44) and calculate Q_c by measuring the Rabi rate. Measured transmon Q_c varies between $3 - 7 \times 10^7$, consistent with the expectation that the simulated $Q_c \approx 4.5 \times 10^7$ can vary by around 50% due to ± 0.5 mm variations in chip positioning within the tunnel. With this knowledge, we can now calculate the 90th percentile of Q_{int} for each transmon and compare them with the predicted values. shown in Fig. 5.6. The predicted ranges are given by one standard deviation above and below the predicted quality factor. Here, we see that the majority of devices measured fall within this predicted range. These measured Q_{int} 's for Ta-based transmons are also similar to those measured in Place et al. [64] and Wang et al. [39]. The fact that the loss model is able to accurately predict transmon relaxation over different processes speaks to its robustness and usefulness in improving our understanding of how internal losses affect device coherence.

Finally, transmon coherence was also measured over long timescales. Both T_2^R and T_2^E fluctuated significantly over long time scales. Interestingly, T_2 varied wildly across different devices regardless of materials or processes used. Many transmons such as the one in Fig. 5.5a had $T_2^R \approx 50 - 100 \mu\text{s}$ and $T_2^E \approx 150 \mu\text{s}$, indicating that low-frequency noise was present that was refocused in the Hahn echo experiment. However, at times the T_2^R fluctuated, seemingly "jumping" to $T_2^R \approx 150 - 200 \mu\text{s}$. Given the significant temporal coherence fluctuations observed in all transmons measured, it is reasonable to assume that the origin of the low-frequency noise may be from TLS polarization noise, as was discussed in Ch. 3.4.1. In rare instances such as the Ta-based transmon whose coherence is shown in Fig. 3.13, $T_2^R \approx T_2^E$, which indicates that high-frequency noise must be the dominant source of dephasing, and is consistent with photon shot noise from

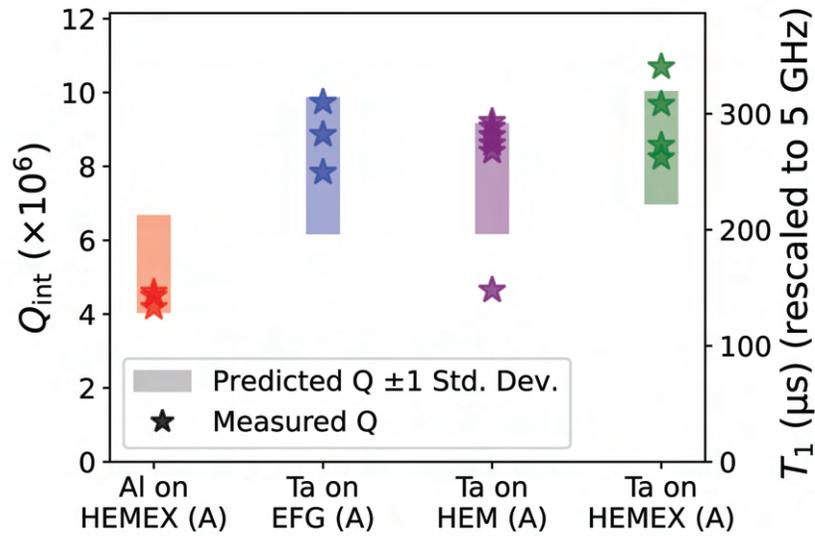


Figure 5.6: **Predicted vs measured transmon quality factors.** Measured transmon Q_{int} compared with predictions. Stars represent the 90th percentile transmon Q_{int} of a distribution formed from repeated coherence measurements over a 10-hour period. Shaded regions represent a predicted range spanning one standard deviation away from predicted transmon Q_{int} . Measured qubit frequencies ranged from 4.5 GHz to 6.7 GHz. Figure adapted from Ganjam et al. [86]

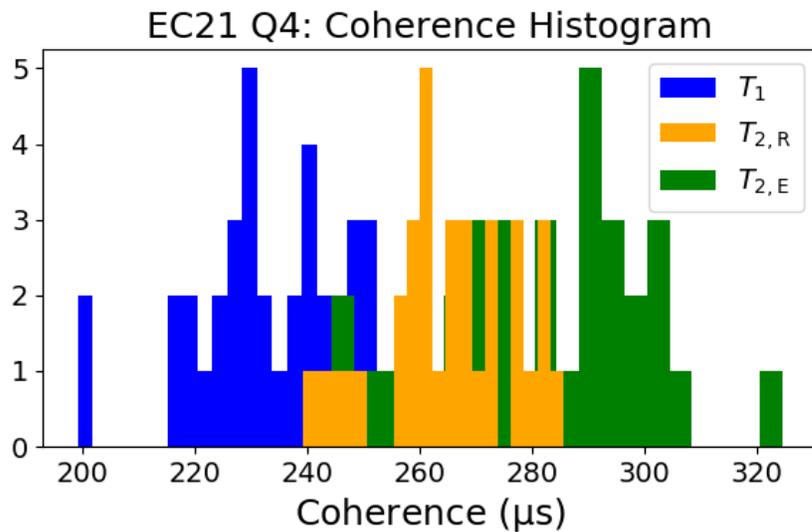


Figure 5.7: **Temporally “stable” transmon.** Histogram of temporal fluctuations for coherence in a temporally “stable” Ta-based transmon device measured over 14 hours.

a readout resonator strongly coupled to a control line at an equilibrium temperature of ≈ 62 mK. In these devices, it may be the case that there are no strongly-coupled TLSs in the region of the junction. Indeed, for this device, not only is the T_2^R high and not significantly limited by low-frequency noise, but the T_1 seems to fluctuate less, as seen in Fig. 5.7. Here, the transmon coherence fluctuates by only $\pm 10\%$, similar to that of a resonator. In no device was $T_2 = 2T_1$ ever achieved with either a Ramsey or Hahn echo experiment. This may be due to inadequate thermalization of control lines to the 20 mK stage of the dilution refrigerator.

5.2 Optimized Geometry to Maximize Coherence in a Quantum Memory

The transmon loss model was highly successful in predicting transmon relaxation (TLS fluctuations notwithstanding). Importantly, the analysis revealed that Ta-based transmons are significantly limited by surface losses near the junction. This motivates a more optimized design choice where a linear resonator is used to encode quantum information[53]. This was discussed in Ch. 2.3.3, this approach implements what is known as a quantum memory. The electromagnetic fields of linear resonators are distributed over a larger area, resulting in reduced surface participation and reduced sensitivity to single TLS fluctuations. As a result, they tend to have a higher Q_{int} regardless of what materials or fabrication processes are employed. Additionally, the lack of a Josephson junction prevents the existence of a region of small area where large electric fields can persist and strongly couple to TLSs. This allows linear resonators to be temporally more stable and have dramatically suppressed pure dephasing. Quantum memories have been demonstrated to great success in 3D cavity resonators, where an ancilla transmon is used to encode a logical qubit and manipulate nonclassical Bosonic states in the resonator[53, 76, 80]. Quantum memories

have also been demonstrated in thin-film resonators[84, 85]; unfortunately, their coherence has been far below their 3D counterparts. However, in the previous chapter I have shown several advancements in materials and fabrication processes that have led to the development of thin-film resonators with $Q_{\text{int}} > 3 \times 10^7$ at single-photon powers. It is therefore possible to optimize the design of a resonator to support a highly coherent on-chip quantum memory within the coaxial architecture. This device would have the advantages of a planar device due to its more compact design and lithographically-defined couplings.

5.2.1 Quantum Memory Device Design

From Table 4.5, we can see that the D2 mode of the tripole striplines generally have the highest internal quality factors. This is due to their relative insensitivity to both surface and package loss. To implement stripline-based quantum memory, a similar principle is followed, resulting in what we call the hairpin stripline, shown in Figs. 5.8a, b. The hairpin stripline is a multimode resonator whose fundamental mode is optimized to balance package and surface loss to maximize its Q_{int} . The insensitivity to package loss is obtained by folding a half-wave (of length $2L$; see Fig. 5.8a) stripline resonator into itself, resulting in a fundamental mode whose field are localized primarily between the two arms of the hairpin. As seen in Fig. 5.8a, the electric fields of the fundamental mode extend from one arm to the other. This mode is in some sense “differential”; the electric field lines do not terminate at the walls of the package. Additionally, the current antinode is positioned at a location where the current flows perpendicularly to the axis of the cylindrical tunnel, resulting in minimal induced current along the walls of the package. As a result, the fundamental mode of the hairpin stripline becomes insensitive to package losses, while the large spacing between the arms dilutes the electric field on the surface, resulting in dramatically reduced surface participation.

Fig. 5.8b shows the circuit layout of the chip, designed to be inserted and measured in

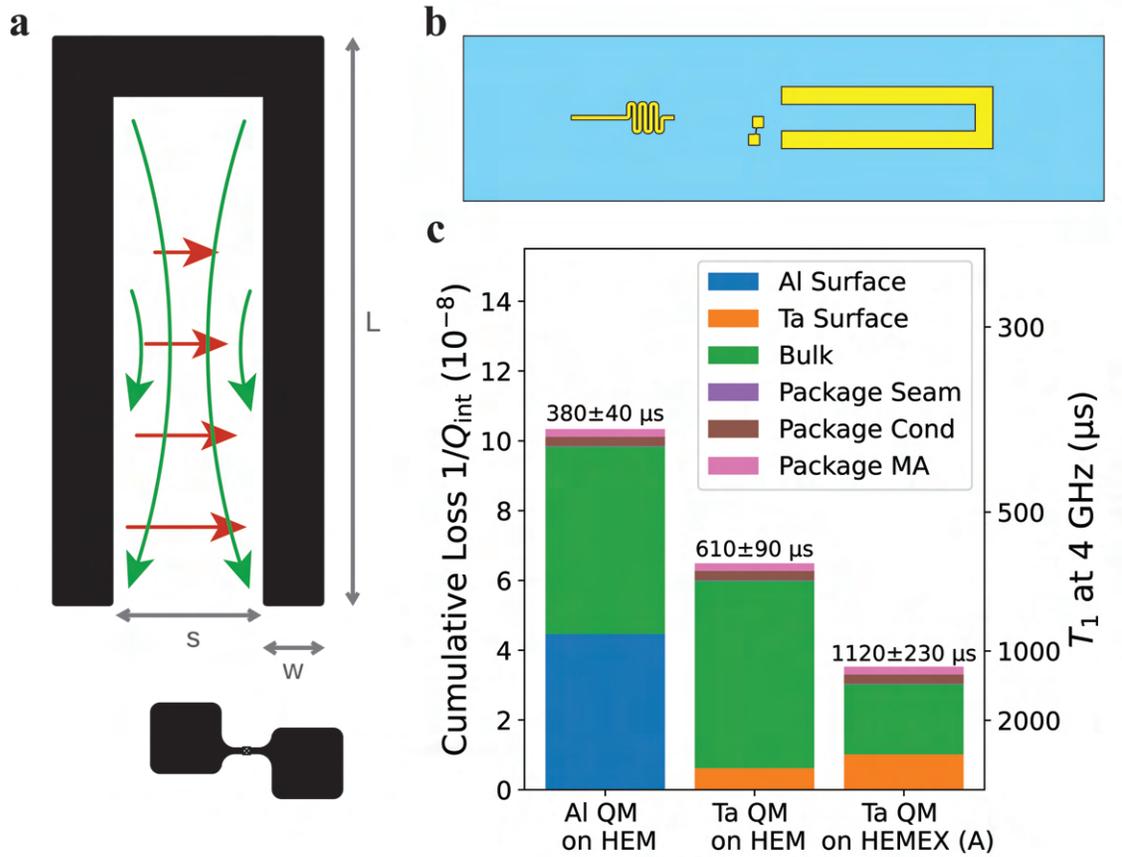


Figure 5.8: **Hairpin stripline quantum memory.** **a** Field behaviors of the memory mode (red arrows) and readout mode (green arrows) are shown. The ancilla transmon has a staggered capacitor pad design that allows coupling to both modes. Optimal design parameters are $w = 800 \mu\text{m}$ and $s = 1200 \mu\text{m}$. The length $L = 9.5 \text{ mm}$ of the hairpin stripline sets storage mode frequency at $\approx 4 \text{ GHz}$. **b** Hairpin stripline quantum memory design. The ancilla transmon couples to the fundamental mode that acts as a storage resonator, and to the higher order mode that acts as a readout resonator. A Purcell filter is used to enhance the external coupling of the readout mode. **c** Predicted loss and expected T_1 for hairpin striplines made using different substrate preparations and different superconducting thin-films. Figure adapted from Ganjam et al. [86]

a cylindrical tunnel package. Once again, the same control line is used to control the qubit, the readout drive, and the quantum memory drive. The chip contains the hairpin stripline, an ancilla transmon, and a stripline Purcell filter. The ancilla transmon is designed to couple to both the fundamental mode, which will be used as the quantum memory, and the second-order mode, which will be used for ancilla qubit readout. The readout mode is a full-wave resonance mode, where the electric field antinodes are located at the ends of the hairpin and have opposite polarity, which results in an electric field that is orthogonal to that of the memory mode. In order to couple to both modes simultaneously, the capacitor pads are staggered with respect to each other to re-orient its dipole moment such that it can dispersively couple to both modes. The Purcell filter is then used to enhance the external coupling of the readout resonator. The hairpin stripline’s parameters are given in Table 5.5.

Table 5.5: Typical hairpin stripline parameters. Table obtained from Ganjam et al. [86]

$\omega_m/2\pi$ (GHz)	3.9-4.0
$\omega_t/2\pi$ (GHz)	5.7-6.8
$\omega_r/2\pi$ (GHz)	9.0-9.3
$\chi_{tt}/2\pi$ (MHz)	201-217
$\chi_{tm}/2\pi$ (MHz)	0.1-0.4
$\chi_{tr}/2\pi$ (MHz)	0.3-0.6
$\kappa_r/2\pi$ (MHz)	0.2-0.5

5.2.2 Hairpin Stripline Participations

To maximize the coherence of the hairpin stripline, it is necessary to use the best materials and fabrication processes (i.e. the ones that yield the lowest loss factors) as well as optimize device design to minimize participations in lossy regions. There are three parameters to be optimized: the width of the stripline w , the spacing between the arms of the hairpin s , and the tunnel radius. Increasing w slightly reduces surface participations,

while increasing s reduces both bulk and surface participation. However, increasing both while keeping the tunnel radius fixed increases package participation, as the arms of the hairpin becomes closer to the walls of the tunnel. For a fixed tunnel radius of 2.5 mm, the optimal values for w and s that minimized the total loss $\sum_i p_i \Gamma_i$ were found to be 800 μm and 1200 μm , respectively. The participations of the optimized design are given in Table 5.6.

Table 5.6: Hairpin stripline participations. Table obtained from Ganjam et al. [86]

p_{surf}	2.4×10^{-5}
p_{bulk}	0.72
$p_{\text{pkg}_{\text{cond}}}$	6.7×10^{-6}
$p_{\text{pkg}_{\text{MA}}}$	5.4×10^{-8}
y_{seam}	3.9×10^{-8}

5.2.3 Predicting Hairpin Stripline Loss

The predictive loss model is applied to the hairpin stripline by using the participations from Table 5.6. This is used demonstrate how optimizing materials and process choices can lead to dramatic coherence improvements. As shown in Fig. 5.8c, employing an Al-based process on unannealed HEM sapphire is not expected to produce remarkable coherence, due to the significance of surface loss. Replacing the aluminum with a tantalum-based process results in a modest improvement; however, when both high-temperature substrate annealing and tantalum processes are employed, the hairpin stripline is expected to reach a T_1 of (1.1 ± 0.2) ms, which rivals the coherence of the popularly used 3d coaxial post cavity[53]. This dramatic improvement is realized by optimizing both materials and geometry and was only possible because of loss characterization work described in Ch 4.

Remarkably, the optimized hairpin stripline design does not seem to be significantly limited by surface loss. In fact, it seems to be dominated by bulk dielectric loss. This breaks the pervasive thought in the field that all on-chip devices are dominated by surface

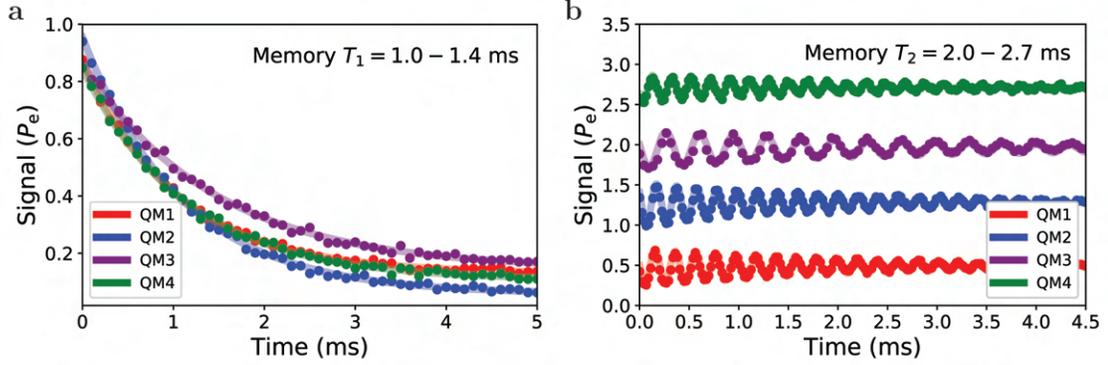


Figure 5.9: **Hairpin stripline quantum memory coherence.** **a** Fock state T_1 measurement of four on-chip quantum memory devices (see Ch. 3.5.5 for measurement methods). Memory T_1 's are 1.05, 1.09, 1.44, and 1.14 ms. **b** Memory T_2 in the Fock ($|0\rangle, |1\rangle$) manifold for the four devices measured in **a** (see Ch. 3.5.5 for measurement methods). Ancilla state as a function of time for QM2-4 are offset vertically by 0.75, 1.5, and 2.25, respectively, for visibility, and were fit to an exponentially decaying sinusoid. Extracted memory T_2 's for QM1-4 are 2.02, 2.00, 2.68, and 2.14 ms. Figure obtained from Ganjam et al. [86]

dielectric or TLS loss. In fact, by properly optimizing a device, one can eliminate its sensitivity to surface loss. Interestingly, this implies that further improvements in the coherence of the hairpin stripline must come from improvements in bulk dielectric loss.

5.2.4 Verifying Hairpin Stripline Coherence Predictions

To confirm the predictions made in the previous section, four hairpin stripline-based quantum memories were fabricated on annealed HEMEX-grade sapphire using a tantalum-based process. The devices were inserted into the same coaxial tunnel packages that were used to measure the transmons and tripole striplines. A single control line was used for readout, qubit, and memory driving. The Q_c for the memory mode was simulated to be $\approx 10^9$. Imprecision in chip positioning can also lead to significant variations in Q_c for this device. However, this variation leads to an expected $Q_c \approx 5 - 15 \times 10^8$, which is still over an order of magnitude higher than the expected $Q_{\text{int}} \approx 25 - 35 \times 10^6$. As a result, the external loss is expected to account for less than 5% of the total loss of the hairpin stripline

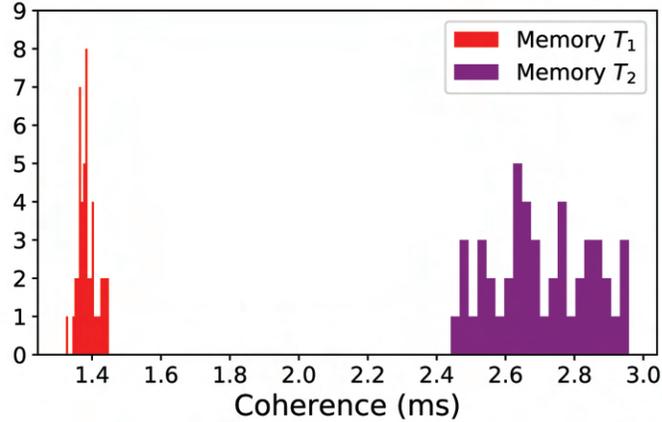


Figure 5.10: **Temporal fluctuations in quantum memory coherence.** Representative histogram of temporal fluctuations in coherence in a hairpin stripline quantum memory device measured over 30 hours. Figure adapted from Ganjam et al. [86]

storage mode.

Quantum memory coherence in the Fock ($|0\rangle, |1\rangle$) manifold was measured using the techniques described in Ch. 3.5.5. The measured T_1 was remarkably consistent with the predictions from Fig. 5.8c. Fock state decay times were measured to be 1 – 1.4 ms, as shown in Fig. 5.9a. Additionally, measured Fock T_2 times approached $2T_1$, indicating that there is very little dephasing, as shown in Fig. 5.9b. This measurement alone bounds the dephasing time to $T_\phi > 24$ ms, which is far higher than what is expected for a transmon and is similar to those of 3D cavity-based quantum memories[76, 147]. Finally, continuous coherence measurements over 20 hours showed minimal temporal fluctuations, as shown in Fig. 5.10. Neither T_1 nor T_2 fluctuated by more than $\pm 10\%$ over long timescales. This reflects a general insensitivity to TLS fluctuations that can be attributed to the much more spatially distributed electric field of the memory mode.

The implementation of a quantum memory in a stripline places the coaxial architecture as an attractive alternative to the cavity-based implementation of Bosonic qubits. For the purposes of scaling up to multiqubit systems, the coaxial architecture can be made more modular and more compact, thanks to the ability to define the quantum memory on

a chip that can be individually fabricated, mass produced on a wafer, and easily replaced if defective. Additionally, since the circuits are lithographically defined, devices and couplings can be engineered with greater precision. Finally, the low pure dephasing makes the hairpin stripline an attractive option to implement noise-biased qubits, which can enable lower error correction thresholds towards the implementation of surface codes of dual-rail qubits[148, 149]. Stripline-based quantum memories therefore provide a promising building block for realizing a quantum processor with bosonic qubits.

Chapter 6

Future Directions and Conclusion

The experiments described in Chs. 4 and 5 led to a much greater understanding of microwave loss mechanisms and how they affect device coherence. However, there are still many unanswered questions. In this chapter, I will discuss some open questions about loss mechanisms in superconducting circuits, as well as outline some promising directions for future exploration, some of which are already being pursued.

6.1 Open Questions About Microwave Losses

Conductor Loss

As discussed in Ch. 4, the conductor loss factor is much larger in bulk superconductors than in thin films. I showed in Ch. 4.4.5 that the R_s measured in 5N5 etched Al is not explained by vortex loss given the estimates of magnetic field strengths we have inside the Cryoperm shield. In fact, the residual resistance seems to be due to a large excess quasiparticle fraction $x_{qp} \approx 3 \times 10^{-4}$ whose origin is not known. Looking at the SRF community that works with niobium cavities, it seems as if there are many lessons to be learned with respect to surface processing in order to reduce the residual resistance of bulk

high-purity Al.

Additionally, a precise measurement of residual resistance in thin-film conductors has not yet been done. With the tripole striplines we are able to extract surface loss, which is a combination of dielectric and conductor loss. There is some evidence from quasiparticle tunneling experiments that x_{qp} in thin-film aluminum can be very low, $\sim 10^{-10}$ [98], which indicates that excess quasiparticles are an insignificant contributor to loss in thin-film devices. This leaves vortex loss, which could be significant; however, we do not know the strength of the cooling field to accurately measure it. In Ch. 4.5.5 I used the surface loss factor to bound R_s to around < 1 n Ω ; this was used to bound the cooling magnetic field strength to < 4 nT, and if in reality it is much lower, then even vortex loss would be an insignificant contributor to the loss. Nevertheless, since conductor losses are dependent on the conducting and superconducting properties of the materials used, it is important to consider residual resistance as new materials are explored. Importantly, the bounds I have placed on residual resistance in thin films were obtained by using the high-power surface loss factor; in other words, I am assuming that conductor loss is power independent. This follows the traditional way of thinking about power dependence in resonators; i.e. TLSs couple to electric fields. However, if magnetic fields can couple to anomalous surface spins, then conductor loss could show saturating behavior at high powers. Uniquely determining the conductor loss factor could help explore this.

TLS Loss

In cQED, TLSs are treated very phenomenologically, and are often times used to describe behavior that we do not understand. The physical nature of a TLS is relatively unknown; it is unclear what molecular or atomic defect, or what dangling bond in an oxide or amorphous surface layer is responsible for giving rise to the behavior that we see in thin-film resonators and transmons. Moreover, the models that we use to describe TLS loss in res-

onators do so in an ensemble sense; it is difficult to make any meaningful statements other than a qualitative comparison of TLS behavior in two different resonators. Understanding where TLSs come from requires a large-scale interdisciplinary effort to bring together cQED specialists and materials scientists to enable an in-depth exploration into lossy interfaces using physical materials characterization tools. Such collaborations are already underway and provide an optimistic outlook towards answering some of these questions.

Frequency Dependence of Loss Factors

In Ch. 4 when describing the FWGMR and TSL experiments, I implicitly assumed that the loss factors were either independent or weakly dependent on frequency. This assumption may not hold, and for some sources of loss there is an explicit frequency dependence. In particular, both quasiparticle and vortex loss are frequency dependent. Additionally, dielectric loss may also be frequency dependent. Creedon et al. [103] and Read et al. [106] both measured similar values for the loss tangent of HEMEX-grade sapphire at very different frequencies; this indicates that bulk dielectric loss could have weak frequency dependence. However, it should be noted that the sapphire samples in these experiments were prepared very differently; in Read et al. [106] the sample was cut into a shape of a wafer that was later cleaved into the desired shape, while Creedon et al. [103] measured a whispering-gallery mode resonator made of HEMEX sapphire. It may be a coincidence that these two experiments measured the same number. Fortunately, the tripole stripline modes were all within a narrow frequency range of 4 – 7 GHz, and the transmons and quantum memories described in Ch. 5 lay within this frequency range as well. However, if these loss factors are significantly frequency dependent, then quality factor predictions for devices of much higher or lower frequencies would accumulate significant systematic error and may no longer be accurate. In the next section, I will propose an experiment to measure the frequency dependence of loss factors.

6.2 Future Directions in Microwave Loss Characterization

6.2.1 Exploring Frequency Dependence with Multimode Resonators

The benefit of using multimode devices to characterize loss is that the higher order modes of the resonator can also be included in the loss analysis. In the tripole stripline, for example, the second order D1, D2, and C modes are at 9 – 12 GHz. Measurement of these modes and subsequent loss analysis can answer questions about the frequency dependence of the loss factors. While these higher order modes were not measured for this project due to the mode frequencies being outside of the current HEMT bandwidth, a future redesigning of the experimental setup could enable the measurement of the higher order modes. Measuring multimode resonators that are sensitive to conductor loss can be particularly illuminating because the different mechanisms of conductor loss have different frequency dependence. Vortex loss, for example, scales as $R_{\text{vortices}} \propto \omega^{1/2}$, whereas surface resistance due to a fixed quasiparticle fraction scales as $R_{\text{qp}} \propto \omega^{3/2}$. Additionally, other proposed mechanisms for residual resistance exist and have different frequency dependence[93]; exploring these mechanisms may answer important questions about the origins of conductor loss.

6.2.2 Measuring Conductor and MA Losses with Flip-Chip Circuits

In Ch. 4.5 I showed how the D1 mode of the tripole stripline is sensitive to surface loss. However, it is not able to distinguish between conductor loss and surface dielectric losses in the three interfaces. Doing this is very difficult, as these participations tend to scale together with geometry; increasing one of these participations tends to also increase the

other. This tends to happen regardless of geometry in planar topologies where the circuit is defined on a single plane (i.e. the surface of the wafer). Some innovative solutions exist that utilize substrate etching to change the profile of the bulk dielectric[105, 107, 108]; however, this requires extra processing that by itself changes the surface loss properties.

In a flip-chip architecture, circuits can be designed to localize electromagnetic fields in ways that are not otherwise possible. As shown in Fig. 6.1a, the circuits are patterned on two separate chips that are then mated with each other using indium bump-bonds. This results in a layout that is still planar, but circuit elements can exist on multiple planes; in some sense, this is a 2.5-dimensional architecture. The advantage here is that provided the gap between the two chips is small enough, effective parallel-plate capacitances can be achieved, allowing for the localization of electric field in the vacuum gap between the two chips. As a result, capacitors realized in this way would have much reduced bulk dielectric, MS, and SA participation, with the trade-off being increased MA participation, which enables the design of a device that is particularly sensitive to MA loss without being sensitive to MS or SA loss.

Losses in Flip-Chip Circuits

The realization of a flip-chip stripline, i.e. two striplines patterned on a top and bottom chip that are positioned on top of each other, results in a multimode resonator whose differential mode localizes the electromagnetic field in the vacuum gap (Fig. 6.1c). This creates an electric field that is perpendicular to the plane of the chips, resulting in high MA participation. However, the currents in the two striplines flow in opposite directions, generating a large magnetic field in the gap that runs parallel to the plane (Fig. 6.1a). As a result, the flip-chip stripline becomes highly sensitive to both MA loss and conductor loss, and less sensitive to MS and SA loss.

The ability to realize a parallel-plate capacitor in a flip-chip circuit provides a spatially-

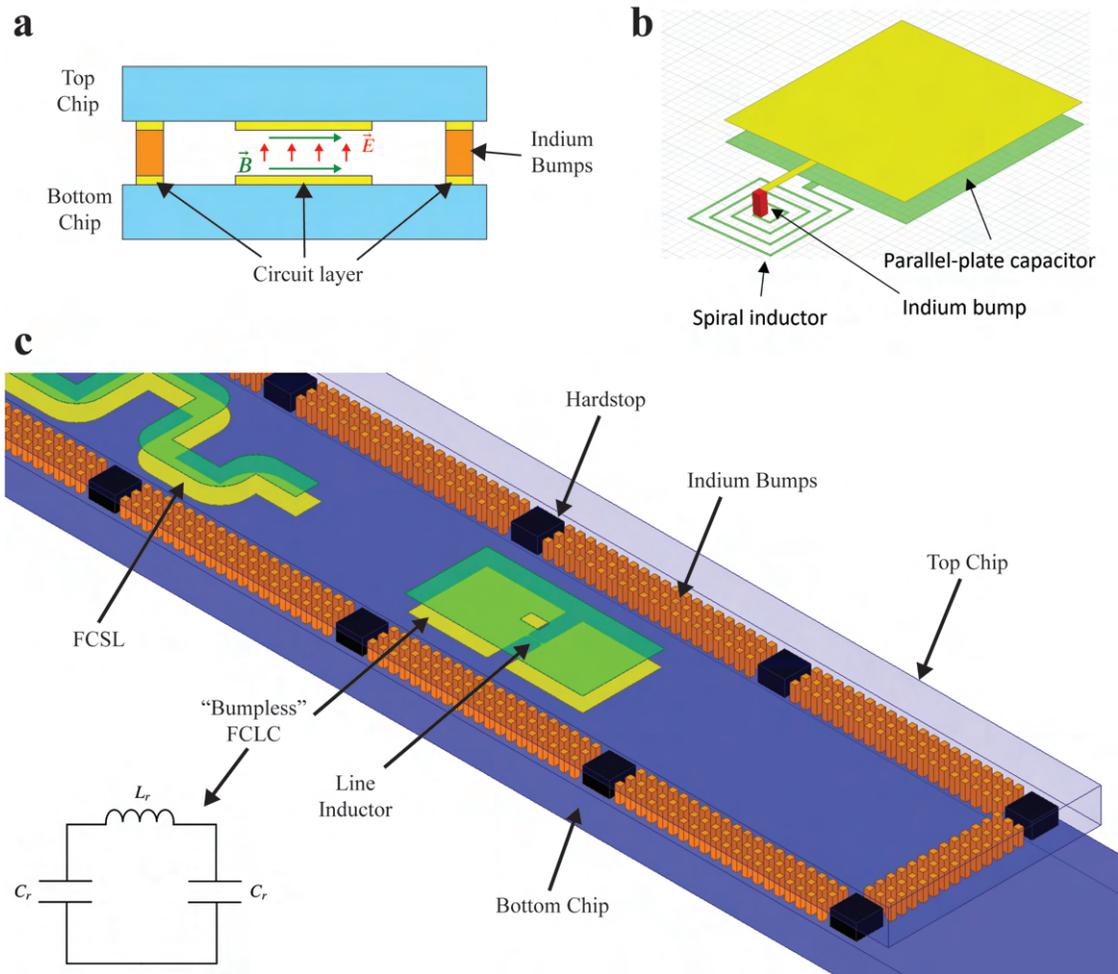


Figure 6.1: **Flip-chip circuit architecture.** **a** Cross-sectional view of a typical flip-chip device. Circuits are patterned on a top and bottom chip that are then bonded together using indium bumps that are deposited on the circuit layer. Electric fields run from the bottom plate to the top plate and are localized in the vacuum gap. In a flip-chip stripline, the magnetic fields run parallel to the plates and are also localized in the gap. **b** Flip-chip lumped element LC circuit (“bumped” FCLC). A parallel-plate capacitor is shunted by a spiral inductor. An indium bump bond connects the center of the spiral to the circuit layer on the top chip. **c** Diagram of a flip-chip loss characterization device. Circuits on the bottom chip are in yellow and circuits on the top chip are in green. A “bumpless” FCLC resonator is patterned by defining two parallel plate capacitors in series with equivalent capacitance $L_r/2$ that are connected to a line inductor. A flip-chip stripline (FCSL) is also patterned on the chip; two meandered striplines (width $400\ \mu\text{m}$) are placed on top of each other. Mechanical support bumps (orange) made of indium bond the two chips together and hardstops (black) made of titanium fix the gap between the chips. The gap between the two chips is $4\ \mu\text{m}$ and is exaggerated in this figure for visibility.

efficient way to realize a large capacitance with a high self-resonant frequency. Likewise, larger and more compact inductances are possible through spiral inductors, which are feasible to make in the flip-chip architecture. This allows for the design of a compact lumped-element resonator whose electric and magnetic fields can be spatially separated (Fig. 6.1b). An appropriately designed LC resonator can therefore localize the electric field within the capacitor, resulting in high MA participation, while distributing the magnetic field throughout the inductor, whose design can be independently tuned to reduce conductor participation.

The implementation of a flip-chip architecture introduces additional sources of loss. To position and hold the two chips together, indium bump-bonds are required. Indium is chosen for its low melting point and high ductility, allowing it to be easily cold-welded to itself. Indium bumps are patterned on an existing circuit layer on both chips, creating a M/In contact, where M is the circuit layer metal. When the two chips are bonded together, the bumps will compress and deform from the bonding force, breaking the oxide and forming an In-In contact between the top and bottom chips. This creates a bump stack-up of M/In-In/M (see Fig. 6.1a), where indium is deposited on the circuit layer of metal M. This stackup contains three distinct contact regions, all of which can be resistive akin to seam loss due to the presence of interfacial oxide or contaminants. This is a similar source of loss to the Ta/Al contact loss described in Ch. 5.1.3. Indium bumps perform the primary role of mechanically adhering the two chips together, but they can also be used as part of a resonator to carry current from one chip to the other, as in the case of a spiral inductor. As a result, the seam conductance associated with indium bump-bonds must be characterized. While this has been done for bump bonds realized through In/In-In/In[82] and Al/Au/In-In/Au/Al stackups[123], it has not been done for other superconducting materials. Additionally, conductor and surface dielectric loss from the indium bumps and hardstops used to define the gap between the chips can contribute to the total loss of flip-

chip circuits. Hardstops can be defined using dielectrics like silicon[150] or using metals such as aluminum or titanium that are harder than indium and will not compress under the application of bonding force.

Ensemble Loss Characterization of Flip-Chip Circuits

Losses in flip-chip circuits can be characterized using an ensemble of multimode devices. These devices are designed to fit in coaxial waveguide tunnels, similar to the tripole striplines, in order to maintain the same package environment. The bottom chip is identical in size to the TSL chips (see Ch. 3.5.1) at $40 \text{ mm} \times 4 \text{ mm}$ and the top chip is shorter in length at $30 \text{ mm} \times 4 \text{ mm}$. The two chips are bonded together using indium bump-bonds that line the edges of the chip (I will call these bumps as “mechanical support bumps”), with regularly spaced hardstops to define a flip-chip gap of $4 \text{ }\mu\text{m}$ after bonding.

A flip-chip lumped-element (“bumped” FCLC) resonator made with a large parallel-plate capacitor ($C \approx 2.8 \text{ pF}$; $1 \text{ mm} \times 1 \text{ mm}$ parallel plate size, $4 \text{ }\mu\text{m}$ gap) and a small spiral inductor ($L \approx 0.17 \text{ nH}$) is designed to have a single current-carrying indium bump-bond to connect the parallel plate on the top chip with the spiral inductor on the bottom chip (Fig. 6.1b). This device would have a very high $y_{\text{seambump}} = (Z_0 w)^{-1}$, where $Z_0 = \sqrt{L/C}$ and $w = 15 \text{ }\mu\text{m}$ is the side-length of the indium bump. Another type of lumped-element resonator has a bump-less design (“bumpless” FCLC; see Fig. 6.1c) where two capacitors in series (equivalent capacitance $C \approx 3.2 \text{ pF}$; $1.5 \text{ mm} \times 1.5 \text{ mm}$ parallel plate size, $4 \text{ }\mu\text{m}$ gap) are connected by a line inductor ($L \approx 0.3 \text{ nH}$). This device would have no current-carrying bumps and would therefore be limited primarily by MA loss; the line inductor would give rise to a spatially diffuse magnetic field, rendering the mode less sensitive to conductor loss. Additionally, the differential mode of the flip-chip stripline (FCSL D) would be sensitive to conductor loss due to its magnetic field localization, and the common mode (FCSL C), due to its diffuse electromagnetic fields, would be sensitive

to losses associated with the hardstops and mechanical support bumps (Fig. 6.1c). Finally, the tripole stripline can be used to extract $\Gamma_{\text{SAMS}} = \frac{p_{\text{SA}}}{p_{\text{SA}}+p_{\text{MS}}}\Gamma_{\text{SA}} + \frac{p_{\text{MS}}}{p_{\text{SA}}+p_{\text{MS}}}\Gamma_{\text{MS}}$, Γ_{bulk} , and g_{seam} . The participation matrix for this loss characterization system is given in Table 6.1.

Table 6.1: Flip-Chip Participation Matrix, assuming 4 μm flip-chip gap and $\lambda = 50$ nm. $p_{\text{pkg}_{\text{MA}}}$ and $p_{\text{pkg}_{\text{cond}}}$ are omitted here for simplicity. p_{mech} is the combined surface dielectric and conductor participation of the mechanical support bumps and hardstops, and was approximated by defining an integration surface that is a 0.5 mm-wide strip at the edge of the chip. The tripole striplines are defined on a single chip and will therefore will have zero sensitivity to losses related to bumps and hardstops. Participatons for the flip-chip devices were calculated using cross-sectional 2D simulations and coarse 3D simulations; a finer simulation using the methods in Appendix C was not done here and would likely change these participations by $\approx 20\%$.

Participation Matrix							
Mode	p_{SAMS}	p_{MA}	p_{bulk}	p_{cond}	$y_{\text{seam}_{\text{bump}}}$ $1/(\Omega\text{m})$	$y_{\text{seam}_{\text{pkg}}}$ $1/(\Omega\text{m})$	p_{mech}
TSL D1	1.1×10^{-3}	5.8×10^{-5}	0.90	6.1×10^{-2}	0	4.0×10^{-8}	0
TSL D2	3.4×10^{-5}	8.3×10^{-7}	0.80	8.0×10^{-4}	0	3.6×10^{-8}	0
TSL C	2.2×10^{-5}	5.3×10^{-7}	0.45	5.5×10^{-4}	0	1.4×10^{-5}	0
“bumped” FCLC	2.4×10^{-5}	1.4×10^{-4}	0.10	5.6×10^{-2}	8.5×10^3	9.0×10^{-9}	6.3×10^{-10}
“bumpless” FCLC	1.7×10^{-5}	1.4×10^{-4}	0.10	4.2×10^{-2}	0	2.1×10^{-7}	1.1×10^{-8}
FCSL D	5.6×10^{-5}	1.3×10^{-4}	0.10	3.4×10^{-1}	0	1.9×10^{-7}	6.0×10^{-9}
FCSL C	1.4×10^{-5}	4.3×10^{-7}	0.66	1.0×10^{-3}	0	1.1×10^{-5}	1.3×10^{-6}

6.2.3 Characterizing Losses in Rhenium

The determination of surface loss being the dominant source of loss in transmons and the exploration into the flip-chip architecture has motivated the search for materials that have lower surface losses. Rhenium (Re) is a refractory metal with $T_c \approx 1.9$ K (see Appendix B.4) and has a hexagonal crystal structure that has excellent lattice matching with the oxygen sublattice of sapphire[151]. This is unlike tantalum or aluminum, whose crystal

structures do not match well with that of sapphire, which may lead to lattice distortions in the MS interface. Since rhenium has better lattice matching, it might reduce MS loss. Additionally, it is thought that Re does not oxidize at room temperature, which could result in dramatically reduced MA loss and yield flip-chip devices that are both highly compact and low loss. This was confirmed using TEM (Fig. B.5a), where no oxide can be observed on the surface of rhenium after being exposed in air for over two months.

Rhenium tripole striplines and flip-chip devices (see Table 6.1) were fabricated on annealed HEMEX sapphire using the methods in Appendix A. Unfortunately, a fabrication error during the indium bumps patterning process resulted in the redeposition of indium onto the wafer, which substantially contaminated the surface of the device. Nevertheless, the devices were packaged and measured in the dilution refrigerator at single-photon powers. As of writing, the FCSL D mode has not yet been measured. Two TSL D1/D2 modes, two TSL C modes, two “bumped” FCLCs, and three “bumpless” FCLCs were measured during the cooldown. The measured quality factors of the devices are given in Table 6.2

Table 6.2: Rhenium Device Quality Factors Measured at $\bar{n} \approx 1$. Error on average Q is determined by the range of the two device Q s.

Mode	Device $Q_{\text{int}} (\times 10^6)$	Average $Q_{\text{int}} (\times 10^6)$
TSL D1	1.71, 1.72	1.72 ± 0.01
TSL D2	15.8, 17.7	16.8 ± 1.0
TSL C	16.4, 14.0	15.2 ± 1.2
“bumped” FCLC	3.00, 3.14	3.07 ± 0.07
“bumpless” FCLC	1.39, 1.17, 1.38	1.31 ± 0.14
FCSL D	-	-
FCSL C	0.019, 0.024	0.022 ± 0.003

The measured quality factors were slightly surprising. For one, the FCSL C mode was

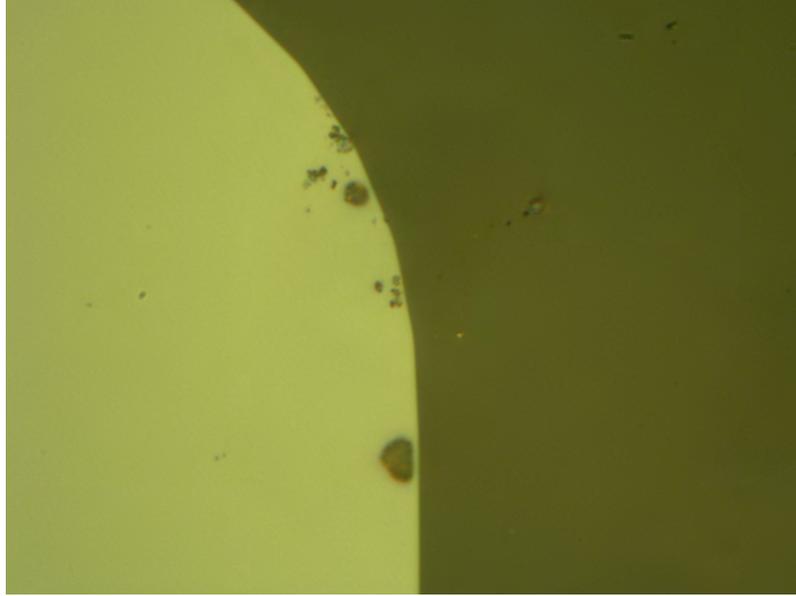


Figure 6.2: **Indium redeposition on rhenium devices.** Sonication during liftoff resulted in microscopic $\sim \mu\text{m}$ fragments of indium that redeposited onto the wafer. This redeposition resulted in a contaminated surface that may have affected the surface loss factors.

surprisingly low Q , which indicates that either mechanical support bumps or the hardstops contribute a significant amount of loss. It is unlikely that this is due to the indium, because micromachined cavities made with indium deposited in the same way had extremely high quality factors[82]. The hardstops, on the other hand, were made of titanium deposited through a mechanical shadow mask and may have poor film quality due to its ill-defined film edges. Additionally, the “bumpless” FCLC had much lower Q than the “bumped” FCLC despite having similar p_{MA} and an additional source of loss. However, the “bumpless” FCLC has higher p_{mech} , which implies that mechanical support loss is contributing somewhat to the total loss of these flip-chip devices.

The device quality factors were used along with the participation matrix in Table 6.1 to extract Γ_{SAMS} , Γ_{MA} , Γ_{bulk} , $g_{\text{seam}_{\text{bump}}}$, $g_{\text{seam}_{\text{pkg}}}$, and Γ_{mech} . Package MA and conductor losses were subtracted from the total internal loss in the same way as in Ch. 4.5.2. Since the conductor loss-sensitive mode was not measured, we will assume that it is negligible using the same arguments as in Ch. 4.5.5. It is important to note here that since the differ-

ent modes are not all from one single device, we cannot do a device-to-device examination of the variation of the loss factors. Rather, we will use the range of internal quality factors measured as the uncertainty around the average value, which will serve as a proxy for device-to-device variation. Loss extraction using the least-squares algorithm revealed an unconstrained fit for g_{seambump} ; likely due to the fact that Γ_{MA} and Γ_{mech} are far too high and inhibits the sensitivity of the “bumped” FCLC mode to the bump loss. Instead, a non-negative least-squares algorithm similar to what was used to analyze the FWGMR in Ch. 4.4 was invoked to set an upper bound on g_{seambump} [112]. A similar set of flip-chip devices for tantalum was also measured and compared with the rhenium devices (a Ta-based “bumped” FCLC device was not made). The resulting loss factors for the Ta-based devices and Re-based devices are shown in Table 6.3.

Table 6.3: Flip-Chip Loss Factors Measured at $\bar{n} \approx 1$.

Materials System	Γ_{SAMS} ($\times 10^{-4}$)	Γ_{MA} ($\times 10^{-4}$)	Γ_{bulk} ($\times 10^{-8}$)	$1/g_{\text{seambump}}$ (Ωm) ($\times 10^{-12}$)	$1/g_{\text{seampkg}}$ (Ωm) ($\times 10^{-3}$)	Γ_{mech}
Rhenium	3.80 ± 0.15	19.8 ± 2.5	5.38 ± 0.46	< 3.4	1.78 ± 0.37	38.0 ± 4.8
Tantalum	3.66 ± 0.71	22.2 ± 3.1	3.55 ± 1.36	-	2.59 ± 1.10	5.54 ± 1.35

Surprisingly, the surface loss factors Γ_{SAMS} and Γ_{MA} for the two materials systems are very similar; this is despite the additional contamination introduced during the fabrication of the Re-based device. It is very possible that if this contamination did not exist, the surface loss factors for Re would be even lower. Additionally, the Re/In-In/Re bump stackup is remarkably low loss, with a seam conductance of lower bound $g_{\text{seambump}} > 2.9 \times 10^{11} (\Omega\text{m})^{-1}$. This is similar to the measurement of $g_{\text{seamTa/Al}} \approx 3.9 \times 10^{11}$ shown in Ch. 5.1.3, which indicates that we can make indium bump-bonds as well as we can deposit two films on top of each other in vacuum. Indium bump loss for an In/In-In/In stackup has been measured in Lei et al. [82]; this measurement is also a bound at $> 2 \times 10^{10}$. However,

that bound was determined by using the mode Q and invoking Eq. (4.3), whereas the bound determined here used a multidimensional loss analysis that accounts for the other sources of loss present in the device. Finally, the measurement of Γ_{mech} is abnormally high in both the tantalum and rhenium-based devices. If we assume the mechanical support loss is conductor loss, we can estimate a surface resistance of $R_s \approx 7.5 \times 10^{-2} \Omega$ in the Rhenium device and $R_s \approx 1.1 \times 10^{-2} \Omega$ in the tantalum device. These surface resistances are alarmingly high, and may indicate that the use of a titanium hardstop is ill-advised for flip-chip circuits. However, if the device contamination issue is resolved and a different material (perhaps aluminum) is used for the hardstop instead of titanium, the Rhenium device may have dramatically improved surface loss that would make it an ideal candidate for flip-chip circuits. On the other hand, if reducing surface contamination does not change the loss factor, it might indicate that the source of MA loss may not be due to the presence or absence of a native oxide.

6.2.4 Rhenium-Based Transmons with Long Relaxation Times

Planar and flip-chip transmons were also fabricated alongside the Re-based loss-characterization devices. These devices had the same level of indium redeposition during fabrication. However, the single planar Re-based transmon (identical design as the Ta-based transmon in Fig. 5.1, except the Ta is replaced with Re) that was measured showed the longest relaxation times of any Ta-based transmon that was measured. This particular device has an average T_1 of around 350 μs with a maximum $T_1 \approx 410 \mu\text{s}$ (Fig. 6.4a). While this relaxation time is still within two standard deviations of what is to be expected from the loss model, it highlights that Rhenium seems to be a promising material that, provided the fabrication issues are resolved, can be an even lower-loss materials platform than tantalum, enabled possibly by its superior lattice matching with sapphire and potentially “cleaner” MA interface.

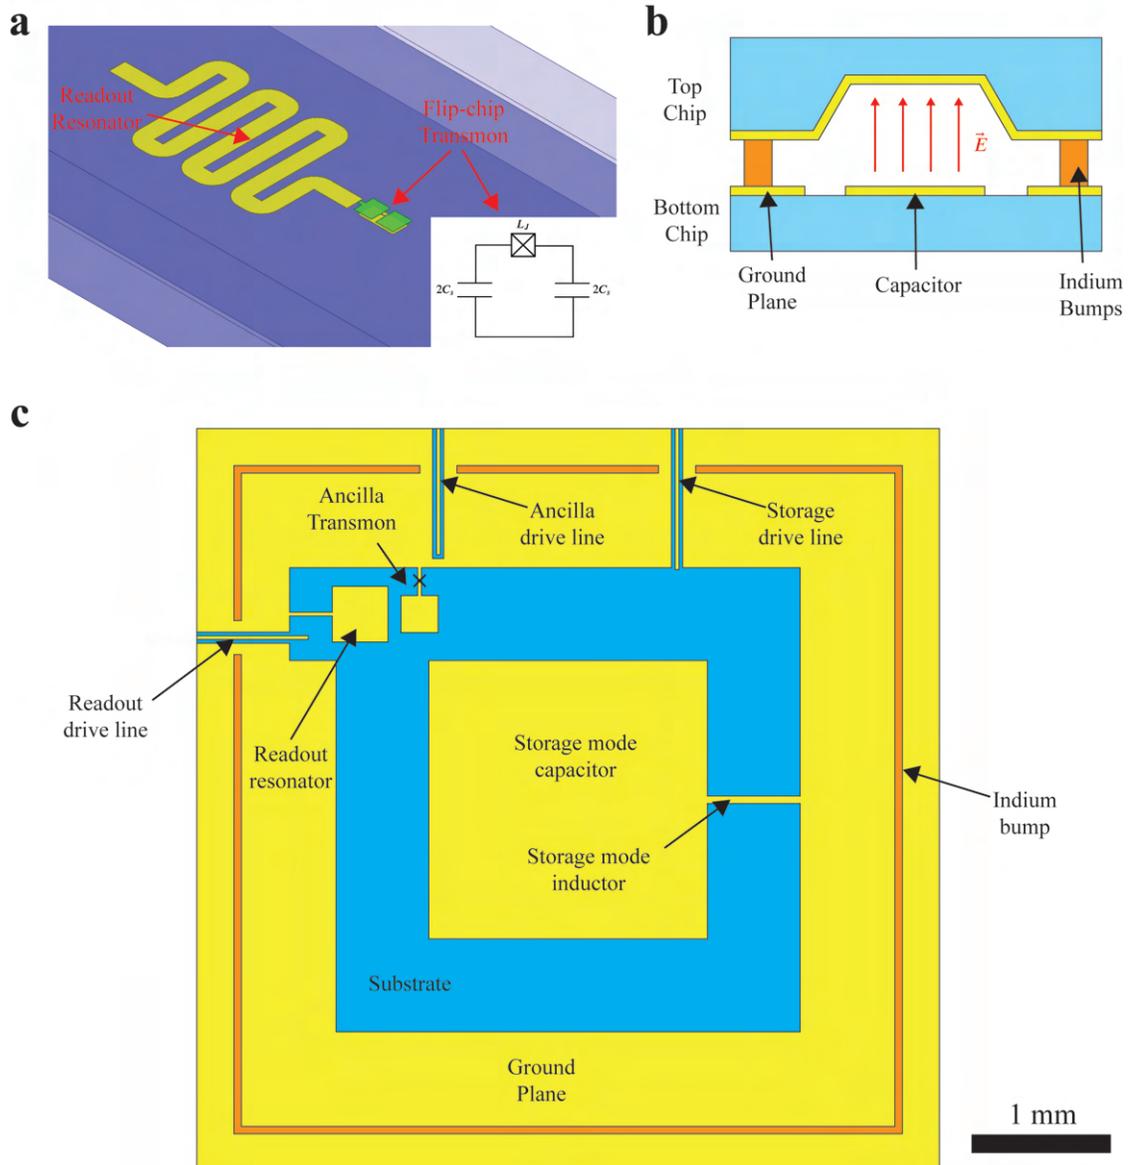


Figure 6.3: Flip-chip transmon and quantum memory. **a** Flip-chip transmon measured in the coaxial tunnel architecture. The green circuit is patterned on the top chip and the yellow circuit is patterned on the bottom chip. The transmon is of a “bumpless” design, with two capacitors ($190\ \mu\text{m} \times 190\ \mu\text{m}$ each) in series with equivalent capacitance C_s that are connected to a Josephson junction. A small antenna extends out from one of the capacitors to increase coupling to the readout resonator. **b** Cross-sectional diagram of a high-coherence flip-chip quantum memory. The electric field is localized in the vacuum gap, made bigger by the micromachined recess in the top chip, which is a silicon substrate that is fully metallized. The metallized top chip is connected to the ground plane on the bottom chip with indium bumps. **c** Top view of a quantum memory module in a planar architecture. Top chip is not shown for visibility; it would look identical to the top chip in **b**. Resonators and transmons utilize parallel-plate capacitances to the ground plane on the top chip.

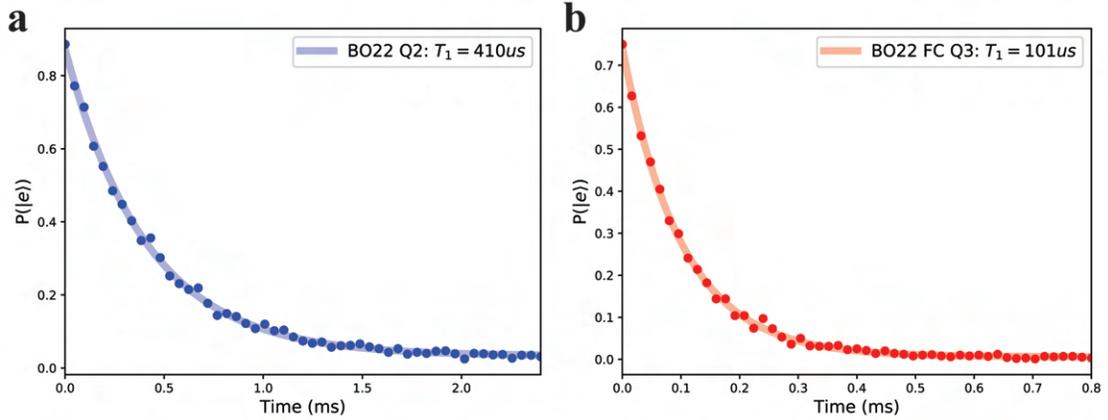


Figure 6.4: **Rhenium transmon** T_1 . **a** Re-based planar transmon. **b** Re-based flip-chip transmon. The reduced T_1 is due to additional p_{MA} , which allows the device to be ten times smaller than the planar transmon at the cost of four times reduced T_1 .

Additionally, flip-chip rhenium-based transmons were measured to have relaxation times of $T_1 \approx 100 \mu\text{s}$ (Fig. 6.4b). These devices use a parallel-plate capacitor to shunt the junction, resulting in a more compact device (see Fig. 6.3a). While this relaxation time is far lower than that of a planar Ta- or Re-based transmon, the flip-chip transmon is over ten times smaller than the planar transmon and has a far more localized electromagnetic field. Since the capacitor is a parallel plate, the flip-chip transmon has much higher MA participation $p_{MA} \sim 10^{-4}$ and around half of the bulk dielectric participation $p_{\text{bulk}} \approx 0.48$. However, the fringing fields of the capacitor still gives rise to significant MS and SA participation. This, in addition to MA loss results in the reduced T_1 . However, if rhenium's Γ_{MA} can be reduced further through improved processing, the flip-chip transmon's coherence can also be significantly improved.

6.3 A Flip-Chip Architecture for Bosonic Qubits

The idea of localizing electric fields in vacuum comes naturally with the flip-chip architecture and resembles the design principles of 3D cavity architectures, where high vacuum

participation minimizes surface participation, resulting in lower loss. With Ta and Re-based processes which consistently seem to have substantially lower surface losses than aluminum, it may be possible to obtain high coherence using flip-chip circuits in a planar architecture in a compact and lithographically scalable way. In such an architecture, on-chip quantum memories can be realized using parallel-plate capacitors to localize the electric field in the vacuum gap. While the FCLC devices measured in the previous section do not seem to have significantly high Q , the dramatic reduction of p_{SAMS} in these devices motivates an optimization of flip-chip device geometry to minimize SAMS loss and MA loss. This optimization comes in the form of increasing the gap between the chips, which can be easily done with micromachining into silicon substrates, as was demonstrated in Lei et al. [82]. By increasing the gap to $100\ \mu\text{m}$, the MA participation can be reduced to $\approx 4 \times 10^{-6}$ while maintaining $p_{\text{SAMS}} \approx 2.5 \times 10^{-5}$ (comparable to the surface participation of the hairpin stripline), which can enable compact flip-chip quantum memories with millisecond coherence times in a $3\ \text{mm} \times 3\ \text{mm}$ footprint (see Fig. 6.3b).

A flip-chip quantum memory can be made to be almost three times smaller than a hairpin stripline, but more importantly, it can be straightforwardly implemented in a planar architecture, as shown in Fig. 6.3c. As described in Ch. 2.5.1, the planar architecture benefits from the inclusion of an on-chip ground plane, which enables much more versatility in the engineering of microwave control lines. In particular, AC flux modulation and other fast-flux control “knobs” are straightforward to implement, which adds an additional degree of control with regards to engineering driven quantum circuit Hamiltonians. Additionally, an on-chip ground plane provides more options for complex on-chip filtering such as stepped impedance filters and lumped-element filters. If properly implemented, such an architecture could pose as a promising platform to implement large-scale quantum processors with bosonic qubits.

6.4 Summary

In this thesis I have given an overview of the work I have done towards microwave loss characterization. In Ch. 3 I discussed the various sources of energy loss in superconducting quantum circuits. In Ch. 4 I described several techniques to characterize and understand energy relaxation mechanisms in microwave resonators, with special attention to a modern, matrix-based approach using multimode resonators. Through these techniques, we can quantitatively answer important questions within the community about where dominant losses originate, and how the choices of materials, fabrication processes, and device geometry can play a role in determining a device's susceptibility to energy loss.

In introducing the FWGMR, I have shown how multimode resonators can serve as useful devices for learning how to improve coherence. In particular, I showed how we can improve surface dielectric and conductor quality in bulk-machined aluminum superconducting cavities by utilizing higher purity stock and surface processing techniques such as chemical etching, diamond-turning, and thin-film coating. The insights gained from this project led to the development of thin-film multimode resonators to characterize losses in on-chip circuits. Here, I showed that depending on resonator geometry, not only surface losses but bulk and package losses can be significant contributors to the total internal loss of a thin-film resonator. I also demonstrated how materials like tantalum and fabrication processes like high-temperature substrate annealing can improve coherence by reducing surface and bulk loss factors, respectively. Through this, I revealed in Ch. 5 how we gain a comprehensive understanding of what limits transmon relaxation, which motivated the design of the hairpin stripline, an optimized resonator that can be used as a quantum memory with coherence times exceeding one millisecond. This advancement in turn enables a compact, scalable, and modular coaxial architecture towards realizing large-scale quantum computing with qubits encoded in the bosonic modes of quantum memories.

The techniques presented in this thesis can also be straightforwardly extended towards understanding losses in other superconducting qubit architectures. While the results presented in this thesis are specific to the material and fabrication processes utilized in these works, the participation ratio model together with matrix loss characterization serves as a powerful and versatile approach towards understanding losses in co-planar waveguide (CPW) and flip-chip architectures, as was detailed in Sec. 6.2.3. One needs only to identify the limiting sources of loss and designing the appropriate set of resonators or resonance modes to characterize them. If this can be done in a systematic and comprehensive way, I expect it would lead to a dramatically improved understanding of relaxation-limiting mechanisms in superconducting quantum circuits and would hopefully lead to the development of large-scale quantum processors with low enough error rates to enable useful computation.

Bibliography

- [1] Feynman, R.P.: Simulating physics with computers. *International Journal of Theoretical Physics* **21**(6–7), 467–488 (1982) 1, 12
- [2] Deutsch, D.: Quantum-theory, the Church-Turing principle and the universal quantum computer. *Proceedings of the royal society of London series A - mathematical physical and engineering sciences* **400**(1818), 97–117 (1985) 1, 6, 13
- [3] Shor, P.W.: Algorithms for quantum computation: discrete logarithms and factoring. *Proceedings 35th Annual Symposium on Foundations of Computer Science*, 124–134 (1994) 2
- [4] Rivest, R.L., Shamir, A., Adleman, L.: A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM* **21**(2), 120–126 (1978) 2
- [5] Bennet, C.H., Brassard, G.: Quantum cryptography: Public key distribution and coin tossing. *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing, Bangalore*, 175–179 (1984) 2, 5
- [6] Bennet, C.H., Brassard, G., Mermin, N.D.: Quantum cryptography without Bell’s theorem. *Physical Review Letters* **68**, 557 (1992) 2
- [7] Grover, L.P.: Quantum mechanics helps in searching for a needle in a haystack. *Physical Review Letters* **79**(2), 325–328 (1997) 2
- [8] Nielsen, M.A., Chuang, I.L.: *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge CB2 8BS, United Kingdom (2016) 3, 5, 12
- [9] DiVincenzo, D.P.: The physical implementation of quantum computation. *Fortschritte der Physik* **48**(9-11), 771–783 (2000) 6
- [10] Monroe, C., Meekhof, D.M., King, B.E., Itano, W.M., Wineland, D.J.: Demonstration of a fundamental quantum logic gate. *Physical Review Letters* **75**(25), 4714–4717 (1995) 7

- [11] Gershenfeld, N.A., Chuang, I.L.: Bulk spin-resonance quantum computation. *Science* **275**(5298), 350–356 (1997) 7
- [12] Cory, D.G., Fahmy, A.F., Havel, T.F.: Ensemble quantum computing by NMR spectroscopy. *Proceedings of the National Academy of Sciences of the United States of America* **94**(5), 1634–1639 (1997) 7
- [13] Wilk, T., Gaëtan, A., Evelin, C., Wolters, J., Miroshnychenko, Y., Grangier, P., Browaeys, A.: Entanglement of two individual neutral atoms using Rydberg blockade. *Physical Review Letters* **104**, 010502 (2010) 7
- [14] Isenhower, L., Urban, E., Zhang, X.L., Gill, A.T., Henage, T., Johnson, T.A., Walker, T.G., Saffman, M.: Demonstration of a neutral atom controlled-not quantum gate. *Physical Review Letters* **104**, 010503 (2010) 7
- [15] Loss, D., DiVincenzo, D.P.: Quantum computation with quantum dots. *Physical Review A* **57**(1), 120–126 (1997) 7
- [16] Martinis, J.M., Nam, S., Aumentado, J., Urbina, C.: Rabi oscillations in a large Josephson-junction qubit. *Physical Review Letters* **89**(11), 117901 (2002) 7
- [17] Wallraff, A., Schuster, D.I., Blais, A., Huang, R.-S., Majer, J., Kumar, S., Girvin, S.M., Schoelkopf, R.J.: Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics. *Nature* **431**, 162–167 (2004)
- [18] Devoret, M.H., Wallraff, A., Martinis, J.M.: Superconducting qubits: A short review. *arXiv preprint arXiv:cond-mat/0411174* (2004) 7
- [19] Blais, A., Huang, R.-S., Wallraff, A., Girvin, S.M., Schoelkopf, R.J.: Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation. *Physical Review A* **69**, 062320 (2004) 7, 35, 38
- [20] Arute, F., Arya, K., Babbush, R., Bacon, D., Bardin, J.C., Barends, R., Biswas, R., Boixo, S., Brandao, F.G., Buell, D.A., *et al.*: Quantum supremacy using a programmable superconducting processor. *Nature* **574**(7779), 505–510 (2019) 8, 52, 53, 54
- [21] Kim, Y., Eddins, A., Anand, S., Wei, K.X., Berg, E., Rosenblatt, S., Nayfeh, H., Wu, Y., Zaletel, M., Temme, K., Kandala, A.: Evidence for the utility of quantum computing before fault tolerance. *Nature* **618**, 500–505 (2023) 8, 52
- [22] Kunihiro, N.: Quantum factoring algorithm: Resource estimation and survey of experiments. *International Symposium on Mathematics, Quantum Theory, and Cryptography* (2020) 10
- [23] Kjaergaard, M., Schwartz, M.E., Braumüller, J., Krantz, P., Wang, J.I.-J., Gustavsson, S., Oliver, W.D.: Superconducting qubits: Current state of play. *Annual Review of Condensed Matter Physics* **11**, 369–395 (2020) 11

- [24] Somoroff, A., Ficheux, Q., Mencia, R.A., Xiong, H., Kuzmin, R., Manucharyan, V.E.: Millisecond coherence in a superconducting qubit. *Physical Review Letters* **130**, 267001 (2023) 11
- [25] Lu, Y., Maiti, A., Garmon, J.W.O., Ganjam, S., Zhang, Y., Claes, J., Frunzio, L., Girvin, S.M., Schoelkopf, R.J.: High-fidelity parametric beamsplitting with a parity-protected converter. *Nature Communications* **14**(5767) (2023) 11, 56, 100
- [26] Shor, P.W.: Scheme for reducing decoherence in quantum computer memory. *Physical Review A* **52**, 2493 (1995) 11
- [27] Steane, A.: Multiple-particle interference and quantum error correction. *Proceedings of the Royal Society A* **452**, 1954 (1996)
- [28] Gottesman, D.: Class of quantum error-correcting codes saturating the quantum Hamming bound. *Physical Review A* **54**, 1862 (1996) 11
- [29] Sivak, V.V., Eickbusch, A., Royer, B., Singh, S., Tsioutsios, I., Ganjam, S., Miano, A., Brock, B.L., Ding, A.Z., Frunzio, L., Girvin, S.M., Schoelkopf, R.J., Devoret, M.H.: Real-time quantum error correction beyond break-even. *Nature* **616**, 50–55 (2023) 11, 43, 56
- [30] Girvin, S.M.: Circuit QED: Superconducting qubits coupled to microwave photons. *Les Houches* 15, 75
- [31] Reagor, M.J.: Superconducting cavities for circuit quantum electrodynamics. PhD thesis (2015) 15, 32, 44, 55, 76, 77, 79, 89, 96
- [32] Ambegaokar, v., Baratoff, A.: Tunneling between superconductors. *Physical Review Letters* **10**, 486 (1963) 19
- [33] Tinkham, M.: *Introduction to Superconductivity*. Dover Publications, Mineola, NY, USA (2015) 19, 75, 83, 167
- [34] Gloos, K., Poikolainen, R.S., Pekola, J.P.: Wide-range thermometer based on the temperature-dependent conductance of planar tunnel junctions. *Applied Physics Letters* **77**, 2915–2917 (2000) 20
- [35] Serniak, K.: Nonequilibrium quasiparticles in superconducting qubits. PhD thesis (2019) 20, 167
- [36] Koch, J., Yu, T.M., Gambetta, J., Houck, A.A., Schuster, D.I., Majer, J., Blais, A., Devoret, M.H., Girvin, S.M., Schoelkopf, R.J.: Charge-insensitive qubit design derived from the cooper pair box. *Physical Review A* **76**(4), 042319 (2007) 21, 22, 23, 24, 245
- [37] Schuster, D.I.: Circuit quantum electrodynamics. PhD thesis (2007) 22, 36, 59

- [38] Kim, Z., Suri, B., Zaretsky, V., Novikov, S., Osborn, K.D., Mizel, A., Wellstood, F.C., Palmer, B.S.: Decoupling a cooper-pair box to enhance the lifetime to 0.2 ms. *Physical Review Letters* **106**, 120501 (2011) 22, 89, 176
- [39] Wang, C., Li, X., Xu, H., Li, Z., Wang, J., Yang, Z., Mi, Z., Liang, X., Su, T., Yang, C., Wang, G., Wang, W., Li, Y., Chen, M., Li, C., Linghu, K., Han, J., Zhang, Y., Feng, Y., Song, Y., Ma, T., Zhang, J., Wang, R., Zhao, P., Liu, W., Xue, G., Jin, Y., Yu, H.: Towards practical quantum computers: transmon qubit with a lifetime approaching 0.5 milliseconds. *npj Quantum Information* **8**(3) (2022) 23, 53, 54, 123, 184
- [40] Blais, A., Grimsmo, A., Girvin, S.M., Wallraff, A.: Circuit quantum electrodynamics. *Reviews of Modern Physics* **93**, 025005 (2021) 24, 34, 35, 36, 39, 59
- [41] Reinhold, P.: Controlling error-correctable bosonic qubits. PhD thesis (2019) 27
- [42] Krantz, P., Kjaergaard, M., Yan, F., Orlando, T.P., Gustavsson, S., Oliver, W.D.: A quantum engineer's guide to superconducting qubits. *Applied Physics Reviews* **6**, 021318 (2019) 28, 33, 34, 35, 38, 59, 63, 116, 117
- [43] Motzoi, F., Gambetta, J.M., Reberstrost, P., Wilhelm, F.K.: Simple pulses for elimination of leakage in weakly nonlinear qubits. *Physical Review Letters* **103**, 110501 (2009) 30
- [44] Gambetta, J.M., Motzoi, F., Merkel, S.T., Wilhelm, F.K.: Analytic control methods for high-fidelity unitary operations in a weakly nonlinear oscillator. *Physical Review A* **83**, 012308 (2011) 30
- [45] Schuster, D.I., Wallraff, A., Blais, A., Frunzio, L., Huang, R.-S., Majer, J., Girvin, S.M., Schoelkopf, R.J.: AC Stark shift and dephasing of a superconducting qubit strongly coupled to a cavity field. *Physical Review Letters* **94**, 123602 (2005) 31
- [46] Devoret, M.H., Girvin, S.M., Schoelkopf, R.J.: Circuit-QED: How strong can the coupling between a Josephson junction atom and a transmission line resonator be? *Annals of Physics* **16**(10–11), 767–779 (2007) 34
- [47] Haroche, S., Raimond, J.-M.: *Exploring the Quantum: Atoms, Cavities, and Photons*. Oxford University Press, New York, USA (2006) 35
- [48] Lamb Jr., W.E., Retherford, R.C.: Fine structure of the hydrogen atom by a microwave method. *Physical Review Journals Archive* **72**, 241 (1947) 37
- [49] Vijay, R., Slichter, D.H., Siddiqi, I.: Observation of quantum jumps in a superconducting artificial atom. *Physical Review Letters* **106**, 110502 (2011) 41
- [50] Hatridge, M., Shankar, S., Mirrahimi, M., Shackert, F., Geerlings, K., Brecht, T., Silwa, K.M., Abdo, B., Frunzio, L., Girvin, S.M., Schoelkopf, R.J., Devoret, M.H.: Quantum back-action of an individual variable-strength measurement. *Science* **339**(6116), 178–181 (2013) 41

- [51] Frattini, N.E., Sivak, V.V., Lingenfelter, A., Shankar, S., Devoret, M.H.: Optimizing the nonlinearity and dissipation of a SNAIL parametric amplifier for dynamic range. *Physical Review Applied* **10**, 054020 (2018) 41
- [52] Gambetta, J., Blais, A., Boissonneault, M., Houck, A.A., Schuster, D.I., Girvin, S.M.: Quantum trajectory approach to circuit QED: Quantum jumps and the Zeno effect. *Physical Review A* **77**, 012112 (2008) 42
- [53] Reagor, M., Pfaff, W., Axline, C., Heeres, R.W., Ofek, N., Sliwa, K., Holland, E., Wang, C., Blumoff, J., Chou, K., Hatridge, M.J., Frunzio, L., Devoret, M.H., Jiang, L., Schoelkopf, R.J.: Quantum memory with millisecond coherence in circuit QED. *Phys. Rev. B* **94**, 014506 (2016) 42, 45, 55, 74, 123, 140, 145, 186, 190, 245
- [54] Ofek, N., Petrenko, A., Heeres, R., Reinhold, P., Leghtas, Z., Vlastakis, B., Liu, Y., Frunzio, L., Girvin, S.M., Jiang, L., Mirrahimi, M., Devoret, M.H., Schoelkopf, R.J.: Extending the lifetime of a quantum bit with error correction in superconducting circuits. *Nature* **536**, 441–445 (2016) 43
- [55] Heeres, R.W., Vlastakis, B., Holland, E., Krastanov, S., Albert, V.V., Frunzio, L., Jiang, L., Schoelkopf, R.J.: Cavity state manipulation using photon-number selective phase gates. *Physical Review Letters* **115**, 137002 (2015) 43, 45
- [56] Leek, P.J., Fink, J.M., Blais, A., Bianchetti, R., Göppl, M., Gambetta, J.M., Schuster, D.I., Frunzio, L., Schoelkopf, R.J., Wallraff, A.: Observation of Berry’s phase in a solid-state qubit. *Science* **339**(6116), 178–181 (2013) 44
- [57] Krastanov, S., Albert, V.V., Shen, C., Zou, C.-L., Heeres, R.W., Vlastakis, B., Schoelkopf, R.J., Jiang, L.: Universal control of an oscillator with dispersive coupling to a qubit. *Physical Review A* **92**, 040303 (2015) 45
- [58] Khaneja, N., Reiss, T., Kehlet, C., Schulte-Herbrüggen, T., Glaser, S.J.: Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms. *Journal of Magnetic Resonance* **172**, 296–305 (2005) 46
- [59] Pozar, D.M.: *Microwave Engineering*. John Wiley & Sons, Inc., Hoboken, NJ, USA (2012) 46, 68, 76, 77
- [60] Nigg, S.E., Paik, H., Vlastakis, B., Kirchmair, G., Shankar, S., Frunzio, L., Devoret, M.H., Schoelkopf, S.M. R. J. Girvin: Black-box superconducting circuit quantization. *Physical Review Letters* **108**, 240502 (2012) 47, 69
- [61] Mineev, Z.K., Leghtas, Z., Mundhada, S.O., Christakis, L., Pop, I.M., Devoret, M.H.: Energy-participation quantization of Josephson circuits. *npj quantum information* **7**(131) (2021) 48
- [62] Schuster, D.I., Houck, A.A., Schreier, J.A., Wallraff, A., Gambetta, J.M., Blais, A., Frunzio, L., Majer, J., Johnson, B., Devoret, M.H., Girvin, S.M., Schoelkopf, R.J.:

- Resolving photon number states in a superconducting circuit. *Nature* **445**, 515–518 (2007) 51, 245
- [63] Barends, R., Kelly, J., Megrant, A., Sank, D., Jeffrey, E., Chen, Y., Yin, Y., Chiaro, B., Mutus, J., Neill, C., O’Malley, P., Roushan, P., Wenner, J., White, T.C., Cleland, A.N., M., M.J.: Coherent josephson qubit suitable for scalable quantum integrated circuits. *Physical Review Letters* **111**, 080502 (2013) 51, 53, 54, 245
- [64] Place, A.P.M., Rodgers, L.V.H., Mundada, P., Smitham, B.M., Fitzpatrick, M., Leng, Z., Premkumar, A., Bryon, j., Sussman, S., Cheng, G., Madhavan, T., Babla, H.K., Jaeck, B., Gyenis, A., Yao, N., Cava, R.J., Leon, N.P., Houck, A.A.: New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds. *Nature Communications* **12**(1779) (2021) 51, 53, 54, 123, 184, 227, 231, 245
- [65] Andresen, S.: QCage.24 User Manual. QDevil ApS, (2022). QDevil ApS 51, 245
- [66] Huang, S., Lienhard, B., Calusine, G., Vepsäläinen, A., Braumüller, J., Kim, D.K., Melville, A.J., Niedzielski, B.M., Yoder, J.L., Kannan, B., Orlando, T.P., Gustavsson, S., Oliver, W.D.: Microwave package design for superconducting quantum processors. *PRX Quantum* **2**, 020306 (2021) 52
- [67] O’Brien, W., Vahidpour, M., Whyland, J.T., Angeles, J., Marshall, J., Scarabelli, D., Crossman, G., Yadav, K., Mohan, Y., Bui, C., Rawat, V., Renzas, R., Vodrahalli, N., Bestwick, A., Rigetti, C.: Superconducting caps for quantum integrated circuits. arXiv preprint arXiv:1708.02219 (2017)
- [68] Foxen, B., Mutus, J.Y., Lucero, E., Graff, R., Megrant, A., Chen, Y., Quintana, C., Burkett, B., Kelly, J., Jeffrey, E., Yang, Y., Yu, A., Arya, K., Barends, R., Chen, Z., Chiaro, B., Dunsworth, A., Fowler, A., Gidney, C., Giustina, M., Huang, T., Klimov, P., Neeley, M., Neill, C., Roushan, P., Sank, D., Vainsencher, A., Wenner, J., White, T.C., Martinis, J.M.: Qubit compatible superconducting interconnects. *Quantum Science and Technology* **3**(1), 014005 (2017)
- [69] Rosenberg, D., Kim, D., Das, R., Yost, D., Gustavsson, S., Hover, D., Krantz, P., Melville, A., Racz, L., Samach, G.O., Weber, S.J., Yan, F., Yoder, J.L., Kerman, A.J., Oliver, W.D.: 3D integrated superconducting qubits. *npj Quantum Information* **3**(1), 42 (2017) 52
- [70] Houck, A.A., Schreier, J.A., Johnson, B.R., Chow, J.M., Koch, J., Gambetta, J.M., Schuster, D.I., Frunzio, L., Devoret, M.H., Girvin, S.M., Schoelkopf, R.J.: Controlling the spontaneous emission of a superconducting transmon qubit. *Physical Review Letters* **101**, 080502 (2008) 53, 61, 71, 123
- [71] Gordon, R.T., Murray, C.E., Kurter, C., Sandberg, M., Hall, S.A., Balakrishnan, K., Shelby, R., Wacaser, B., Stabile, A.A., Sleight, J.W., Brink, M., Rothwell,

- M.B., Rodbell, K.P., Dial, O., Steffen, M.: Environmental radiation impact on lifetimes and quasiparticle tunneling rates of fixed-frequency transmon qubits. *Applied Physics Letters* **120**, 074002 (2022) 53, 54
- [72] Houck, A.A., Koch, J., Devoret, M.H., Girvin, S.M., Schoelkopf, R.J.: Life after charge noise: recent results with transmon qubits. *Quantum Information Processing* **8**, 105–115 (2009) 54
- [73] Paik, H., Schuster, D.I., Bishop, L.S., Kirchmair, G., Catelani, G., Sears, A.P., Johnson, B.R., Reagor, M.J., Frunzio, L., Glazman, L.I., Girvin, S.M., Devoret, M.H., Schoelkopf, R.J.: Observation of high coherence in josephson junction qubits measured in a three-dimensional circuit QED architecture. *Physical Review Letters* **107**, 240501 (2011) 54, 55, 57, 123, 171
- [74] Wang, C., Axline, C., Gao, Y.Y., Brecht, T., Chu, Y., Frunzio, L., Devoret, M.H., Schoelkopf, R.J.: Surface participation and dielectric loss in superconducting qubits. *Applied Physics Letters* **107**, 162601 (2015) 54, 88, 127, 130, 175, 242
- [75] Chou, K.S., Blumoff, J.Z., Wang, C.S., Reinhold, P.C., Axline, C.J., Gao, Y.Y., Frunzio, L., Devoret, M., Jiang, L., Schoelkopf, R.J.: Deterministic teleportation of a quantum gate between two logical qubits. *Nature* **561**(7723), 368 (2018) 55, 56, 145, 245
- [76] Milul, O., Guttel, B., Goldblatt, U., Hazanov, S., Joshi, L.M., Chausovsky, D., Kahn, N., Çiftyürek, E., Lafont, F., Rosenblum, S.: A superconducting quantum memory with tens of milliseconds coherence time. *arXiv preprint arXiv:2302.06442* (2023) 55, 123, 186, 192
- [77] Romanenko, A., Pilipenko, R., Zorzetti, S., Frolov, D., Awida, M., Belomestnykh, S., Posen, S., Grassellino, A.: Three-dimensional superconducting resonators at $T < 20$ mk with photon lifetimes up to $\tau = 2$ s. *Physical Review Applied* **13**, 034032 (2020) 55, 76, 144
- [78] Gargiulo, S. O. ans Oleschko, Prat-Camps, J., Zanner, M., Kirchmair, G.: Fast flux control of 3D transmon qubits using a magnetic hose. *Applied Physics Letters* **118**, 012601 (2021) 56
- [79] Chapman, B.J., Graaf, S.J., Xue, S.H., Zhang, Y., Teoh, J., Curtis, J.C., Tsunoda, T., Eickbusch, A., Read, A.P., Koottandavida, A., Mundhada, S.O., Frunzio, L., Devoret, M.H., Girvin, S.M., Schoelkopf, R.J.: High-on-off-ratio beam-splitter interaction for gates on bosonically encoded qubits. *PRX Quantum* **4**, 020355 (2023) 56
- [80] Chakram, S., Oriani, A.E., Naik, R.K., Dixit, A.V., He, K., Agrawal, A., Kwon, H., Schuster, D.I.: Seamless high-Q microwave cavities for multimode circuit quantum electrodynamics. *Physical review letters* **127**(10), 107701 (2021) 56, 123, 186

- [81] Zhou, C., Lu, P., Praquin, M., Chien, T.-C., Kaufman, R., Cao, X., Xia, M., Mong, R., Pfaff, W., Pekker, D., Hatridge, M.: Realizing all-to-all couplings among detachable quantum modules using a microwave quantum state router. *npj quantum information* **9**(54) (2023) 56
- [82] Lei, C.U., Krayzman, L., Ganjam, S., Frunzio, L., Schoelkopf, R.J.: High coherence superconducting microwave cavities with indium bump bonding. *Applied Physics Letters* **116**, 154002 (2020) 57, 81, 93, 94, 95, 97, 124, 127, 142, 176, 200, 204, 205, 209, 234
- [83] Krayzman, L.: Thin-film 3D resonators for superconducting quantum circuits. PhD thesis, Yale University (2023) 57, 76, 79, 80, 144, 176
- [84] Axline, C., Reagor, M., Heeres, R., Reinhold, P., Wang, C., Shain, K., Pfaff, W., Chu, Y., Frunzio, L., Schoelkopf, R.J.: An architecture for integrating planar and 3D cQED devices. *Applied Physics Letters* **109**(4), 042601 (2016) 57, 58, 59, 70, 72, 101, 126, 169, 187, 245
- [85] Mineev, Z.K., Serniak, K., Pop, I.M., Leghtas, Z., Sliwa, K., Hatridge, M., Frunzio, L., Schoelkopf, R.J., Devoret, M.H.: Planar multilayer circuit quantum electrodynamics. *Physical Review Applied* **5**(4), 044021 (2016) 59, 187
- [86] Ganjam, S., Wang, Y., Lu, Y., Banerjee, A., Lei, C.U., Krayzman, L., Kisslinger, K., Zhou, C., Li, R., Jia, Y., Liu, M., Frunzio, L., Schoelkopf, R.J.: Surpassing millisecond coherence times in on-chip superconducting quantum memories by optimizing materials, processes, and circuit design. arXiv preprint arXiv:2308.15539 (2023) 59, 85, 95, 102, 104, 114, 132, 147, 148, 149, 150, 151, 155, 157, 162, 164, 171, 172, 175, 177, 179, 180, 182, 183, 185, 188, 189, 190, 191, 192, 235, 236, 237, 245
- [87] Purcell, E.M., Torrey, H.C., V., P.R.: Resonance absorption by nuclear magnetic moments in a solid. *Physical Review Journals Archive* **69**, 37 (1946) 61, 71
- [88] Göppl, M., Fragner, A., Baur, M., Bianchetti, R., Filipp, S., Fink, J.M., Leek, P.J., Puebla, G., Steffen, L., Wallraff, A.: Coplanar waveguide resonators for circuit quantum electrodynamics. *Journal of Applied Physics* **104**, 113904 (2008) 70, 245
- [89] Geerlings, K., Shankar, S., Edwards, E., Frunzio, L., Schoelkopf, R.J., Devoret, M.H.: Improving the quality factor of microwave compact resonators by optimizing their geometrical parameters. *Applied Physics Letters* **100**, 192601 (2012) 70, 245
- [90] Reagor, M., Paik, H., Catelani, G., Sun, L., Axline, C., Holland, E., Pop, I.M., Masluk, N.A., Brecht, T., Frunzio, L., Devoret, M.H., Glazman, L., Schoelkopf, R.J.: Reaching 10 ms single photon lifetimes for superconducting aluminum cavities. *Applied Physics Letters* **102**(19), 192604 (2013) 70, 80, 81, 82, 101, 123, 125, 137, 140, 143, 245

- [91] Axline, C.J.: Building blocks for modular circuit QED quantum computing. PhD thesis (2018) 71, 77
- [92] Van Duzer, T., Turner, C.W.: Principles of Superconductive Devices and Circuits. Elsevier, New York, USA (1981) 75, 83
- [93] Gurevich, A.: Theory of rf superconductivity for resonant cavities. *Superconductor Science and Technology* **30**, 034004 (2017) 76, 78, 79, 80, 82, 83, 85, 143, 166, 167, 197
- [94] Zmuidzinas, J.: Superconducting microresonators: Physics and applications. *Annu. Rev. Condens. Matter Phys.* **3**(1), 169–214 (2012) 76, 78, 79
- [95] Gao, J.: The physics of superconducting microwave resonators. PhD thesis, California Institute of Technology (2008) 77, 78, 79, 80, 82, 83, 91, 92, 143, 167
- [96] Houzet, M., Serniak, K., Catelani, G., Devoret, M.H., Glazman, L.I.: Photon-assisted charge-parity jumps in a superconducting qubit. *Physical Review Letters* **123**, 107704 (2019) 83
- [97] Diamond, S., Fatemi, V., Hays, M., Nho, H., Kurilovich, P.D., Connolly, T., Joshi, V.R., Serniak, K., Frunzio, L., Glazman, L.I., Devoret, M.H.: Distinguishing parity-switching mechanisms in a superconducting qubit. *PRX Quantum* **3**, 040304 (2022) 123, 143, 166
- [98] Connolly, T., Kurilovich, P.D., Diamond, S., Nho, H., Böttcher, C.G.L., Glazman, L.I., Fatemi, V., Devoret, M.H.: Coexistence of nonequilibrium density and equilibrium energy distribution of quasiparticles in a superconducting qubit. *arXiv preprint arXiv:2302.12330* (2023) 83, 103, 123, 143, 166, 176, 195
- [99] McEwen, M., Faoro, L., Arya, K., Dunsworth, A., Huang, T., Kim, S., Burkett, B., Fowler, A., Arute, F., Bardin, J.C., Bengtsson, A., Bilmes, A., Buckley, B.B., Bushnell, N., Chen, Z., Collins, R., Demura, S., Derk, A.R., Erickson, C., Giustina, M., Harrington, S.D., Hong, S., Jeffrey, E., Kelly, J., Klimov, P.V., Kostritsa, F., Laptev, P., Locharla, A., Mi, X., Miao, K.C., Montazeri, S., Mutus, J., Naaman, O., Neeley, M., Neill, C., Opremcak, A., Quintana, C., Redd, N., Roushan, P., Sank, D., Satzinger, K.J., Shvarts, V., White, T., Yao, Z.J., Yeh, P., Yoo, J., Chen, Y., Smelyanskiy, V., Martinis, J.M., Neven, H., Megrant, A., Lev, I., Barends, R.: Resolving catastrophic error bursts from cosmic rays in large arrays of superconducting qubits. *Nature Physics* **18**, 107–111 (2022) 83
- [100] Thorbeck, T., Eddins, A., Lauer, I., McClure, D.T., Carroll, M.: Two-level-system dynamics in a superconducting qubit due to background ionizing radiation. *PRX Quantum* **4**, 020356 (2023) 83, 94, 95, 182
- [101] Catelani, G., Li, K., Axline, C.J., Brecht, T., Frunzio, L., Schoelkopf, R.J., Glazman, L.I.: Ac losses in field-cooled type I superconducting cavities. *Superconductor Science and Technology* **35**, 065016 (2022) 84, 85, 123

- [102] Wang, C., Gao, Y.Y., Pop, I.M., Vool, U., Axline, C., Brecht, T., Heeres, R.W., Frunzio, L., Devoret, M.H., Catelani, G., Glazman, L.I., Schoelkopf, R.J.: Measurement and control of quasiparticle dynamics in a superconducting qubit. *Nature Communications* **5**(5836) (2014) 84, 85
- [103] Creedon, D.L., Reshitnyk, Y., Farr, W., Martinis, J.M., Duty, T.L., Tobar, M.E.: High Q-factor sapphire whispering gallery mode microwave resonator at single photon energies and millikelvin temperatures. *Applied Physics Letters* **98**, 222903 (2011) 85, 90, 123, 161, 196
- [104] Chu, Y., Axline, C., Wang, C., Brecht, T., Gao, Y.Y., Frunzio, L., Schoelkopf, R.J.: Suspending superconducting qubits by silicon micromachining. *Applied Physics Letters* **109**(11), 112601 (2016) 127
- [105] Calusine, G., Melville, A., Woods, W., Das, R., Stull, C., Bolkhovsky, V., Braje, D., Hover, D., Kim, D.K., Miloshi, X., Rosenberg, D., Sevi, A., Yoder, J.L., Dauler, E.A., Oliver, W.D.: Analysis and mitigation of interface losses in trenched superconducting coplanar waveguide resonators. *Applied Physics Letters* **112**(6), 062601 (2018) 85, 88, 124, 129, 131, 152, 198
- [106] Read, A.P., Chapman, B.J., Lei, C.U., Curtis, J.C., Ganjam, S., Krayzman, L., Frunzio, L., Schoelkopf, R.J.: Precision measurement of the microwave dielectric loss of sapphire in the quantum regime with parts-per-billion sensitivity. *Physical Review Applied* **19**(3), 034064 (2023) 85, 90, 91, 106, 123, 131, 145, 149, 158, 161, 196, 227
- [107] Woods, W., Calusine, G., Melville, A., Sevi, A., Golden, E., Kim, D.K., Rosenberg, D., Yoder, J.L., Oliver, W.D.: Determining interface dielectric losses in superconducting coplanar-waveguide resonators. *Physical Review Applied* **12**(1), 014012 (2019) 85, 124, 129, 131, 152, 198
- [108] Melville, A., Calusine, G., Woods, W., Serniak, K., Golden, E., Niedzielski, B.M., Kim, D.K., Sevi, A., Yoder, J.L., Dauler, E.A., Oliver, W.D.: Comparison of dielectric loss in titanium nitride and aluminum superconducting resonators. *Applied Physics Letters* **117**(12), 124004 (2020) 129, 131, 152, 198
- [109] McRae, C.R.H., Wang, H., Gao, J., Vissers, M.R., Brecht, T., Dunsworth, A., Pappas, D.P., Mutus, J.: Materials loss measurements using superconducting microwave resonators. *Review of Scientific Instruments* **91**, 091101 (2020)
- [110] Crowley, K.D., McLellan, R.A., Dutta, A., Shumiya, N., Place, A.P.M., Le, X.H., Gang, Y., Madhavan, T., Bland, M.P., Chang, R., Khedkar, N., Feng, Y.C., Umbarkar, E.A., Gui, X., Rodgers, L.V.H., Jia, Y., Feldman, M.M., Lyon, S.A., Liu, M., Cava, R.J., Houck, A.A., Leon, N.P.: Disentangling losses in tantalum superconducting circuits. *Physical Review X* **13**, 041005 (2023) 85, 88, 92, 123, 124, 127, 158, 230, 236

- [111] Wenner, J., Barends, R., Bialczak, R., Chen, Y., Kelly, J., Lucero, E., Mariantoni, M., Megrant, A., O'Malley, P., Sank, D., Vainsencher, A., Wang, H., White, T.C., Yin, Y., Zhao, J., Cleland, A.N., Martinis, J.M.: Surface loss simulations of superconducting coplanar waveguide resonators. *Applied Physics Letters* **99**(11), 113513 (2011) 88
- [112] Lei, C.U., Ganjam, S., Krayzman, L., Banerjee, A., Kisslinger, K., Hwang, S., Frunzio, L., Schoelkopf, R.J.: Characterization of microwave loss using multimode superconducting resonators. *Physical Review Applied* **20**, 024045 (2023) 88, 101, 123, 124, 132, 133, 135, 139, 140, 141, 146, 153, 205, 246
- [113] Dunne, L.J., Axelsson, A.-K., Alford, N.M., Breeze, J., Aupi, x., Brandas, E.J.: Quasi-classical fluctuation–dissipation description of dielectric loss in oxides with implications for quantum information processing. *International Journal of Quantum Chemistry* **106**, 986–993 (2006) 90
- [114] Martinis, J.M., Cooper, K.B., McDermott, R., Steffen, M., Ansmann, M., Osborn, K.D., Cicak, K., Oh, S., Pappas, D.P., Simmonds, R.W., Yu, C.C.: Decoherence in josephson qubits from dielectric loss. *Physical Review Letters* **95**, 210503 (2005) 92, 94, 123, 175
- [115] Gao, J., Daal, M., Vayonakis, A., Kumar, S., Zmuidzinas, J., Sadoulet, B., Mazin, B.A., Day, P.K., Leduc, H.G.: Experimental evidence for a surface distribution of two-level systems in superconducting lithographed microwave resonators. *Applied Physics Letters* **92**, 152505 (2008) 94, 123, 158
- [116] Wang, H., Hofheinz, M., Wenner, J., Ansmann, M., Bialczak, R.C., Lenander, M., Lucero, E., Neeley, M., O'Connell, A.D., Sank, D., Weides, M., Cleland, A.N., Martinis, J.M.: Improving the coherence time of superconducting coplanar resonators. *Applied Physics Letters* **95**, 233508 (2009)
- [117] Pappas, D.P., Vissers, M.R., Wisbey, D.S., Kline, J.S., Gao, J.: Two level system loss in superconducting microwave resonators. *IEEE Transactions on Applied Superconductivity* **21**(3), 871–874 (2011) 94, 123
- [118] Burnett, J., Faoro, L., Wisby, I., Gurtovoi, V.L., Chernykh, A.V., Mikhailov, G.M., Tulin, V.A., Shaikhaidarov, R., Antonov, V., Meeson, P.J., Tzalenchuk, A.Y., Lindström, T.: Evidence for interacting two-level systems from the $1/f$ noise of a superconducting resonator. *Nature Communications* **5**(4119) (2014) 92
- [119] Bruno, A., De Lange, G., Asaad, S., Enden, K., Langford, N., DiCarlo, L.: Reducing intrinsic loss in superconducting resonators by surface treatment and deep etching of silicon substrates. *Applied Physics Letters* **106**(18), 182601 (2015) 94
- [120] Lisenfeld, J., Bilmes, A., Megrant, A., Barends, R., Kelly, J., Klimov, P., Weiss, G., Martinis, J.M., Ustinov, A.V.: Electric field spectroscopy of material defects in transmon qubits. *npj Quantum Information* **5**(105) (2019) 94

- [121] Klimov, P.V., Kelly, J., Chen, Z., Neeley, M., Megrant, A., Burkett, B., Barends, R., Arya, K., Chiaro, B., Chen, Y., Dunsworth, A., Fowler, A., Foxen, B., Gidney, C., Giustina, M., Graff, R., Huang, T., Jeffrey, E., Lucero, E., Mutus, J.Y., Naaman, O., Neill, C., Quintana, C., Roushan, P., Sank, D., Vainsencher, A., Wenner, J., White, T.C., Boixo, S., Babbush, R.: Fluctuations of energy-relaxation times in superconducting qubits. *Physical Review Letters* **121**, 090502 (2018) 94, 182
- [122] Brecht, T., Reagor, M., Chu, Y., Pfaff, W., Wang, C., Frunzio, L., Devoret, M.H., Schoelkopf, R.J.: Demonstration of superconducting micromachined cavities. *Applied Physics Letters* **107**(19), 192603 (2015) 95, 97, 124, 127, 145, 174
- [123] Brecht, T., Chu, Y., Axline, C., Pfaff, W., Blumoff, J.Z., Chou, K., Krayzman, L., Frunzio, L., Schoelkopf, R.J.: Micromachined integrated quantum circuit containing a superconducting qubit. *Physical Review Applied* **7**(4), 044018 (2017) 200
- [124] Brecht, T.: Micromachined quantum circuits. PhD thesis (2017) 95
- [125] Constantin, M., Yu, C.C., Martinis, J.M.: Saturation of two-level systems and charge noise in josephson junction qubits. *Physical Review B* **79**, 094520 (2009) 98
- [126] Yan, F., Gustavsson, S., Kamal, A., Birenbaum, J., Sears, A.P., Hover, D., Gudmundsen, T.J., Rosenberg, D., Samach, G., Weber, S., Yoder, J.L., Orlando, T.P., Clarke, J., Kerman, A.J., Oliver, W.D.: The flux qubit revisited to enhance coherence and reproducibility. *Nature Communications* **7**(12964) (2016) 99, 118
- [127] Koch, R.H., DiVincenzo, D.P., Clarke, J.: Model for $1/f$ flux noise in SQUIDs and qubits. *Physical Review Letters* **98**, 267003 (2007) 100
- [128] Braumüller, J., Ding, L., Vepsäläinen, A.P., Sung, Y., Kjaergaard, M., Menke, T., Winik, R., Kim, D., Niedzielski, B.M., Melville, A., Yoder, J.L., Hirjibehedin, C.F., Orlando, T.P., Gustavsson, S., Oliver, W.D.: Characterizing and optimizing qubit coherence based on SQUID geometry. *Physical Review Applied* **13**, 054079 (2020) 100
- [129] Leong, K., Mazierska, J.: Precise measurements of the Q factor of dielectric resonators in the transmission mode-accounting for noise, crosstalk, delay of uncalibrated lines, coupling loss, and coupling reactance. *IEEE Transactions on Microwave Theory and Techniques* **50**(9), 2115–2127 (2002) 108, 109
- [130] Probst, S., Song, F.B., Bushev, P.A., Ustinov, A.V., Weides, M.: Efficient and robust analysis of complex scattering data under noise in microwave resonators. *Review of Scientific Instruments* **86**(2), 024706 (2015) 108, 109, 110
- [131] Khalil, M.S., Stoutimore, M.J.A., Wellstood, F.C., Osborn, K.D.: An analysis method for asymmetric resonator transmission applied to superconducting devices. *Journal of Applied Physics* **111**, 054510 (2012) 110

- [132] Dial, O., McClure, D.T., Poletto, S., Keefe, G.A., Rothwell, M.B., Gambetta, J.M., Abraham, D.W., Chow, J.M., Steffen, M.: Bulk and surface loss in superconducting transmon qubits. *Superconductor Science and Technology* **29**, 044001 (2016) 123
- [133] Gambetta, J.M., Murray, C.E., Fung, Y.-K.-K., McClure, D.T., Dial, O., Shanks, W., Sleight, J.W., Steffen, M.: Investigating surface loss effects in superconducting transmon qubits. *IEEE Transactions on Applied Superconductivity* **27**(1), 1–5 (2016)
- [134] Serniak, K., Hays, M., Lange, G., Diamond, S., Shankar, S., Burkhart, L.D., Frunzio, L., Houzet, M., Devoret, M.H.: Hot nonequilibrium quasiparticles in transmon qubits. *Physical Review Letters* **121**, 157701 (2018) 123, 143, 166
- [135] Martinis, J.M., Megrant, A.: UCSB final report for the CSQ program: Review of decoherence and materials physics for superconducting qubits. arXiv preprint arXiv:1410.5793 (2014) 123
- [136] Deng, H., Song, Z., Gao, R., Xia, T., Bao, F., Jiang, X., K, H.-S., Li, Z., Ma, X., Qin, J., Sun, H., Tang, C., Wang, T., Wu, F., Yu, W., Zhang, G., Zhang, X., Zhou, J., Zhu, X., Shi, Y., Zhao, H.-H., Deng, C.: Titanium nitride film on sapphire substrate with low dielectric loss for superconducting qubits. *Physical Review Applied* **19**, 024013 (2023) 127, 158
- [137] Bevington, P.R., Robinson, D.K.: *Data Reduction and Error Analysis For the Physical Sciences*. McGraw-Hill, New York, USA (2003) 128
- [138] Richter, P.H.: Estimating errors in least-squares fitting. TDA Progress Report, 42–122 (1995) 128
- [139] Quintana, C., Megrant, A., Chen, Z., Dunsworth, A., Chiaro, B., Barends, R., Campbell, B., Chen, Y., Hoi, I.-C., Jeffrey, E., Kelly, J., Mutus, J.Y., O’Malley, P.J.J., Neill, C., Roushan, P., Sank, D., Vainsencher, A., Wenner, J., White, T.C., Cleland, A.N., Martinis, J.M.: Characterization and reduction of microfabrication-induced decoherence in superconducting quantum circuits. *Applied Physics Letters* **105**(6), 062601 (2014) 160
- [140] Earnest, C.t., Béjanin, J.H., McConkey, T.G., Peters, E.A., Korinek, A., Yuan, H., Mariantoni, M.: Substrate surface engineering for high-quality silicon/aluminum superconducting resonators. *Superconductor Science and Technology* **31**, 125013 (2018) 160
- [141] Khattak, C.P., Schmid, F., Smith, M.B.: Correlation of sapphire quality with uniformity and optical properties. *Window and Dome Technologies and Materials V* **3060**, 250–257 (1997) 161, 227
- [142] Halloran, M.H., Condon, J.H., Graebner, J.E., Kunzler, J.E., Hsu, F.S.L.: Experimental study of the Fermi surfaces of niobium and tantalum. *Physical Review B* **1**(2), 366 (1970) 166

- [143] Bobrov, N.L., Rybal'chenko, L.F., Fisun, V.V., Yanson, I.K.: Point-contact function of electron-phonon interaction in tantalum. *Fizika Nizkikh Temperatur* **13**(6), 611–626 (1987) 166
- [144] Schreier, J.A., Houck, A.A., Koch, J., Schuster, D.I., Johnson, B.R., Chow, J.M., Gambetta, J.M., Majer, J., Frunzio, L., Devoret, M.H., Girvin, S.M., Schoelkopf, R.J.: Suppressing charge noise decoherence in superconducting charge qubits. *Physical Review B* **77**, 180502 (2008) 175
- [145] Shalibo, Y., Rofe, Y., Shwa, D., Zeides, F., Neeley, M., Martinis, J.M., Katz, N.: Lifetime and coherence of two-level defects in a josephson junction. *Physical Review Letters* **105**, 177001 (2010)
- [146] Stoutimore, M.J.A., Khalil, M.S., Lobb, C.J., Osborn, K.D.: A josephson junction defect spectrometer for measuring two-level systems. *Applied Physics Letters* **101**, 062602 (2012) 175
- [147] Rosenblum, S., Reinhold, P., Mirrahimi, M., Jiang, L., Frunzio, L., Schoelkopf, R.J.: Fault-tolerant detection of a quantum error. *Science* **361**(6399), 266–270 (2018) 192
- [148] Teoh, J., Winkel, P., Babla, H.K., Chapman, B.J., Claes, J., Graaf, S.J., Garmon, J.W.O., Kalfus, W.D., Lu, Y., Maiti, A., Sahay, K., Thakur, N., Tsunoda, T., Xue, S.H., Frunzio, L., Girvin, S.M., Puri, S., Schoelkopf, R.J.: Dual-rail encoding with superconducting cavities. *arXiv preprint arXiv:2212.12077* (2022) 193
- [149] Chou, K.S., Shemma, T., McCarrick, H., Chien, T.-C., Teoh, J.D., Winkel, P., Anderson, A., Chen, J., Curtis, J., Graaf, S.J., Garmon, J.W.O., Gudwelski, B., Kalfus, W.D., Keen, T., Khedkar, N., Lei, C.U., Liu, G., Lu, P., Lu, Y., Maiti, A., Mastalli-Kelly, L., Mehta, N., Mundhada, S.O., Narla, A., Noh, T., Tsunoda, T., Xue, S.H., Yuan, J.O., Frunzio, L., Aumentado, J., Puri, S., Girvin, S.M., Moseley Jr., S.H., Schoelkopf, R.J.: Demonstrating a superconducting dual-rail cavity qubit with erasure-detected logical measurements. *arXiv preprint arXiv:2307.03169* (2023) 193
- [150] Niedzielski, B.M., Kim, D.K., Schwartz, M.E., Rosenberg, D., Calusine, G., Das, R., Melville, A.J., Plant, J., Racz, L., Yoder, J.L., Ruth-Yost, D., Oliver, W.D.: Silicon hard-stop spacers for 3D integration of superconducting qubits. *IEEE International Electron Devices Meeting (IEDM)*, 31–313134 (2019) 201
- [151] Ratter, K.: Epitaxial rhenium, a clean limit superconductor for superconducting Qbits. PhD thesis (2017) 202, 238
- [152] Altoé, M.V.P., Banerjee, A., Berk, C., Hajr, A., Schwartzberg, A., Song, C., Alghadeer, M., Aloni, S., Elowson, M.J., Kreikebaum, J.M., Wong, E.K., Griffin,

S.M., Rao, S., Weber-Bargioni, A., Minor, A.M., Santiago, D.I., Cabrini, S., Siddiqi, I., Ogletree, D.F.: Localization and mitigation of loss in niobium superconducting circuits. *PRX Quantum* **3**(2), 020312 (2022) 230

- [153] Kamal, A., Yoder, J.L., Yan, F., Gudmundsen, T.J., Hover, D., Sears, A.P., Welanders, P., Orlando, T.P., Gustavsson, S., Oliver, W.D.: Improved superconducting qubit coherence with high-temperature substrate annealing. arXiv preprint arXiv:1606.09262 (2016) 236

Appendix A

Device fabrication for Thin-Film Devices

A.1 Substrate Preparation

Substrate preparation includes an initial cleaning with strong acids followed by an optional annealing treatment. All devices were fabricated on c-plane sapphire substrates grown using either the edge-fed film growth (EFG) method or heat-exchange method (HEM). HEMEX wafers were additionally graded HEM wafers based on superior optical properties[106, 141]. The substrate preparation process is as follows:

1. Initial piranha clean: 2:1 $\text{H}_2\text{SO}_4:\text{H}_2\text{O}_2$ for 20 min[64]. This clean removed organic residues and most metal contaminants. It is recommended to use PTFE beakers for this process.
2. Rinse in flowing DI water for 20 min. This long rinse is necessary to fully remove acid residues from the wafer; these residues can damage the sapphire during high-temperature processes.
3. (Optional) Anneal in a FirstNano EasyTube 6000 furnace at 1200 °C in an oxygen-rich environment. The annealing processes is as follows:
 - (a) Preheat furnace to 200 °C and purge with nitrogen. A purge is defined as flowing a volume of gas equivalent to five times the volume of the annealing chamber.
 - (b) Load wafer, then purge wafer with pure oxygen.
 - (c) Heat furnace to 1200 °C with a linear ramp of 400 °C/hr while continuously flowing oxygen.
 - (d) Once the furnace reaches 1200 °C, turn off gas flows, anneal for 1 hr.
 - (e) Cool down by turning off the heaters and flowing a 4:1 mixture of $\text{N}_2:\text{O}_2$ gas. The wafer passively cools over approximately 6 hours.

A.2 Aluminum/Josephson Junction Process

Aluminum-based thin-film circuits are patterned additively using electron-beam lithography and a liftoff process. The fabrication of Al resonators and Josephson junctions are done the same way, only the specific patterns are different. The Al process can be performed immediately after substrate preparation or after the tantalum process in the case of tantalum-based transmons (Appendix A.3). The aluminum process is as follows:

1. Dehydration bake: put wafer on a hotplate set to 180 °C for 5 min. This step is to remove any moisture left on the wafer.
2. Spin electron-beam resist (EBR) bilayer:
 - (a) Spin MMA (8.5) MAA EL13 at 4000 rpm (700 nm).
 - (b) Soft bake on hotplate at 180 °C for 5 min.
 - (c) Spin 950K PMMA A4 at 4000 rpm (200 nm).
 - (d) Soft bake on hotplate at 180 °C for 5 min.
3. Deposit 15nm Al with electron-beam evaporation at a rate of 0.2 nm/s. This applies an anti-charging layer on top of the resist to eliminate charging effects during electron-beam lithography (EBL).
4. Perform EBL with Raith EBPG 5200+ to expose pattern. Undercuts are dosed at $\approx 300 \mu\text{C}/\text{cm}^2$, large area features are dosed at $\approx 800 \mu\text{C}/\text{cm}^2$, fine features are dosed at $\approx 1400 \mu\text{C}/\text{cm}^2$.
5. Remove anti-charging layer: immerse wafer in Microposit MF312 developer for 80 sec.
6. Develop in 3:1 IPA:H₂O at 6 °C for 2 min.
7. Load wafer into the load-lock of the electron-beam evaporator (Plassys UMS300).
8. Perform Ar ion beam clean at 400 V, 20 mA to remove the tantalum oxide (in the case of depositing junctions on a Ta-based transmon) and other surface residues prior to aluminum deposition.
 - (a) tilt wafer by ± 45 degrees, perform cleaning for 34 sec. at each angle. This is done to clean under dolan bridges, under artificial undercuts, and on the Ta sidewalls in the case of the Ta-based processes.
9. Transfer wafer to the evaporation chamber without breaking vacuum, tilt the wafer by -25 degrees, deposit 20 nm Al at 1 nm/s.
10. Transfer wafer to oxidation chamber, perform static oxidation using an 85:15 Ar:O₂ mixture at 30 Torr for 10 min. This step grows the junction oxide.

11. Transfer wafer back to evaporation chamber, tilt the wafer by +25 degrees, deposit 30 nm Al at 1 nm/s.
12. Transfer wafer back to oxidation chamber, perform static oxidation using an 85:15 Ar:O₂ mixture at 100 Torr for 5 min. This step caps the surface of the bare aluminum with pure aluminum oxide.
13. Liftoff by immersing the wafer in NMP at 90 °C for 1 hr.
14. Sonicate 2 min each in N-Methylpyrrolidone (NMP), acetone, isopropanol, and DI water. Dry with nitrogen.

For Al-based transmons, Josephson junctions are patterned with the rest of the circuit.

A.3 Tantalum Process

For tantalum-based devices, tantalum was deposited after the substrate preparation process and patterned using subtractive lithography via reactive ion etching (RIE). The process is as follows:

1. Deposit tantalum by DC magnetron sputtering in a Lesker CMS-18 sputtering system:
 - (a) Heat wafer to 400 °C at a rate of 100 °C/min.
 - (b) Dehydration bake: Bake for 15 min.
 - (c) Heat wafer to 800 °C at a rate of 100 °C/min.
 - (d) Thermalization bake: Bake for 10 min. This is to thermalize the wafer to 800 °C prior to deposition.
 - (e) Deposit 150 nm Ta by sputtering using an Ar pressure of 6 mTorr, a power of 300 W, and a deposition rate of 2.5 Å/sec.
 - (f) Bake for 10 min. at 800 °C. This step is intended to anneal the film.
 - (g) Cool down to 500 °C at a rate of 10 °C/min. This is to prevent substrate damage due to the differential contraction of the Ta film and the sapphire surface.
 - (h) Cool down to room temperature at a rate of 100 °C/min. This step takes 2-3 hrs.
2. Spin photoresist (PR):
 - (a) Spin S1827 at 4000 rpm ($\approx 2.5 \mu\text{m}$).
 - (b) soft bake on hotplate at 115 °C for 3 min.
 - (c) Remove edge bead using a Q-tip dipped in acetone. This step is to allow the photomask to directly contact the resist with no spacing. This step is unnecessary if exposing with a maskless aligner.

3. Expose pattern using a glass photomask and a Suss MJB4 contact aligner. $\approx 300 \text{ mJ/cm}^2$ broadband exposure dose.
4. develop resist in Microposit MF319 developer for $\approx 1 \text{ min.}$ with gentle agitation.
5. Hard bake wafer on hotplate at $120 \text{ }^\circ\text{C}$ for 1 min. This step improves the chemical stability of the resist in preparation for RIE.
6. Descum with oxygen plasma using an AutoGlow 200 at 150 W and 300 mTorr O_2 for 2 min. This step removes resist residue in developed regions.
7. Etch the tantalum in an Oxford 80+ Reactive Ion Etcher using SF_6 with a flow rate of 20 sccm, a pressure of 10 mTorr, and an RF power of 50 W. This recipe gives a Ta etch rate of $\approx 100 \text{ nm/min}$ with excellent etch selectivity to the PR.
8. Remove PR: sonicate 2 min each in N-Methylpyrrolidone (NMP), acetone, isopropanol, and DI water. Dry with nitrogen.
9. Acid clean: 2:1 $\text{H}_2\text{SO}_4:\text{H}_2\text{O}_2$ for 20 min. This step is to remove any remaining organic residue. It is recommended to use PTFE beakers for this process.
10. Oxide strip: Transene 10:1 buffered oxide etch (BOE) for 20 min. It is recommended to use PTFE beakers for this process. This step removes the tantalum oxide. The long 20 min. immersion has been shown to improve surface quality by causing a thinner oxide to regrow upon exposure to air[110, 152].
11. Rinse in flowing DI water for 20 min.

For Ta-based transmons, the junctions are subsequently patterned using the process in Appendix A.2.

A.4 Rhenium Process

The rhenium process is very similar to the tantalum process. Rhenium was deposited after the substrate preparation process and patterned using subtractive lithography via reactive ion etching (RIE). Notably, rhenium is deposited at a slower rate and a higher temperature in order to increase grain size. Additionally, piranha solution etches rhenium, so an alternate cleaning using diluted sulfuric acid is used. The process is as follows:

1. Deposit rhenium by DC magnetron sputtering in a Lesker CMS-18 sputtering system:
 - (a) Heat wafer to $400 \text{ }^\circ\text{C}$ at a rate of $100 \text{ }^\circ\text{C/min.}$
 - (b) Bake for 15 min.
 - (c) Heat wafer to $700 \text{ }^\circ\text{C}$ at a rate of $100 \text{ }^\circ\text{C/min.}$

- (d) Heat wafer to 900 °C at a rate of 20 °C/min. Instead of doing a thermalization bake, the wafer is heated up slowly to the deposition temperature to allow the wafer to thermalize as it is heating. This minimizes the radiative heating of the walls of the deposition chamber which would result in outgassing and potential film contamination during deposition.
 - (e) Deposit 150 nm Re by sputtering using an Ar pressure of 2 mTorr, a power of 100 W, and a deposition rate of 0.87 Å/sec.
 - (f) Cool down to 500 °C at a rate of 20 °C/min. Since Re has far better lattice matching with sapphire than Ta, cooling down faster does not introduce substrate damage.
 - (g) Cool down to room temperature at a rate of 100 °C/min. This step takes 2-3 hrs.
2. Spin photoresist (PR):
 - (a) Spin S1827 at 4000 rpm ($\approx 2.5 \mu\text{m}$).
 - (b) soft bake on hotplate at 115 °C for 3 min.
 - (c) Remove edge bead using a Q-tip dipped in acetone. This step is to allow the photomask to directly contact the resist with no spacing. This step is unnecessary if exposing with a maskless aligner.
 3. Expose pattern using a glass photomask and a Suss MJB4 contact aligner. $\approx 300 \text{ mJ/cm}^2$ broadband exposure dose.
 4. develop resist in Microposit MF319 developer for ≈ 1 min. with gentle agitation.
 5. Hard bake wafer on hotplate at 120 °C for 1 min. This step improves the chemical stability of the resist in preparation for RIE.
 6. Descum with oxygen plasma using an AutoGlow 200 at 150 W and 300 mTorr O₂ for 2 min. This step removes resist residue in developed regions.
 7. Etch the tantalum in an Oxford 80+ Reactive Ion Etcher using SF₆ with a flow rate of 20 sccm, a pressure of 10 mTorr, and an RF power of 50 W. This recipe gives a Ta etch rate of $\approx 12 \text{ nm/min}$ with excellent etch selectivity to the PR.
 8. Remove PR: sonicate 2 min each in N-Methylpyrrolidone (NMP), acetone, isopropanol, and DI water. Dry with nitrogen.
 9. Acid clean: 4:1 H₂SO₄:H₂O heated to 100 °C on a hotplate for 20 min. This step is to remove any remaining organic residue. Piranha cannot be used as it will attack the Re film. Because heating of the solution is required, PTFE beakers are not used due to its low melting point and low thermal conductivity. Since there is evidence that borosilicate glassware can contaminate the wafer when hot sulfuric acid is used[64], it is recommended to use quartz beakers for this process.

10. Oxide strip: Transene 10:1 buffered oxide etch (BOE) for 5 min. It is recommended to use PTFE beakers for this process. This step removes any oxide that may have formed from the oxidizing acid of the previous step.
11. Rinse in flowing DI water for 20 min.

For Re-based transmons, the junctions are subsequently patterned using the process in Appendix A.2.

A.5 Flip-Chip Process

The flip-chip process is used to pattern indium bumps and hardstops on top of an existing circuit layer patterned on a wafer. Hardstops are deposited through a stainless steel shadowmask. Indium bumps are deposited using negative-tone optical resist and a liftoff process. This step is performed after the circuit layer and Josephson junctions are patterned with a Ta, Re, or Al process (Appendix A.3, A.4, or A.2). Then the following process is used:

1. Hardstop deposition: deposit 2 μm Ti via electron-beam evaporation using a Plassys UMS300. Deposition is performed through a stainless steel shadowmask that was laser-cut with $\approx 25 \mu\text{m}$ accuracy. The mask was positioned $\approx 1 \text{ mm}$ above the wafer during deposition.
2. Spin photoresist (PR):
 - (a) Spin AZnLOF2070 at 2000 rpm ($\approx 6 \mu\text{m}$).
 - (b) soft bake on hotplate at 110 $^{\circ}\text{C}$ for 3 min.
 - (c) Remove edge bead using a Q-tip dipped in acetone. This step is to allow the photomask to directly contact the resist with no spacing. This step is unnecessary if exposing with a maskless aligner.
3. Expose pattern using a glass photomask and a Suss MJB4 contact aligner. $\approx 300 \text{ mJ}/\text{cm}^2$ broadband exposure dose.
4. Post-exposure bake (PEB) wafer on hotplate at 120 $^{\circ}\text{C}$ for 1 min. This step tunes the undercut resist profile after development.
5. develop resist in Microposit MF319 developer for $\approx 2 \text{ min.}$ with gentle agitation.
6. Deposit indium with thermal evaporation in a Lesker PVD75 deposition system:
 - (a) Ar ion beam clean at 100 V for 2 min. This removes surface residues prior to deposition.
 - (b) Deposit 100 nm In at 0.2 nm/sec. This step deposits a seed layer of In to improve film adhesion.

- (c) Deposit 3.9 μm In at 2 nm/sec. A total of 4 μm In is deposited.
- 7. Liftoff by immersing the wafer in Technistrip NI555 at 80 °C for 2 hrs. At this point, the indium was not fully lifted off, and the wafer should have been left in the liftoff bath for 2-4 more hrs. However, after only two hours the wafer was removed and the next step was performed.
- 8. Sonicate 5 min each in N-Methylpyrrolidone (NMP), acetone, isopropanol, and DI water. Dry with nitrogen. This step completed the liftoff; however, the lifted-off indium redeposited onto the wafer. If the previous step was carried out for a longer period of time, this issue may have been averted.

A.6 Wafer Dicing

After circuit patterning, the wafer is diced into individual chips. The process is as follows:

1. Spin photoresist (PR). This serves as a protective layer for the patterned circuit in preparation for the dicing process:
 - (a) Spin S1827 at 2000 rpm ($\approx 4 \mu\text{m}$).
 - (b) soft bake on hotplate at 80 °C for 10 min. At this point, Josephson junctions are patterned on the wafer. These junctions are at risk of shorting if subjected to high temperatures for prolonged periods of time. As a result, the soft bake is conducted at lower temperature for a longer period of time.
2. Dice with ADT ProVectus 7100 dicer.
3. Remove PR: sonicate 2 min each in N-Methylpyrrolidone (NMP), acetone, isopropanol, and DI water. Dry with nitrogen.

After dicing, planar devices can be packaged and measured. Flip-chip devices must be bonded using a wafer bonder, the procedure for which is described in Appendix A.7.

A.7 Flip-chip Bonding

Flip-chip devices were bonded using a SET FC 150 automatic flip-chip bonder. This tool has an arm that holds the top chip and a chuck that holds the bottom chip and can align the two chips with an accuracy of $< 5 \mu\text{m}$. To ensure that the top and bottom chips are parallel to each other after bonding, the chips are bonded a second time after rotating the chip by 180 degrees. A single bond sequence is as follows:

1. Scrubbing along the x -axis by $\pm 3 \mu\text{m}$ for 30 sec. When the two chips make contact, before any force is applied, the chuck will move transversely a few microns back and forth to perform a sort of “rubbing” action with the intention of breaking through the indium oxide.

2. Heat the chips to 105 °C with a ramp time of 30 sec. The scrubbing and heating steps occur at the same time. Heating the chips soften the indium, allowing it to deform easier.
3. Apply 50 kg force with a ramp time of 30 sec. while maintaining a temperature of 105 °C
4. Maintain 50 kg bonding force and a temperature of 105 °C for 60 sec.
5. Turn off heaters and let the chips cool down while maintaining 50 kg of bonding force for 120 sec. This lets the indium cool down and settle into its new bonded configuration while force is still added to maintain the flip-chip gap.

It is worth noting that this process is different from what was used in Lei et al. [82]. In particular, I do not use the OES Ontos7 to perform an atmospheric plasma treatment that utilizes helium, nitrogen ,and hydrogen plasma to remove the indium oxide on the bumps and replace it with a more brittle nitride. Instead, I used heat and scrubbing to break through the oxide. The result was ultra-high seam quality, as shown in Table 6.3.

Appendix B

Materials Characterization

Here, I will describe some of the materials characterization work done in collaboration with scientists at the Center for Functional Nanomaterials (CFN) at Brookhaven National Laboratory (BNL). DC transport measurements were made using a Quantum Design PPMS DynaCool, X-ray diffractometry (XRD) was done using a Rigaku Miniflex II XRD, and Transmission electron microscopy (TEM) was done using a FEI Talos F200X; all were done at BNL. Atomic force microscopy (AFM) was done at Yale using a Bruker Dimension Fastscan AFM. Samples for TEM were pre-prepared by sputtering ≈ 20 nm of gold on top of them to protect them from the sample preparation process. The gold was sputtered several weeks after sample fabrication; as a result, the MA region is representative of one that has been exposed to air.

B.1 Sapphire Annealing

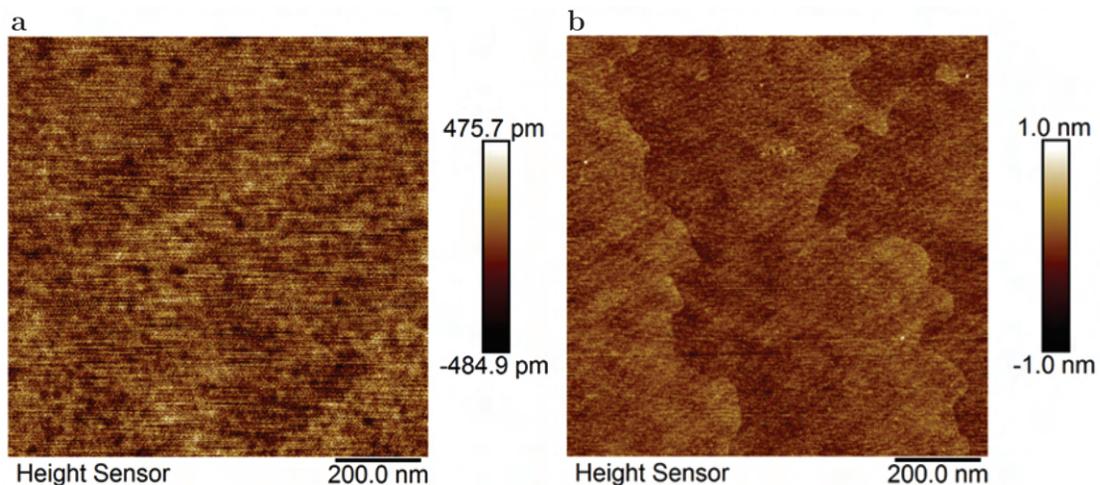


Figure B.1: **Atomic force microscopy of sapphire surface** before annealing (a) and after annealing (b). Figure obtained from Ganjam et al. [86]

Atomic Force Microscopy (AFM) on sapphire substrates was conducted before and after annealing of an EFG sapphire substrate. Surfaces before annealing had sub-nanometer roughness and uniform surface topology with no distinguishable features (Fig. B.1a). After annealing, surfaces were atomically flat and displayed a terraced structure typically seen for annealed c-plane sapphire (Fig. B.1b)[110, 153]. The terraces have step height of around 220 pm, approximately equal to the inter-atomic spacing in the c-axis ($c/6 = 216$ pm), and width 420 nm related to the miscut angle of the wafer, which in this case is calculated to be approximately 0.03° .

B.2 Aluminum Film Characterization

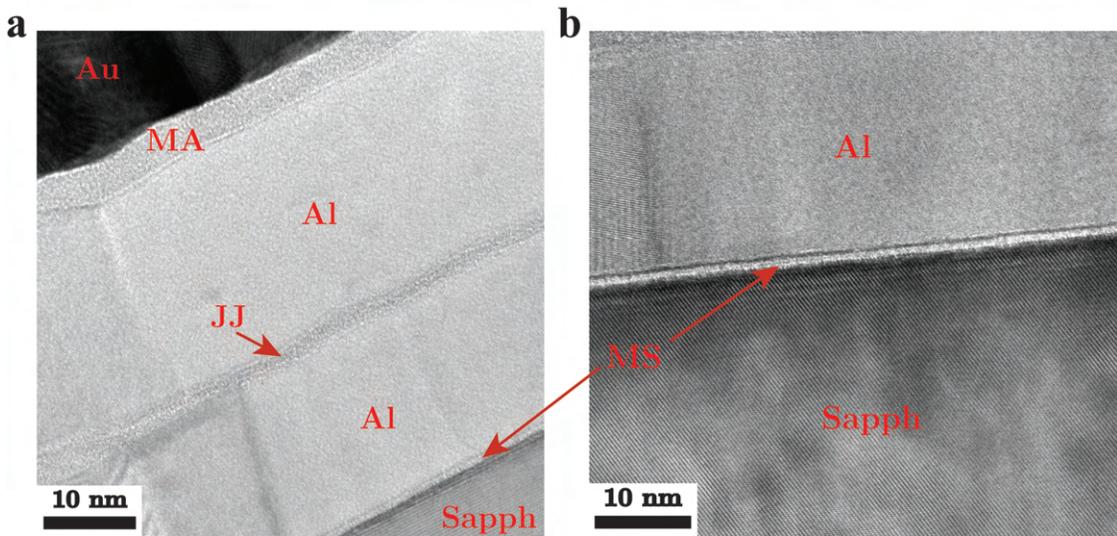


Figure B.2: **TEM of Al/AIO_x/Al film.** Brightfield TEM of the MS and MA interfaces of a typical Al/AIO_x/Al film. The MA interface appears amorphous and has general stoichiometry AlO_x. The lower and upper Al layers are the two Josephson junction electrodes, and the AlO_x between them act as the tunnel barrier. TEM done by Kim Kisslinger at Brookhaven National Laboratory. Figure adapted from Ganjam et al. [86]

A Al/AIO_x/Al junction fabricated using the methods of Appendix A.2 was measured using TEM. The aluminum/sapphire interface has a thin (≈ 2 nm) amorphous region (Fig. B.2b). The MA interface of aluminum has an ≈ 5 nm oxide layer (Fig. B.2a). Between the two layers of aluminum lies the Josephson junction oxide, which is approximately ≈ 2 nm and looks amorphous, similar to the aluminum MA and MS interfaces.

B.3 Tantalum Film Characterization

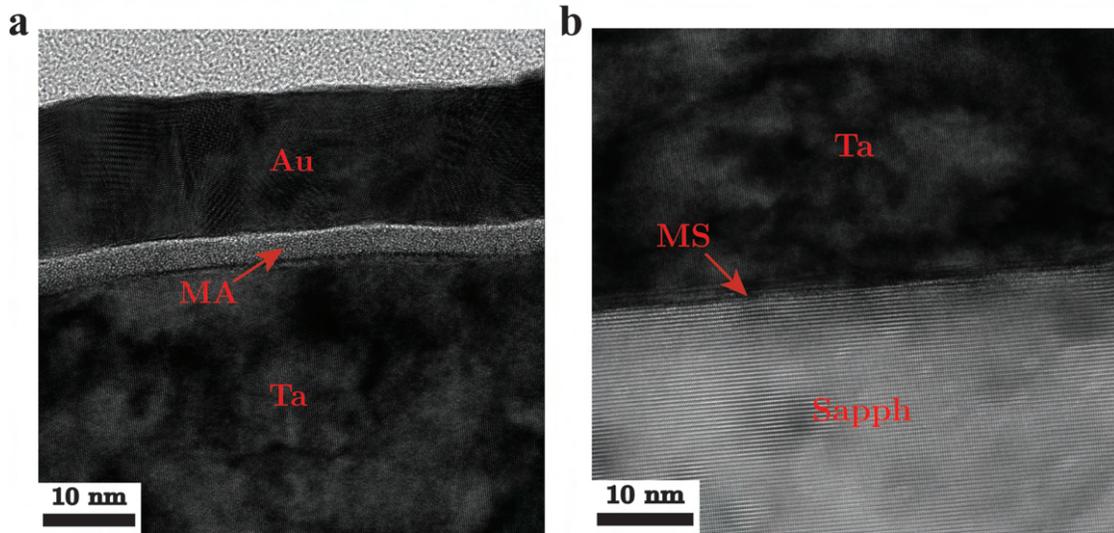


Figure B.3: **TEM of a tantalum film.** Brightfield TEM of the (a) MA and (b) MS interfaces of a typical Ta thin film. Like the MA interface of the Al films, the MA interface of the Ta film also appears amorphous, with general stoichiometry TaO_x . TEM done by Kim Kisslinger at Brookhaven National Laboratory. Figure adapted from Ganjam et al. [86]

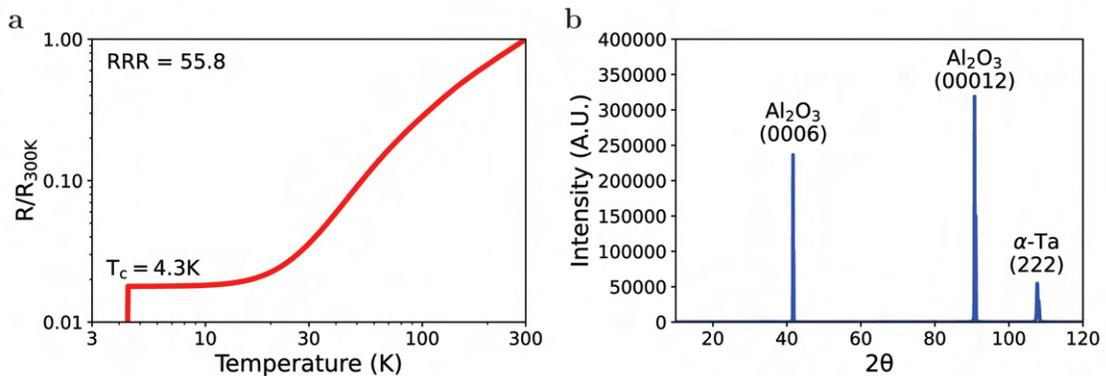


Figure B.4: **Tantalum film characterization** (a) DC resistance as a function of temperature. The sharp drop in resistance at $T = 4.3\text{ K}$ indicates the emergence of the superconducting state. (b) X-ray diffractometry (XRD) of a Ta film. This particular film was entirely in the (111) orientation; other films were either entirely (110) or a mixture of (111) and (110) (not shown). Resistance measurements were done by Ruoshui Li and Mingzhao Liu at Brookhaven National Laboratory. XRD was done by Chenyu Zhou, Yichen Jia, and Mingzhao Liu at Brookhaven National Laboratory. Figure obtained from Ganjam et al. [86]

A Ta film fabricated using the methods of Appendix A.3 was measured using TEM. Unlike the MS region of the Al films, the MS region of the Ta film seems free of amorphous material and displays nearly epitaxial growth (Fig: B.3b). Additionally, the metal-air (MA) interface of tantalum has a thin (≈ 3 nm) oxide layer, which seems to be thinner than that of the aluminum film.

Room temperature resistivity of the as-deposited tantalum films was measured to be $\rho^{300\text{ K}} \approx 14.2 \times 10^{-8} \Omega\text{m}$ and varied by only around 10% from film to film. Resistance as a function of temperature was measured for multiple samples, all of which had $T_c > 4.17$ K and $\text{RRR} > 15$, with our best sample having $T_c = 4.3$ K and $\text{RRR} = 55.8$ (Fig. B.4a). XRD of the Ta films confirmed the dominant presence of α -Ta growing in either the (111) or (110) orientation, while the β phase was not observed (Fig. B.4b).

The ability to deposit Ta films that are consistently in the α -phase with high T_c , RRR, and epitaxial film growth is due to the clean deposition system we have at Yale. The Lesker CMS-18 sputtering system used to deposit these films is dedicated for only depositing refractory materials (Ta, Re, Nb, Ti, and their alloys and nitrides). The system underwent multiple bakeout processes to reduce the base pressure to $\approx 5 \times 10^{-10}$ Torr at room temperature and $\approx 5 \times 10^{-9}$ Torr at the deposition temperature of 800 °C. This allows us to consistently grow very high quality films.

B.4 Rhenium Film Characterization

A Re film fabricated using the methods of Appendix A.4 was measured using TEM. Similar to the Ta films, the rhenium/sapphire MS interface is free from amorphous material and a distinct boundary exists between the two materials (Fig. B.5b). Additionally, there is no oxide or amorphous material between the gold and the Re, which indicates that no oxide is present (Fig. B.5a), which is confirmed by EDS (Fig. B.5c). Remarkably, this sample lay exposed to air for over two months before being covered with gold and sent to BNL for the TEM measurement, which indicates that Re is quite robust to oxidation at room temperature.

Room temperature resistivity of the as-deposited rhenium film was measured to be $\rho^{300\text{ K}} \approx 19.0 \times 10^{-8} \Omega\text{m}$. Too few rhenium thin-films were measured to assess the consistency of the deposition process. However, another Re film was grown at 800 °C and a deposition rate of $\approx 0.25 \text{ \AA}/\text{sec}$; this film had $\rho^{300\text{ K}} \approx 22.3 \times 10^{-8} \Omega\text{m}$, which differs from the film grown at 900 °C by $< 20\%$. Resistance as a function of temperature was measured for one sample (Fig. B.6), and is shown to have $T_c \approx 1.95$ K, consistent with other studies[151].

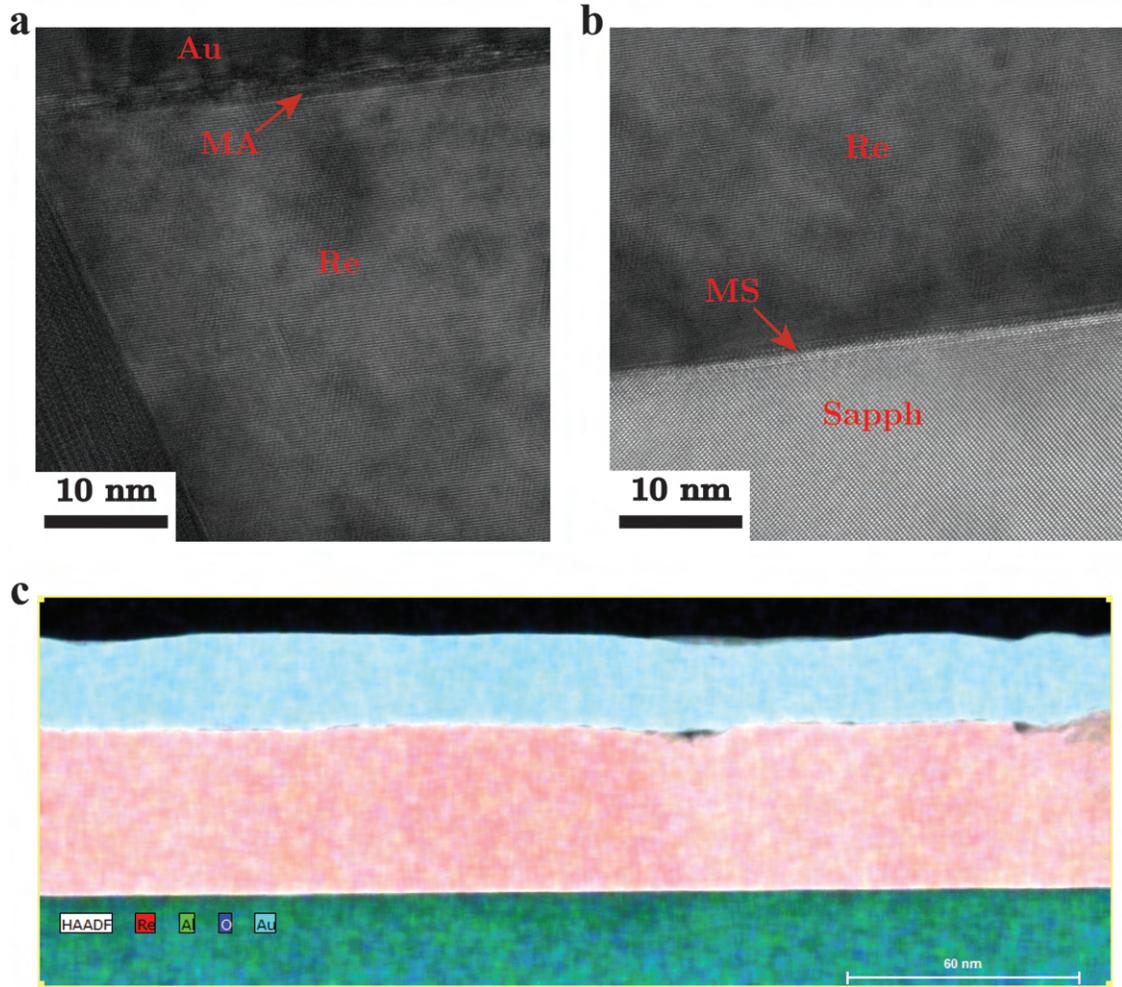


Figure B.5: **TEM of a rhenium film.** Brightfield TEM of the (a) MA and (b) MS interfaces of a Re thin film. Unlike the MA interface of the Al and Ta films, the MA interface of the Re film appears to contain no oxide. Similar to the tantalum film, the MS interface contains no amorphous material. **c** High-angle annular dark-field scanning transmission electron microscopy (HAADF-STEM) image overlaid with an elemental map formed by energy-dispersive X-ray spectroscopy (EDS). This shows a clear lack of oxygen in the MA interface. TEM done by Kim Kisslinger at Brookhaven National Laboratory.

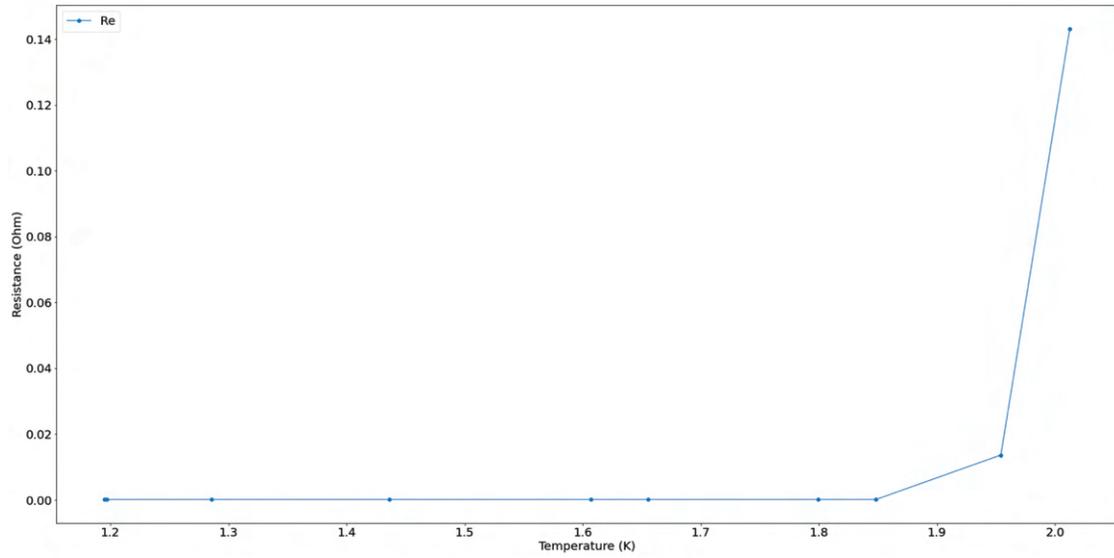


Figure B.6: **Rhenium film** T_c . Low temperature DC resistance of a rhenium film, showing $T_c \approx 1.95$ K. RRR was not measured for this sample. Resistance measurements were done by Ruoshui Li and Mingzhao Liu at Brookhaven National Laboratory.

Appendix C

Surface Participation Ratio Simulations for Thin Films

Expressions for calculating participation ratios were given in Ch. 3.3. For the most part, these expressions are straightforward to apply in finite-element solvers such as Ansys HFSS. The surface participations of thin-films, however, are slightly more complex due to the large aspect ratios involved in the simulation. Proper calculation of surface participation ratios requires considering the thickness of the films as well as the 3 nm-thick lossy surface regions. In the full 3D electromagnetic simulation of a \sim cm-scale device, the \sim nm-scale regions that need to be simulated requires the accurate solving of Maxwell's equations over 7 orders of magnitude in length. This is too computationally intensive to do without employing some tricks.

For one, thin-films can be approximated as infinitely thin conducting sheets. This approximation is reasonable in the case where the general electromagnetic mode structure of the system needs to be understood. In such cases, these sheets are often simulated to be perfect conductors; in other words, there is no field penetration through the conductor. Therefore, for all intents and purposes, a thick film is the same as a thin film for the purposes of determining large-scale electromagnetic field distributions within the simulation space. This reasoning also motivates the omission of surface dielectric regions in these simulations, as their small thickness is not expected to significantly change the global electromagnetic field behavior.

For the purposes of calculating participations, the surface regions can be approximated as regions whose electric field is constant along the thickness axis. In most regions, this approximation is appropriate, as the fields are not expected to diverge rapidly over such small distances. As a result, calculating these participations can be as simple as integrating the electric field over the surface of that interface and multiplying it by the assumed interface thickness (3 nm). This is what is done for calculating p_{MA} for bulk superconductors, after the vacuum electric field is rescaled according to the displacement field continuity relation (see Ch. 3.3.2). However, in thin-film resonators, currents tend to flow along the edges of the film; as a result, charge density accumulates there, resulting in highly divergent fields whose field behavior must be accounted for at sub-micron length scales. Additionally,

infinitely thin sheets cannot capture this behavior, as the charge density is undefined at the edge. As a result, calculating the surface participations in this way can underestimate them by up to an order of magnitude. At the same time, a full-thickness simulation of the edge region with explicitly defined surface regions requires nanometer-scale meshing; simulating such a system is far too computationally inefficient.

To address this issue, the method detailed in Wang et al. [74] is utilized, which I will briefly review here. In this method, the 3D global simulation still approximates the thin films as infinitely thin sheets; however, this simulation is paired with a 2D cross sectional electrostatic simulation that explicitly simulates the actual thickness of the films and interface regions at the edges of the films. By determining how the electric field scales at the edges, surface participations can be accurately computed in the 3D simulation by rescaling the edge-field to capture the divergent field behavior.

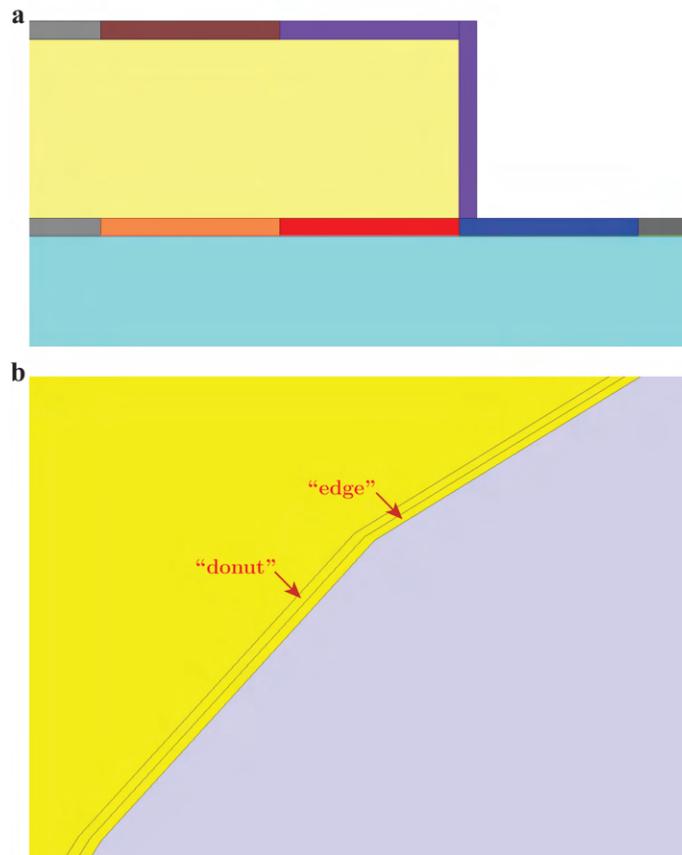


Figure C.1: **Simulating edge-field behavior of thin films.** **a** The field behavior is simulated in a 2D cross-sectional electrostatic simulation. The “edges” of each interface are denoted by purple (MA edge), red (MS edge), and blue (SA edge). These regions do not converge in 3D simulations due to diverging charge densities. The “donut” region is defined as the region that is adjacent to the “edge” regions but converges in both 2D and 3D simulations. **b** The “edge” and “donut” regions in a 3D electromagnetic simulation.

Fig. C.1a shows the 2D cross-section of the thin film and the three interfaces. The

“edges” of each interface are denoted by purple (MA edge), red (MS edge), and blue (SA edge). These regions do not converge in 3D simulations due to diverging charge densities. The “donut” region is defined as the region that is adjacent to the “edge” regions but converges in both 2D and 3D simulations. Assuming a linear scaling between the “donut” field and the “edge” field, a scale factor can be calculated that describes the energy in the edge regions:

$$F_i = \frac{U_{\text{edge},i}}{U_{\text{donut},i}} \quad (\text{C.1})$$

where $i = \text{MS, MA}$. Since the SA edge region is located near the MA and MS regions, a similar linear scaling factor can be determined:

$$F_{\text{SA}} = \frac{U_{\text{edge},i}}{U_{\text{donut},i}} \quad (\text{C.2})$$

where $i = \text{MS, MA}$ once again; in other words, F_{SA} can be determined relative to either the MS “donut” or the MA “donut” regions. This scale factor can then be used in the 3D simulation to determine the “edge” participation by calculating the “donut” participation (Fig. C.1b):

$$p_{\text{edge},i} = F_i p_{\text{donut},j} \quad (\text{C.3})$$

where $i = \text{SA, MS, MA}$ and $j = i$ for $i \neq \text{SA}$, in which case $j = \text{MS, MA}$ depending on how F_{SA} is defined.

To simulate the cross-section in Fig. C.1a, the appropriate boundary conditions and “edge”/“donut” lengths must be chosen. In order for the electrostatic simulation to be an accurate approximation of the field behavior in an AC electromagnetic simulation, the electrostatic simulation here has to be over a small enough region such that the field behavior is independent of far-field conditions. In order to satisfy this requirement, the length of the “edge”/“donut” regions must be much smaller than the distance between this film and another conductor. In the simplest case of a stripline in a coaxial tunnel, the other conductor is the tunnel itself, and a voltage is applied from the film to the walls of the tunnel. In this example, $x_0 \ll g$, where $x_0 \sim 1 \mu\text{m}$ is the length of the “donut” plus the “edge” region, and $g \sim 1 \text{mm}$ is the distance between the film and the other conductor. Additionally, since the divergent field behavior at the edges becomes more pronounced for skinnier traces, the width w of the thin-film strip must also be considered. Since x_0 needs to be small enough such that the “edge” and “donut” fields are linearly related, the size of these regions has another constraint where $x_0 \ll w$. In practice, for a $1 \mu\text{m}$ -wide thin-film strip, x_0 should be $\approx 0.1 \mu\text{m}$, otherwise the linear scaling may not be a good enough approximation.

With this scaling method, it can be determined that the majority of surface participation in thin-film devices comes from the edge regions. This also provides an intuitive explanation for why Γ_{surf} is a good predictor of surface losses in transmons and hairpin striplines. Since the surface energies of the three interfaces are all localized in close proximity to each other, they are likely all affected by a process in similar ways.

Since charge density accumulates at the film edges, so too does the current. As a result, the scale factors can also be used to calculate p_{cond} . Since current can flow on the upper face of the conductor, the sidewall, and the bottom face, the surface magnetic field of the top edge and sidewall of the conductor can be calculated using F_{MA} . Likewise, the surface magnetic field at the bottom edge of the film can be calculated using the MS “donut” region and F_{MS} .

The requirement that $x_0 \ll w$ means that the above method will not be effective in simulating the participations near the Josephson junction, where $w \sim 100$ nm. This would require $x_0 \approx 1 - 10$ nm which would be too small to efficiently simulate. Instead, a local 3D electrostatic simulation of the region within $10 \mu\text{m}$ of the junction that includes the 3 nm-thick surface dielectric regions is performed, where the electric fields are simulated by applying a voltage from one junction electrode to the other (Fig. C.2). This simulates the electric fields that are induced due to the parasitic capacitance of the electrodes as well as the junction capacitance. Since the junction oxide is explicitly included, its capacitance can be explicitly calculated using the junction area and oxide thickness as determined by TEM (see Fig. B.2a). In the global simulation, this near-junction region is simulated as a lumped inductor L_J in parallel with a lumped capacitance C_J ; as a result, the energy in this lumped capacitor is used to rescale the energy in the electrostatic simulation. As discussed in Ch. 5.1.2, the region within 100 nm of the junction is omitted from the participation ratio calculation due to the assumption of it being lossless. Despite this, the near-junction regions contributes almost a third of the total participation of the transmon, which explains why even the tantalum-based transmon is dominated by surface loss near the junction (see Ch. 5.1.4).

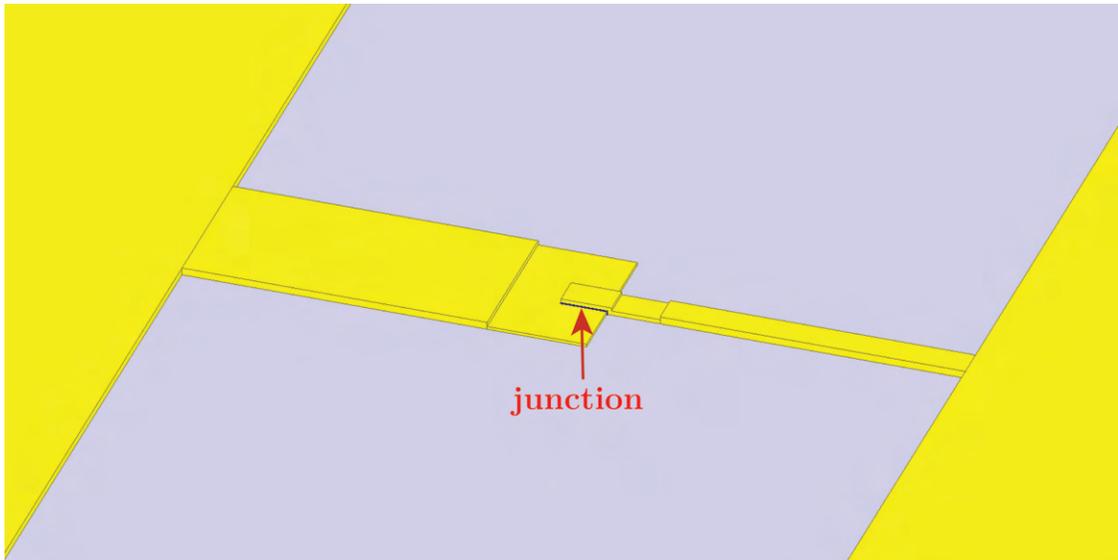


Figure C.2: **Simulating near-JJ fields.** 3D electrostatic simulation of the near-JJ region. The simulation attempts to reproduce the designed Josephson junction precisely. Electric fields are simulated by applying a voltage across the junction electrodes. 3 nm-thick interface regions are explicitly included in the simulation (not shown)

Appendix D

Copyright Permissions

- Figs. 2.3b and 2.4b,c reproduced or adapted from Koch et al. [36] under license No. RNP/23/OCT/070990 granted by the American Physical Society.
- Figs. 2.11a,b,c reproduced or adapted from Schuster et al. [62] under license No. 5645110181564 granted by Springer Nature.
- Fig. 2.11d reproduced or adapted from Barends et al. [63] under license No. RNP/23/OCT/070987 granted by the American Physical Society.
- Fig. 2.11e reproduced or adapted from Place et al. [64] under the Creative Commons CC BY license.
- Fig. 2.11f,g reproduced or adapted with permission from Andresen [65].
- Fig. 2.12a reproduced or adapted from Reagor et al. [53] under license No. RNP/23/OCT/070988 granted by the American Physical Society.
- Fig. 2.12b reproduced or adapted from Chou et al. [75] under license No. 5645110849519 granted by Springer Nature.
- Figs. 2.13 and 3.3c reproduced or adapted from Axline et al. [84] under license No. 5645110989162 granted by AIP Publishing.
- Fig. 3.3a reproduced or adapted from Göppl et al. [88] under license No. 5645111092737 granted by AIP Publishing.
- Fig. 3.3b reproduced or adapted from Geerlings et al. [89] under license No. 5645120203306 granted by AIP Publishing.
- Fig. 3.3d reproduced or adapted from Reagor et al. [90] under license No. 5645120317610 granted by AIP Publishing.
- Figs. 3.8, 3.9, 4.6, 4.7, 4.8, 4.9, 5.1, 5.2, 5.3, 5.4, 5.5, 5.6, 5.8, 5.9, 5.10, B.1, B.2, B.3, and B.4; tables 4.3, 4.4, 4.5, 4.7, 5.1, 5.2, 5.3, 5.4, 5.5, and 5.6 reproduced or adapted from Ganjam et al. [86] under the Creative Commons CC BY license.

- Figs. 4.1, 4.2, 4.3, 4.4, and 4.5 reproduced or adapted from Lei et al. [112] under license No. RNP/23/OCT/070989 granted by the American Physical Society.