THE UNIVERSITY OF CHICAGO

MILLIMETER-WAVE SUPERCONDUCTING QUANTUM DEVICES

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY
ALEXANDER VLADIMIR ANFEROV

CHICAGO, ILLINOIS
MARCH 2024

To Sophie and Miso

"If at first you don't succeed Mr. Wint...?

Try, try again Mr. Kid."

- Diamonds are Forever 1971

TABLE OF CONTENTS

LIST OF PUBLICATIONS

This thesis is based in part on the following works:

[1] Alexander Anferov, Aziza Suleymanzade, Andrew Oriani, Jonathan Simon, and David I. Schuster. Millimeter-Wave Four-Wave Mixing via Kinetic Inductance for Quantum Devices. *Physical Review Applied*, 13(2):024056, February 2020. doi:10.1103/PhysRevApplied.13.024056.

[2] Alexander Anferov, Shannon P. Harvey, Fanghui Wan, Kan-Heng Lee, Jonathan Simon, and David I. Schuster. Low-loss Millimeter-wave Resonators with an Improved Coupling Structure. *arXiv*, November 2023. doi:10.48550/arXiv.2311.01670.

[3] Alexander Anferov, Kan-Heng Lee, Fang Zhao, Jonathan Simon, and David I. Schuster. Improved Coherence in Optically-Defined Niobium Trilayer Junction Qubits. *arXiv*, June 2023. doi:10.48550/arXiv.2306.05883.

[4] Alexander Anferov, Shannon P. Harvey, Fanghui Wan, Jonathan Simon, and David I. Schuster. Superconducting Qubits Above 20 GHz Operating Over 200 mk. In preparation.

[5] Alexander Anferov, Shannon P. Harvey, Fanghui Wan, Jonathan Simon, and David I. Schuster. A Millimeter-wave Artificial Atom Cooled with Helium-4. In preparation.

[6] Aziza Suleymanzade, Alexander Anferov, Mark Stone, Ravi K. Naik, Andrew Oriani, Jonathan Simon, and David Schuster. A tunable high-Q millimeter wave cavity for hybrid circuit and cavity QED experiments. *Appl. Phys. Lett.*, 116(10), March 2020. ISSN 0003-6951. doi:10.1063/1.5137900. (Cited on page 68).

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

In my time as a graduate student, I have been lucky to have benefited from the support of a great community of people.

First and foremost, I owe a lot to my thesis advisor David Schuster, who has always provided a great wealth of ideas, and importantly helped me grow as a scientist, teaching me what questions to ask, to see the benefits of the devices we created, and how to step back from details and simplify experimental problems. I'm also very grateful to Dave for being extremely supportive in trying out novel fabrication ideas, as well as the chance to work with cutting edge millimeter-wave and microwave electronics– it feels like it has been a unique opportunity to play with with technology of the future.

I would also like to thank Jon Simon for providing much needed fresh perspective, and his out of the box thinking, which has helped provide new context and fresh ideas for our millimeter-wave experiments.

I would like to thank Aziza Suleymanzade, for introducing me to millimeter waves and being a great partner in our millimeter-wave experiments, and for always showing me new ways to think about the implications of our experiments. It has also been a pleasure to work with Ash Kumar, as well as Jasmine Kalia and Lavanya Teneya who have helped think about and provide high-frequency knowledge (and sometimes hardware!). It has also been a pleasure to work with Shannon Harvey and Wendy Wan, who helped give our experiments continued momentum.

I have leared a great deal by working with Kan-Heng Lee, who has helped broaden my understanding of fabrication chemistry, and has been a joy to run experiments with. Peter Duda has been instrumental to exploring some of the intricate fabrication challenges, as well as the Pritzker Nanofabrication Facility community: I have enjoyed discussions with Joel Grebel and Anna Mukhortova, and owe a lot to Samuel Khaler for maintaining the etching equipment. I am very grateful to Abigail Shearrow and Akash Dixit, who have made

the cleanroom an enjoyable experience, along with Nate Earnest and Carolyn Zhang: all of whom collectively turned the pit into a wonderful environment.

The Schuster lab has proven to be an incredible community. Srivatsan Chakram helped me become less of a villager at superconducting qubits. Gerwin Koolstra has helped inspire immaculate attention to detail, and made the sub basement a more welcoming place. I'm particularly grateful to have the oppurtunity to work with Andrew Oriani, who aside from providing cryogenic expertise, and invaluable engineering and machining knowledge, has helped hold all of us to high standards in experiments, and has helped me see our experiments in the context of fundamental physics. I have greatly enjoyed our conversations (both serious and silly) with Gabrielle Roberts, who has also helped define our lab community and make it a more fun place. I'm also grateful to Brennan Dizdar and Chris Wang who have helped keep the Schuster lab community alive in Chicago.

To the greater Uchicago research community- Andrew Lingenfeldter has been particularly nice and helpful in growing our theoretical understanding of nonlinear systems. Samuel Whitley has been a valuable bridge between microwave and spin qubits, and also taught me a lot about visual presentation. I'd also like to thank the Anderson lab for providing a second friendly community.

I would very much like to thank my wife Sophie Anferov, not only for being a wonderful partner, but also for helping me understand the chemistry at play throughout our research, as well as reading through my onslaught of drafts. None of this would have been possible without my family who have supported me throughout everything, along with Miso (who is always allowed to steal my chair), and finally my friends: particularly my fellow herbs (Alex and Federico), as well as Allison and Percy, all of whom have helped keep life outside of lab fun.

ABSTRACT

To continue achieving ever faster computation speeds, future computer processors may need to increase their operating frequency to achieve clock speeds beyond several GHz. Quantum computing offers an alternate approach by leveraging quantum mechanical superposition to make each clock operation more efficient, allowing the processor to solve certain problems much more efficiently. Current quantum processors operate slower than their classical counterparts, with the fastest quantum operations at microwave frequencies and utilizing superconducting artificial atoms (qubits)–a promising platform for quantum experiments studying light-matter interactions in the strong coupling regime. Increasing qubit frequency to the millimeter-wave range ($\sim$100 GHz) offers a straightforward way to increase quantum computing speed for any qubit design. Crucially, millimeter-wave frequencies also have greatly reduced sensitivity to thermal noise, and whereas microwave qubits require extremely low temperatures ($<$50 mK) and isotopic enrichment of ^{3}He and ^{4}He in order to reduce sources of decoherence, millimeter-wave qubits can operate at significantly higher temperatures near 1 K. These temperatures are achievable using simpler methods such as direct ^{4}He evaporation, which translates to orders of magnitude higher cooling power. This is transformative for scaling up superconducting quantum processors by significantly increasing the number of qubit control channels possible in a single cryostat, enabling direct integration of qubits with superconducting digital processors, and allowing for more energy efficient possibilities for quantum communication between cryostats. In this thesis, we introduce millimeter-wave superconducting devices as a platform for quantum experiments. We develop a robust niobium trilayer Josephson junction with improved quantum coherence properties capable of operating at higher frequencies and temperatures than conventional aluminum junctions. Based on this technology we explore the thermal resilience of qubits with higher and higher frequency, finally demonstrating a 72 GHz millimeter-wave qubit cooled entirely with ^{4}He.

CHAPTER 1

INTRODUCTION

For the last 120 years, classical computing has scaled exponentially [120], most recently visualized by the number of transitors in circuit doubling every two years (referred to as Moore's law). As a result, current processors also get denser with each year by reducing transistor size. However, as shown in Figure 1.1, processor clock speed has plateaued around 2005, and has remained approximately constant since as processsors reached the limits of thermal dissipation. Computation speed has meanwhile continued to increase at the expense of transitor count through synchronous operation, and increased instruction count per clock cycle. Current commercial transistors have reached dimensions of 20 nm [18], and are approaching fundamental limits to transistor size, making heat management increasingly difficult.

An alternative approach takes advantage of the minimal dissipation in superconductors,

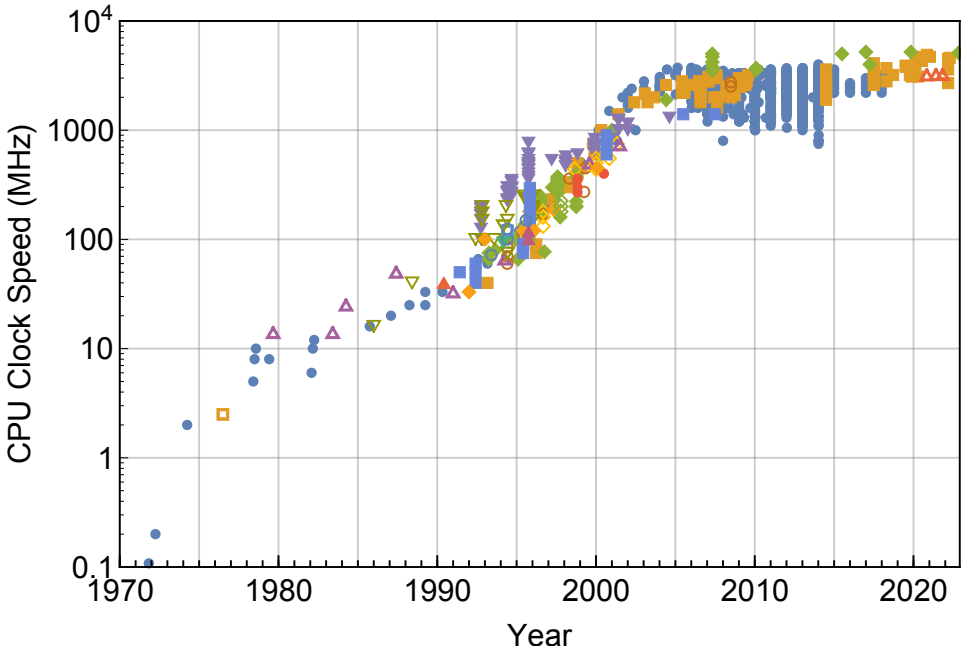

Figure 1.1: Processor clock speed in MHz over time. After an initial exponential increase, the clock speeds plateau after 2005 reaching a limit of 4-5 GHz. (manufacturers not labelled for brevity). Data from [51, 52]

in which currents can flow without resistance. By manipulating quantized units of flux, superconducting processors are capable of immense clock speeds above 100 GHz [226, 237] and have proven invaluable for high-speed computation tasks such as real-time signal processing [102]. While the future may show the full speed benefits of superconducting digital logic have in a large-scale classical processor, high frequencies have already proven useful for high-speed communications [244], capable of reaching data transfer speeds above a terabyte per second [105]. As it stands however, increasing transistor count has proven to be a more efficient and less costly route to increasing computation speed.

Another direction for improving computation speed is through the development of quantum computing, which leverages quantum mechanical superposition of information in a single bit to perform much more complex operations. This special behavior has the possibility to allow quantum computers to solve different types of problems that would be otherwise incredibly challenging to do via traditional methods. Interestingly, some of the original quantum experiments utilized high frequency atomic transitions [35], precise control over miniaturized circuits along with the availability of microwave components popularized the microwave frequency for superconducting quantum devices [111]. These microwave quantum bits (qubits) are a promising platform for realizing a quantum processor, and in part due to the large interaction strengths available in superconducting circuits they have some of the fastest operating speeds[1], achieving effective clock speeds up to 5 MHz (when taking into account 25 ns qubit gates [15] and 100 ns measurement and reset [218]). By comparison, other promising quantum computing platforms such as neutral atoms or trapped ions [36] rely on weaker interactions, resulting in comparatively slow gate speeds even compared to classical computers. Solid state spin qubits [212], use DC pulses to perform gate operations comparable to microwave qubits, however require fairly slow measurement and preparation protocols, though they do have high frequency transitions.

1. See Ref. [139] for a colloquial discussion of qubit speeds for different quantum processor architectures.

The strengths of quantum computing lie not in their speed but rather in the efficiency of each gate operation, which thanks to quantum mechanics can provide exponential improvements [206] for particular problems. So in the immediate future the key to realizing the full potential of quantum computing lies in increasing the interactions a qubit can perform before the losing its quantum information [55]. Looking towards the future however, increasing the operating clock speed of a processor is still a straightforward method to speed up computation speeds– so increasing frequencies of a quantum processor beyond the microwave range can offer even faster computation while still taking advantage the benefits of quantum computing. At the same time higher frequency quantum devices also allow quantum experiments to explore new ranges of frequencies, energy scales and interactions.

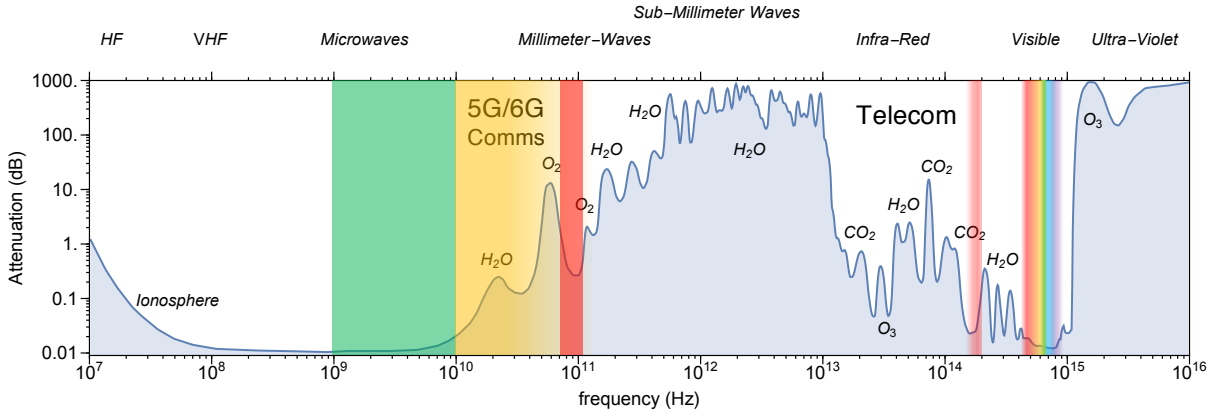

Figure 1.2: Electromagnetic spectrum in terms of frequency, along with atmospheric attenuation which highlighting regions with low attenuation where communication technology helped develop instrumentation and components.

These higher frequencies from 30–300 GHz are referred to as millimeter-waves: a fascinating regime between microwaves and optics, sharing some aspects of both frequency ranges. Aside from high-speed communication, millimeter-waves are also used extensively in astronomy in studying the cosmic microwave background [70, 83, 172]. This frequency range is particularly alluring for quantum experiments due to the higher energy scales of millimeter-wave photons, which directly translates to lower sensitivity to environment noise,

be it thermal photons, or cosmic radiation. As shown in figure 1.3 an increase in frequency exponentially reduces the number of thermal photons in the system. In this way, translating superconducting quantum hardware to higher frequencies directly reduces qubit dephasing from thermal photons [43] and reduces the need for artificial cooling methods [84, 135, 233, 245] to reset qubits to their ground state before experiments begin.

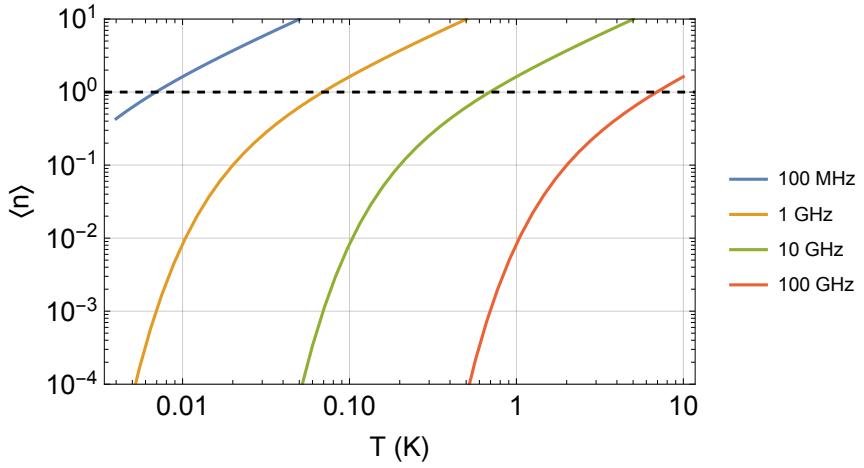

Figure 1.3: Average number of photons in a harmonic oscillator as a function of temperature, plotted for varying resonant frequencies.

Most importantly, higher frequencies increase the operating temperatures for the quantum system, as shown in Figure 1.3. This is extremely significant since the cryogenic refrigeration power scales with temperature squared. In other words a factor of 10 increase in qubit frequency proved 100 times more cooling power! In this way increased qubit frequencies could help managing the heat load from expanding numbers of qubit control lines [117] required for scaling up current superconducting quantum processors beyond hundreds of qubits. Thermally-resilient qubits could also reduce hardware overhead for microwave quantum links [181] (which currently can require a third dilution refrigerator for cooling the interconnect [136]), and could enable direct integration with superconducting digital logic [126, 131, 147], forming an integrated hybrid classical-quantum processor.

Superconducting qubits have also found uses in hybrid systems, coupling different types

of quantum emitters [44, 254]. Many of these experiments pose additional challenges by exposing qubits to magnetic fields or direct optical illumination [96, 119, 156, 240], so more resilient qubits could help reduce hybrid experiment complexity and improve performance, and enable even more diverse applications. Additionally, qubits built from higher-T_c superconductors are less sensitive to decoherence from quasiparticles [144], and have increased nonequilibrium quasiparticle recombination rates [125], enabling higher repetition rates for quantum transduction experiments.

Our aim is thus to increase frequency of microwave superconducting qubits (typically 4–8 GHz) by a factor of 20 into the millimeter-wave range (70–100 GHz). To achieve this we will need to realize a strongly nonlinear quantum system with low enough loss to be resolvable at the single photon level.

1.1 Thesis Organization

This thesis is organized as follows.

In Chapter 2 will first explore the key superconductor properties that underpin the workings of our devices.

In Chapter 3 will then discuss linear circuit elements in the context of millimeter-wave frequencies, with particular attention on resonators, which we will use throughout the rest of the thesis.

Following this in Chapter 4 we will build on our understanding of linear systems and use a general language to describe nonlinearity, focusing on the requirements for a millimeter-wave qubit.

In Chapter 5 we use devices with kinetic inductance to explore weak nonlinearity at millimeter-wave frequencies and high temperatures.

We discuss methods of increasing this nonlinearity strength in Chapter 6 with high-kinetic inductance microwave resonators.

In Chapter 7 we explore niobium trilayer Josephson junctions as a stronger source of high-temperature high-frequency nonlinearity, verifying their behavior in the established microwave qubit platform (1–6 GHz).

We begin to scale the qubit technology up in frequency in Chapter 8, where we explore the thermal resilience of K Band qubits (up to 24 GHz).

In Chapter 9 we revisit the millimeter-wave measurement and coupling methods from earlier, and demonstrate that millimeter-wave resonator loss can be comparable to conventional microwave devices.

Finally, in Chapter 10 we demonstrate a millimeter-wave qubit with single-photon resolved level transitions, and characterize its coherence properties in the time domain, fulfilling our goal.

CHAPTER 2

SUPERCONDUCTORS AND THEIR PROPERTIES

For alternating currents, dissipation is a significant problem as the frequency of a signal increases. In particular, due to the skin effect, higher frequency current is carried by an increasingly thin layer of the metal surface, making microwave and especially millimeter-wave circuits incredibly sensitive to surface resistance. This is particularly problematic for quantum circuits, for which we would like to eliminate every last bit of dissipation in order to preserve as much coherence as possible. For this reason, superconductors have been particularly transformative for quantum circuits since they have zero DC electrical resistance, and only tiny amounts of frequency-dependent loss.

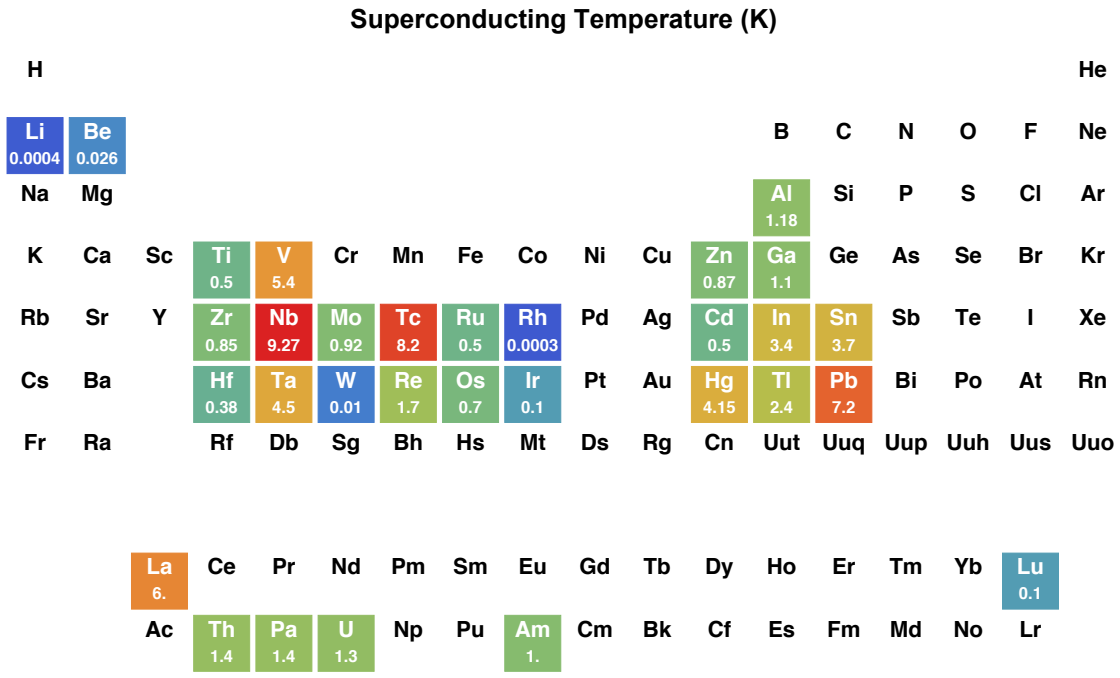

Figure 2.1: Periodic table highlighting elements with superconducting transitions (labeled in Kelvin) and colored by transition temperature T_C. Data taken from [68].

When cooled below a particular critical temperature T_c, many metals (and alloys) exhibit superconductivity as their electrons condense into a Bose-Einstein condensate of Cooper

pairs [48]. The many electron pairs can collectively be described by a single wavefunction with a well defined phase ϕ: this is a powerful effect that gives us a macroscopic quantum degree of freedom, and allows us to build and easily manipulate artificial quantum systems. Despite being a very special property, superconductivity at ambient pressure is surprisingly[1] common among the metallic elements [68], as shown in Figure 2.1. Of these, the group V elements stand out since their transition temperatures are uniformly higher. Already we see that niobium (Nb), which has the highest transition temperature for a pure element, is a particularly tantalizing material for its superconducting properties, and because of this shows up frequently in this work.

Material	λ_0 (nm)	ξ_0 (nm)	$2\Delta_0/kT_c$	T_c (K)
Al	16	1500	3.4	1.2
In	25	400	3.5	3.3
Sn	28	300	3.6	3.7
Pb	28	110	4.1	7.2
Nb	50	39	3.7	9.2
Nb_3Sn	50	5	4.4	18

Table 2.1: Superconducting properties, ordered by transition temperature T_c, for several popular superconducting metals and alloys. Data from [16]

Of course as with other properties, superconducting characteristics of pure materials can also be enhanced by alloying them with other elements. In particular, nitrogen can enhance the superconducting properties of metals: notably niobium nitride [] and titanium nitride [], both of which have significantly increased critical temperatures than the base metal, as shown in Table 2.1. Superconductivity is found in even more complex materials, such as in cuprates, nickelates and ferrites with transition temperatures exceeding liquid nitrogen temperatures [68]. While some of these materials have actually been explored as millimeter-wave resonator materials [20], in this work we will only consider conventional superconductors with Cooper

1. You don't normally think of actinides like uranium as a superconductor, but the transition temperature of one of its metallic phases is higher than that of aluminum!

pairs formed through s-wave pairing, specifically materials that we can easily deposit and pattern into circuits.

The nature, and microscopic mechanisms of superconductivity are fairly extensive topics for which Refs. [16, 224] are particularly useful. In the following sections we will cover a few particular properties of superconductors that are relevant to high-frequency circuits.

2.1 Magnetic Field Explusion

2.1.1 Drude Conductivity and London Equations

For normal metals, resistance arises from dissipation of conduction electrons through scattering events. Denoting the average time between scattering as τ, the Drude mode[16] predicts the conductivity of a material:

$$\sigma_n = \frac{n_n e^2 \tau}{m(1+\omega^2\tau^2)} - i\frac{n_n e^2 \omega\tau^2}{m(1+\omega^2\tau^2)} \tag{2.1}$$

Where the conduction carrier density n_n in this case is the free electron density (for most metals on the order of $\sim 10^{29}$), e is the carrier charge and m is the charge carrier mass[2]. τ is typically on the order of 10^{-14}s, so even at 100 GHz ($\omega \sim 10^{11} s^{-1}$) $\omega\tau$ is negligible. This simplifies the expression for conductivity:

$$\sigma_n \simeq \frac{n_n e^2 \tau}{m} \tag{2.2}$$

In this model, a superconductor would have no scattering, so $\tau \to \infty$. The superconducting conductivity is then:

$$\sigma_s = \frac{n_s e^2}{\frac{m}{\tau} + m\omega^2\tau} - i\frac{n_s e^2 \omega}{\frac{m}{\tau^2} + m\omega^2} \simeq 0 - i\frac{n_s e^2}{m\omega} \tag{2.3}$$

2. For a superconductor, charge is carried by pairs of electrons, so $m = 2m_e$ and $e = 2e$

Thus, if we define a normal σ_1 and complex σ_2 conductivity from $\sigma_s = \sigma_1 + i\sigma_2$, then for a superconductor we have $\sigma_2 = n_s e^2/m\omega$. From ohms law ($J = \sigma E$) applied to an electric field oscillating at ω we arrive at the first London equation:

$$\frac{\partial}{\partial t}\vec{J} = -i\frac{n_s e^2}{m\omega}\frac{\partial}{\partial t}\vec{E} = \frac{n_s e^2}{m}\vec{E} \tag{2.4}$$

Taking the curl of both sides of Equation 2.4 and combining with the Maxwell equation ($\nabla \times \vec{E} = \partial\vec{B}/\partial t$) yields the second London equation:

$$\frac{\partial}{\partial t}\nabla \times \vec{J} = -\frac{n_s e^2}{m}\frac{\partial \vec{B}}{\partial t} \tag{2.5}$$

$$\nabla \times \vec{J} = -\frac{n_s e^2}{m}\vec{B} \tag{2.6}$$

2.1.2 Meissner Effect

By taking the curl of Ampere's law ($\nabla \times \vec{B} = \mu_0 \vec{J}$) we obtain:

$$\nabla \times \nabla \times \vec{B} = \mu_0 \nabla \times \vec{J} \tag{2.7}$$

$$-\nabla^2 \vec{B} = \mu_0 \nabla \times \vec{J} \tag{2.8}$$

We can combine this result with Equation 2.6 to obtain:

$$\nabla^2 \vec{B} = \frac{\mu_0 n_s e^2}{m}\vec{B} = \frac{1}{\lambda_L^2}\vec{B} \tag{2.9}$$

Where the eigenvalue of the equation is defined as the London length $\lambda_L = \sqrt{m/\mu_0 n_s e^2}$. This equation highlights the fact a uniform non-zero magnetic field cannot exist within a superconductor: this magnetic field expulsion property is known as the Meissner effect. To satisfy Equation 2.9, magnetic fields inside a superconductor must be of the type $\vec{B}(x) = \vec{B}_0 e^{-x/\lambda_L}$. Thus, an applied external magnetic field will only penetrate a supercon-

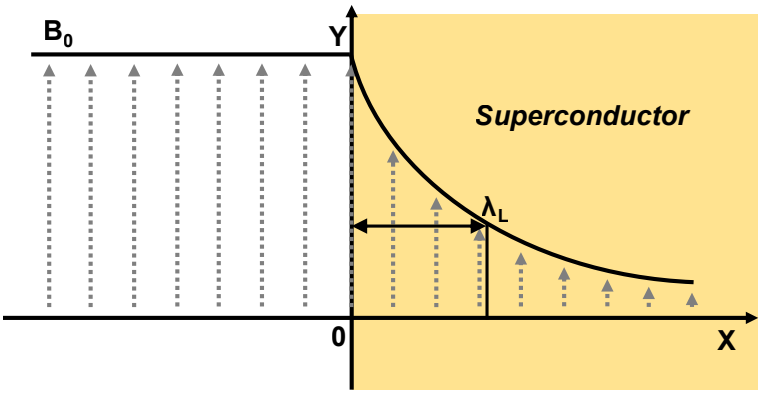

Figure 2.2: Illustration of the Meissner Effect. An external magnetic field outside of a superconductor results in a magnetic field that decays exponentially on the length scale λ_L.

ductor by a distance λ_L, as illustrated in Figure 2.2. For this reason λ_L is also called the London penetration depth. Table 2.1 lists the zero-temperature London lengths for various superconductors.

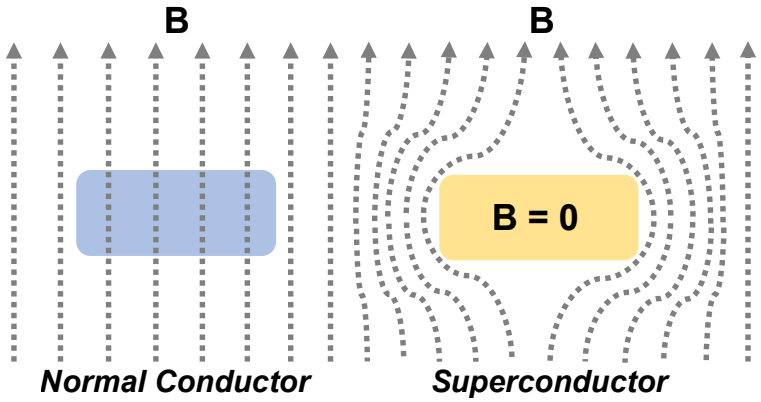

Figure 2.3: The Meissner Effect in action: an external applied magnetic field B, which would penetrate a normal conductor is expelled from a superconductor, inside which $B = 0$ (other than a small distance λ_L from the surface).

Practically, this means that the magnetic field inside a superconductor is zero except for a thin shell of thickness λ_L on ths surface. One result of of this effect is that a superconductor will expel an applied external magnetic field, as shown in Figure 2.3. It follows that the surface (more accurately a shell of thickness of λ_L) of the superconductor must be carrying

a circulating current that exactly opposes this magnetic field.

2.2 Kinetic Inductance

2.2.1 Currents in a superconducting wire

Previously we determined that an magnetic field applied on a superconductor generated an opposing surface current on a shell of thickness λ_L. We will find that the reverse also holds true. Consider a wire with circular cross section, with radius R carrying a current I, as shown in Figure 2.4. In the normal conducting case, Ampere's law means that a homogeneous current density J_n would yield a magnetic field inside the wire $B(r) = \mu_0 r I / 2\pi R^2$. However as a qualitative argument, we know that if $R \gg \lambda_L$, then we must have $B = 0$ inside most of the wire, except for a small shell on the surface.

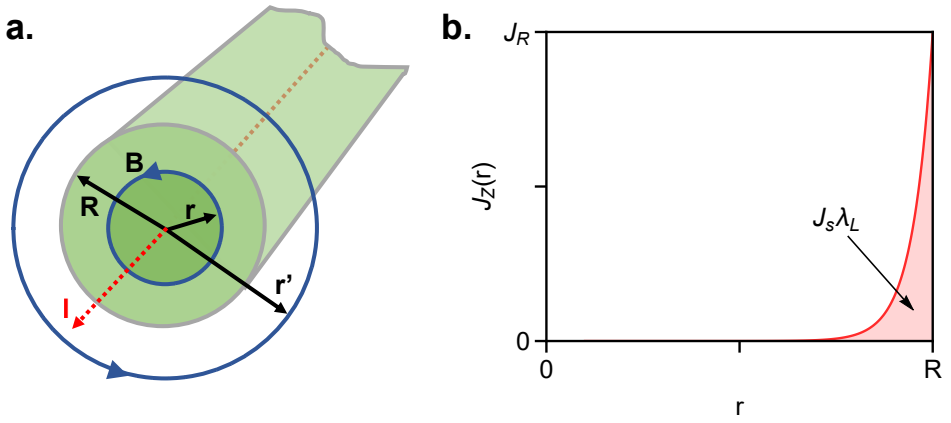

Figure 2.4: A cylindrical wire carrying a current generates a magnetic field inside the wire. For a normal conductor this current depends linearly on radius, however for a superconductor the current is localized on the surface in order to satisfy the London equations.

Indeed if we solve Equation 2.9 in cylindrical coordinates, we can approximate the solution for the magnetic field (which is a Bessel J function) near the surface with the exponential form:

$$\frac{B_\theta(r)}{B_\theta(R)} = J_1(ir/\lambda_L) \approx e^{-(r-R)/\lambda_L} \tag{2.10}$$

Applying Ampere's law on the surface-approximated solution, we must have:

$$J_z(r) = \frac{1}{\mu_0}\nabla \times B(r) \approx \frac{I}{2\pi R\lambda_L}\left(1+\frac{\lambda}{r}\right)e^{(r-R)/\lambda_L} \tag{2.11}$$

In other words, like the magnetic field, the supercurrent density is exponentially localized within a London length of the surface! As suggested by its prefactor, integrating Equation 2.11 across the entire cylindrical area of the wire gives exactly the same result as a fixed current density J_s over a shell of thickness λ_L on the surface S. [3]

$$I = \int_0^R 2\pi\rho J_z(\rho)d\rho = J_s\lambda_L S \tag{2.12}$$

2.2.2 Kinetic energy of a supercurrent

From the complex conductivity derived previously, we conclude the response of the superconducting charge carriers is purely reactive and not resistive. From Ohm's law this also suggests that current (or velocity of charge carriers) from an alternating electric field will be out of phase, as is the case in an inductor. For a supercurrent carried by some density of charge carriers n_s, we can write out the density of its kinetic energy:

$$\partial_V U_k = \frac{1}{2}n_s m v_s^2 \tag{2.13}$$

Where we have defined the velocity of superconducting charge carriers v_s. This velocity is related to the current density $J_s = n_s e v_s$, which allows us to rewrite the energy density in terms of the London length:

$$\partial_V U_k = \frac{1}{2}\frac{m}{n_s e^2}J_s^2 = \frac{1}{2}\mu_0\lambda_L^2 J_s^2 \tag{2.14}$$

3. It is fairly straightforward (and maybe simpler than our example) to convince yourself that this is also the case in rectangular coordinates as well.

Now if we equate this with the inductive energy for a wire section of length l, we obtain

$$\frac{1}{2}L_k I^2 = \frac{l}{2}\mu_0 \lambda_L^2 \int_A J_s^2 \tag{2.15}$$

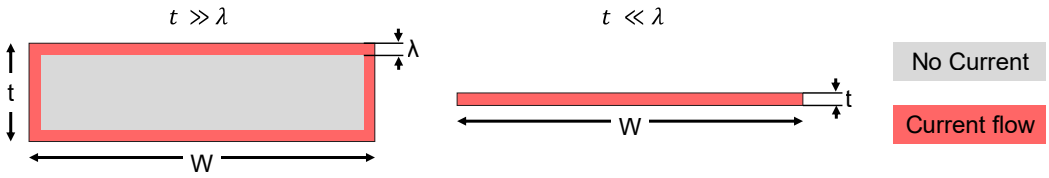

Figure 2.5: Cross sections of a rectangular wire in the thick and thin-film limits (with respect with the London length). In the thin-film limit, the entire area of the wire carries the current, while in the thick-film limit only the surface does.

To proceed with this integral, we should consider the geometry of the current carrying conductor. Consider a rectangular wire of width W and thickness t in the limits shown in Figure 2.5. If the film is thick, or $t \gg \lambda$, then we $\int_A J_s = I/(2W + 2t)\lambda_L \approx I/2W\lambda_L$. However if the film is thin, then the current area is simply that of the wire, so $\int_A J_s = I/Wt$. Thus, we can solve for the kinetic inductance L_k in these two limits:

$$\frac{L_k}{l} = \begin{cases} \mu_0 \frac{\lambda_L}{2W} & t \gg \lambda_L \\ \mu_0 \frac{\lambda_L^2}{Wt} & t \ll \lambda_L \end{cases} \tag{2.16}$$

A practical concept for films of uniform thickness t is sheet inductance $L_\square$ (similar to sheet resistance $R_\square = \rho/t$). Knowing this property allows us to quickly estimate the inductance of a wire from its width w and length l from $L = L_\square l/w$. The kinetic sheet inductance in both limits is then:

$$L_\square(H/\square) = \begin{cases} \mu_0 \frac{\lambda_L}{2} & t \gg \lambda_L \\ \mu_0 \frac{\lambda_L^2}{t} & t \ll \lambda_L \end{cases} \tag{2.17}$$

The thin-film case is particularly interesting, since we see an enhancement of λ_L/t in the kinetic inductance. For extremely thin films and materials with high London lengths this scaling factor can become quite significant!

2.2.3 Total Inductance

Thus far we have been sneaky by only considering the kinetic energy of the electrons. In reality the energy of the charge carriers also contains a magnetic (or geometric) term [154]:

$$\partial_V U_k = \frac{1}{2} L_m I^2 + \frac{1}{2} n_s m v_s^2 \tag{2.18}$$

So the total inductance for a thin film will be $L_T = L_m + L_k$.

We can also more accurately integrate the current density across the rectangular wire for dimensions in between the two limits $t \gg \lambda_L$ and $t \ll \lambda_L$. A full analytical expression can be derived for both [42, 57, 154], from which we can obtain:

$$L_{m,\square} = \frac{\mu_0 \lambda_L}{4} \left(\coth \frac{t}{2\lambda_L} - \left(\frac{t}{2\lambda_L} \right) \operatorname{csch}^2 \frac{t}{2\lambda_L} \right) \tag{2.19}$$

$$L_{k,\square} = \frac{\mu_0 \lambda_L}{4} \left(\coth \frac{t}{2\lambda_L} + \left(\frac{t}{2\lambda_L} \right) \operatorname{csch}^2 \frac{t}{2\lambda_L} \right) \tag{2.20}$$

Combined, we arrive at the simplified expression for sheet inductance of a wire [57]:

$$L_\square = \frac{\mu_0 \lambda_L}{2} \coth \left(\frac{t}{2\lambda_L} \right) \tag{2.21}$$

In the limit $t \ll \lambda_L$, the function approaches $\lambda \coth(t/2\lambda) \to \lambda^2/t$ and we recover the thin-film expression in Equation 2.17, while in the thick limit $t \gg \lambda_L$, the coth term goes to 1 leaving the thick-film expression in Equation 2.17.

An interesting note from this result is the source of inductance (inductive or kinetic) for varying thickness limits. In Figure 2.6 we plot L_k and L_m as a function of $t/2\lambda_L$. For high

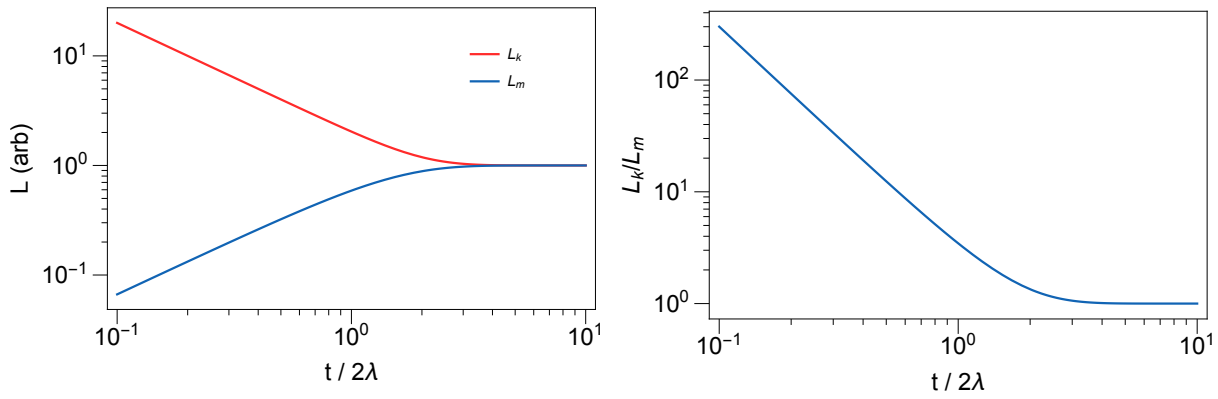

Figure 2.6: Kinetic and magnetic (or geometric) inductance, plotted as a function of film thickness relative to the London length. While for thick films the two are comparable, for thin films we find that kinetic inductance becomes significantly higher.

values (or the thick limit) the two values are similar, however as we dive deeper into the thin film limit, not only does the overall inductance rise, but L_k becomes significantly more dominant. This highlights the power of film thickness for generating incredible amounts of inductance with a simple rectangular wire!

2.2.4 Nonlinearity in kinetic inductance

One very important property of kinetic inductance is the fact that the inductance outlined above actually increases with the current carried in the wire, or in other words the superconductors exhibit a nonlinear response. For low temperatures $T \ll T_c$ we can expand the inductance in terms of current [269]:

$$L_k(I) \simeq L_k(0)\left(1 + \frac{I^2}{I_*^2}\right) \tag{2.22}$$

To satisfy time-reversal symmetry the inductance can't have first order dependence on current, so we have a quadratic dependence on current. The scale of the nonlinearity is set by I_*, which is expected to be the same order as the depairing current [57, 269]. Thus the

nonlinear portion of inductance will be:

$$\delta L_k = \frac{L_k(0)I^2}{I_*^2} = \frac{L_k(0)J^2}{J_*^2} = \frac{L_k(0)\kappa_* J^2}{J_c^2} \tag{2.23}$$

Which is written in terms of the depairing current (defined in terms of units we will examine later)

$$J_*^2 = \frac{J_c^2}{\kappa_*} = \frac{1}{\kappa_*}\left(\frac{\pi}{\hbar}\sigma_n N(0)\Delta(0)^3\right) \tag{2.24}$$

The empirical proportionality constant κ_* is of order unity (found to be between 0.6–0.95 [14, 269]).

The important thing to note is that the energy scale of the nonlinear inductance is set by a combination of the linear inductance and material properties: leveraging both yields a simple nonlinearity that scales with I^2. We will be making good use of this property!

2.3 Vortices

2.3.1 Wavefunction and Velocity

Thus far we have treated the supercurrent charge carriers somewhat ambiguously with an effective charge and mass. In reality, superconductivity is a result of electrons condensing into Cooper pairs [48], so the charge carriers we have discussed thus far actually have mass $2m$ and charge $2e$. More importantly, by condensing into pairs, the Fermionic electrons are transformed into Bosons: the magic of this property is that we can now collectively describe the condensate with a single wavefunction ψ which has a phase ϕ:

$$\psi = \psi_0 e^{i\phi}, \quad |\psi|^2 = n_s/2 \tag{2.25}$$

Here, the expectation of the wavefunction is defined as the electron density. In this picture, a phase gradient ($\nabla\phi$) will result in a supercurrent. If we combine this with the effects of

any applied magnetic field ($\nabla \times A = B$) in the quantum mechanical picture, we have the momentum, and thus velocity of the Cooper pairs [115]:

$$\vec{J_s} = n_s e \vec{v_s} = -\frac{n_s e^2}{m}\left(\vec{A} - \frac{\hbar}{2e}\nabla\phi\right) \tag{2.26}$$

This is actually derived from the Landau-Ginzburg theory of superconductivity [73] (touched on in next section) which establishes that the superconducting pairs interact on a length scale ξ. In other words, their wavefunction can't vary on a smaller scale than ξ. Thus since $\nabla\phi < 1/\xi$, the velocity must stay below

$$v_s < \frac{\hbar}{2m\xi} \tag{2.27}$$

As we will explore later, this also connects the coherence length to an energy scale.

2.3.2 Flux Quantization

Having established a general form for the Cooper pair wavefunction, we can quickly demonstrate how a Superconductor can quantize magnetic flux. Consider a loop of superconductor enclosing a region E penetrated by a magnetic field, as shown in Figure 2.7 (This center region could be either vacuum, insulator or normal metal). For simplicity, let us assume the superconductor dimensions are greater than λ_L and the core doesn't magnetize ($\mu = \mu_0$).

We pick a loop l of radius ρ **inside** the superconductor covering area S. Integrating the current around this loop gives (dropping a few prefactors):

$$\oint_l \left(\vec{A} - \frac{\hbar}{2e}\nabla\phi\right) dl = \int_0^{2\pi} 2\pi\rho d\theta \left(\vec{A}(\rho,\theta) - \frac{\hbar}{2e}\nabla\phi(\theta)\right) \tag{2.28}$$

$$= \int_S \nabla \times A dS - \frac{\hbar}{2e}(\phi(2\pi) - \phi(0)) \tag{2.29}$$

Since the phase ϕ can't be discontinuous, we must have $\phi(2\pi) - \phi(0) = 2n\pi$ for integer n.

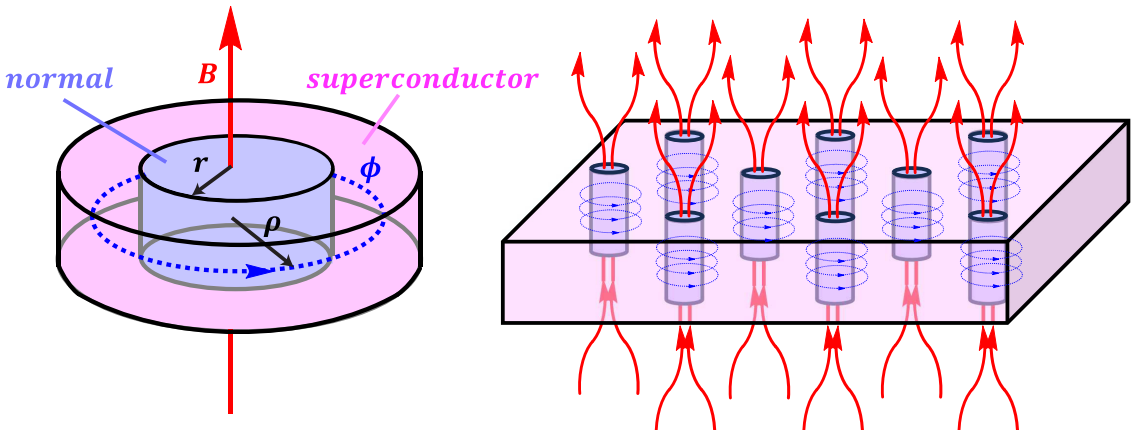

Figure 2.7: Superconducting vortex.

Now since we picked a loop inside the bulk of the superconductor, we must have $\vec{J}_s = 0$. Combining this with the definition of magnetic flux through the center $\oint_S BdS = \oint_E BdE = \Phi$, the above simplifies to:

$$0 = \Phi - \frac{\hbar}{2e} 2n\pi \tag{2.30}$$

Thus we have a quantized unit of magnetic flux through a superconducting loop:

$$\Phi_0 = \frac{h}{2e} \tag{2.31}$$

2.3.3 Vortices in the superconductor

Having established the magnetic flux inside the superconducting loop, we can repeat the above process on the **surface** of the superconductor. Lets define the inside radius of the hole as r. On the inside of the superconducting loop, we must have:

$$\oint J_s(r) = n\Phi \tag{2.32}$$

$$\oint e n_s v_s(r) = n\pi r^2 B \tag{2.33}$$

Rewriting the magnetization field $H = B/\mu_0$, we find

$$v_s(r) = \frac{n\pi r^2 H}{2\pi r \lambda_L e n_s \mu_0} \tag{2.34}$$

Evidently, at some applied magnetic field, our Cooper pairs will exceed the critical velocity in Equation 2.27 which would lead to a breakdown in superconductivity. Now let us consider the case where the core is infinitesimally thin, and the magnetic field penetrating it is just strong enough to induce the critical velocity in the electrons. We can rearrange in terms of the velocity, recovering familiar units:

$$\frac{H_c}{e\lambda_L \mu_0 n_s} \simeq \frac{\hbar}{2m\xi} \tag{2.35}$$

$$H_c \simeq \frac{\hbar}{2\xi e \lambda_L} = \frac{\Phi_0}{2\pi\xi\lambda_L} \tag{2.36}$$

Thus a sufficiently high magnetic field can penetrate a superconductor through tiny regions surrounded by a ring of current, or superconducting Abrikosov vortices. Our derivation of this was quite simplistic: in actuality the true critical field will be [115]:

$$H_{c1} = \frac{\Phi_0}{4\pi\lambda_L^2} \ln\left(\frac{\lambda_L}{\xi}\right) \tag{2.37}$$

Something we haven't touched on at all here is the fact that the critical field gets reduced for films thinner than λ_L [193]. This makes vortices even more prevalent in thin-film superconductors (which from the last section is exactly the kind of films we would like to work with). While a very neat effect, Vortices are actually quite problematic for microwave superconducting circuits since they present several opportunities for dissipation. Examining the simple picture of a infinitesimally thin vortex, nothing stops it from moving around in the horizontal plane if the superconductor is uniform. In fact flowing current in the horizontal plane of the superconducting sheet will apply a Lorentz force on each vortex, moving them

around[4].

This causes two primary issues. First, since the core of the vortex is no longer superconducting, a travelling vortex[5] must break Cooper pairs that enter into its center into normal electrons (or quasiparticles), and it follows that the vortex leaves behind it a wake of broken Cooper pairs. These take some time to recombine into cooper pairs [125], leading to an excess of quasiparticles in the superconductor [144]. Second, the vortex energy depends on the **local** superconductor properties. This means vortices are now able to interact and worse, scatter with variations in an inhomogenous superconductor, leading to a new kind of dissipation [174]! This is all not to mention direct inductive coupling to vortex motion which can also cause energy decay.

In this work, we will attempt to avoid thinking about the vortex problem as much as possible by reducing the magnetic fields penetrating our thin superconducting films. We accomplish this by using magnetic shielding (discussed later) to block out external magnetic fields (including Earth's magnetic field) and also simply avoiding any kind of circuit with loops that would be sensitive to global magnetic fields. Now the only vortices we should have to worry about are the ones generated by our own currents!

2.4 Coherence Length and Proximity Effect

So far our treatment of superconductors has been fairly simple. The Landau Ginzburg theory of superconductivity [73], which phenomenologically predicts a magnetic penetration depth λ (similar to what we wrote out above with the London equations), also takes into account a coherence length ξ setting the scale on which the superfluid density varies, which can help us explain a few more key effects. We will keep this section very brief, so references [16, 115, 224] are a good reference for more details.

4. Vortices can reach speeds of km/s!

5. If the travelling vortex needs to visit every Cooper pair, what is the optimal path it should take? ;)

Using the same superconducting charge carrier wavefunction ψ as in last section, the Landau Ginzburg differential equations dictate the behavior of ψ:

$$\alpha(T)\psi + \frac{\beta(T)}{2}|\psi|^2\psi - \frac{\hbar^2}{2m^*}\left(\nabla + \frac{2ie^*}{\hbar c}A\right)^2\psi = 0 \tag{2.38}$$

$$J_s = \frac{c}{4\pi}\nabla \times H = \frac{e\hbar}{2im^*}(\psi^*\nabla\psi - \psi\nabla\psi^*) - \frac{e^{*2}}{m^*c}\psi^*\psi A \tag{2.39}$$

Here $\alpha(T) \approx \alpha(T - T_c)$ and $\beta(T) \approx \beta$ are parameters used to describe the material. Similar to before, the wavefunction amplitude gives the charge carrier density $n_s = |\psi|^2 = -\alpha(T)/\beta(T)$. Taking the curl of Equation 2.39 gives us back the characteristic London length for magnetic fields in the superconductor:

$$\nabla \times J_s = -\frac{e^{*2}}{mc}|\psi|^2 H = -\frac{n_s e^{*2}}{2m^*}H = \frac{-1}{\mu_0\lambda_L^2}H \tag{2.40}$$

Now, consider Equation 2.38 with no magnetic field ($A = 0$). This yields:

$$\alpha(T)\psi + \frac{\beta(T)}{2}|\psi|^2\psi - \frac{\hbar^2}{2m^*}\nabla^2\psi = 0 \tag{2.41}$$

$$\psi = \frac{\hbar^2}{2m^*\alpha(T)}\nabla^2\psi \tag{2.42}$$

This gives us our second length scale, the coherence length ξ which governs the variation of ψ in space.

$$\xi(T)^2 = \frac{\hbar^2}{2m^*|\alpha(T)|} \approx \frac{\xi(0)^2}{(T_c - T)} \tag{2.43}$$

Later on we can define this length explicitly using BCS theory of weak coupling.

2.4.1 Superconductor Boundaries

Now consider Equation 2.38 at the interface between a superconductor and normal metal, which will have the boundary condition [53, 224]:

$$\left(\frac{\hbar}{i}\nabla - \frac{e^*}{c}A\right)\psi_{\hat{n}} = \frac{i\hbar}{b}\psi_{\hat{n}} \tag{2.44}$$

For some positive constant b. Thus if we ignore the vector potential, we must have a nonzero derivative at the boundary! By extension, the wavefunction must also be nonzero at the boundary, so ψ must effectively "leak" into the normal metal. Indeed, extrapolating ψ based on the derivative condition at $x = 0$ we can estimate that the effect will extend by some non-zero distance b into the normal metal, as illustrated in Figure 2.8.

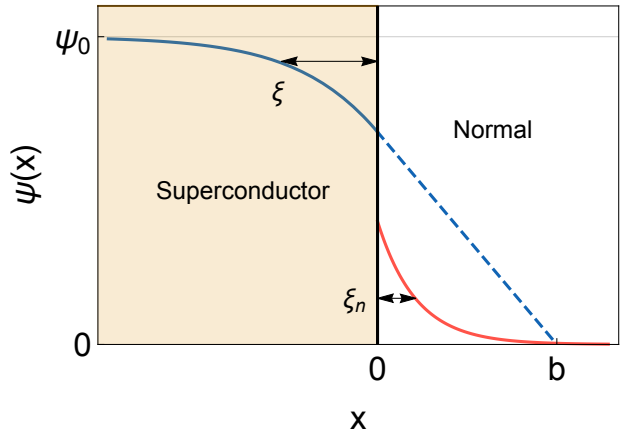

Figure 2.8: Illustration of the proximity effect at the boundary between a superconductor and normal metal, showing that superconductivity extends some distance into the normal metal.

As another way to think about what happens in the normal metal, since $|\psi|^2$ is nonzero at the boundary, we must have some fraction of electrons that are still paired: these will remain paired until they scatter in the normal metal. So the effective coherence length in the normal metal will depend on its mean free path l (similar to how regular conductivity is defined). Generally speaking, the wavefunction in the normal metal (in one dimension) will

have the form:

$$\psi(x) = \psi_0 e^{-x/\xi_n} \tag{2.45}$$

This normal coherence length ξ_n can be approximated in two limits [53, 224]:

$$\xi_n \approx \begin{cases} \frac{\hbar v_F}{2\pi kT} & l \gg \xi_n \text{ (clean)} \\ \left(\frac{\hbar v_F l}{6\pi kT}\right)^{1/2} & l \ll \xi_n \text{ (dirty)} \end{cases} \tag{2.46}$$

On the superconductor side, the positional dependence of ψ is well defined by Equation 2.38, which gives solutions of the form $\psi(x) = \tanh((x - x_0)/\sqrt{2}\xi)$, or in other words the length scale on which ψ changes is the superconducting coherence length ξ. In this case though, it means that while superconductivity extends into the normal metal, ψ also begins to be suppressed in the superconductor. Large proximity lengths ξ_n will reduce the amount ψ decreases at the boundary, but really we should make sure the superconductor is much thicker than ξ to minimize the impact on ψ inside the superconductor.

2.4.2 Aluminum Proximitized by Niobium

Let's consider a concrete example of a layer of aluminum (t_{Al}) in contact with a layer of superconducting niobium (t_{Nb}). Suppose the system is kept at 1.2 K (just above $T_{c,\mathrm{Al}}$, so we will consider Al a normal metal). From Table 2.1 we find that $\xi_{\mathrm{Nb}}(1.2\text{ K}) \approx \xi_{0,\mathrm{Nb}} = 39$ nm. tThis dictates that $t_{\mathrm{Nb}} \gg 39$ nm. Conservatively, this means that a >80 nm thick Nb layer should suffice. For aluminum, $v_F = 2 \times 10^6$ m/s and $l \sim 20$ nm. Thus at 1.2 K, we are in the clean limit and $\xi_n \sim 2\mu$m. So an aluminum film a few tens of nanometers thick will be completely proximitized by the niobium, and in the niobium ψ is virtually unaffected!

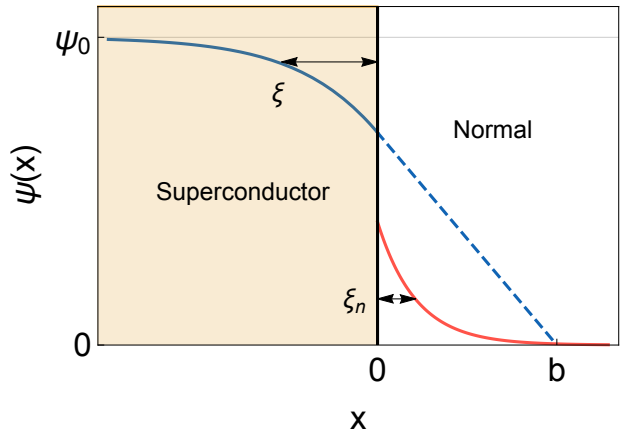

Figure 2.9: Illustration of the Josephson effect: two superconductors separated by a very thin barrier $d \ll \xi$ have two separate phases on either side

2.4.3 Josephson effect

Now instead of a superconductor near a metal, consider instead a superconductor separated from an identical superconductor by a thin insulating[6] barrier of thickness d, much thinner than the coherence length in the superconductor ξ. Treating the barrier as infinitesimally thin, we will assume the wavefunctions $\psi = |\psi|e^{i\phi}$ on either side will have differing values $\psi_1(=0) \neq \psi_2(x=0)$. Similarly to the boundary condition between a normal metal and superconductor, for two superconductors, we will have [53, 224]:

$$\left(\frac{\hbar}{i}\nabla - \frac{e^*}{c}A\right)\psi_1\Big|_{x=0} = a\psi_2\Big|_{x=0} \tag{2.47}$$

For some real constant a. To simplify the problem we will consider a gauge where $A = 0$. Consider the first superconductor at the boundary ($x = 0$): the supercurrent density will be

6. As it turns out you can recover a similar effect with more caveats with a conducting, ferromagnetic or even a different superconductor as the barrier material

given by Equation 2.39:

$$J_s(x=0) = \frac{\hbar e}{2im}\left(\psi_1^* \frac{\partial \psi_1}{\partial x} - \psi_1 \frac{\partial \psi_1^*}{\partial x}\right) \tag{2.48}$$

$$= \frac{\hbar e a}{2im}\left(\psi_1^* \psi_2 - \psi_1 \psi_2^*\right) \tag{2.49}$$

$$= \frac{\hbar e a n_s}{2im}\left(e^{i(\phi_2-\phi_1)} - e^{i(\phi_1-\phi_2)}\right) \tag{2.50}$$

$$= J_c \sin(\phi_2 - \phi_1) \tag{2.51}$$

Thus the current tunneling through the barrier is determined by the two phases of the wavefunctions on either side! From this we can write the total current through a Josephson junction as

$$I = I_c \sin(\varphi) \tag{2.52}$$

This nonlinear dependence will come in handy later.

2.5 BCS theory and Superconducting Energy Gap

2.5.1 Quasiparticle Excitations

In a conductor at low temperatures, Fermionic electrons fill the lowest energy states, with the highest of such being the Fermi energy ϵ_f. Working in terms of momentum, the highest energy electrons can be pictured as filling a layer on the Fermi surface $\xi_k = \epsilon_k - \epsilon_f$, where $\epsilon_k = (\hbar k)^2/2m$ is the energy of a particle with momentum $\hbar k$. A plot of ξ_k is shown with dashed lines in Figure 2.10. Below the superconducting critical temperature T_c, two electrons with exact opposite momentum experience an attractive pairing potential[7] W [16, 115, 224]. Since this interaction lowers the total energy, a pair of coupled electrons with momemtum k taken from the Fermi surface will result in an energy below the Fermi surface. We call this

7. Usually labelled V, but we call it W for clarity.

energy difference Δ, the superconducting energy gap.

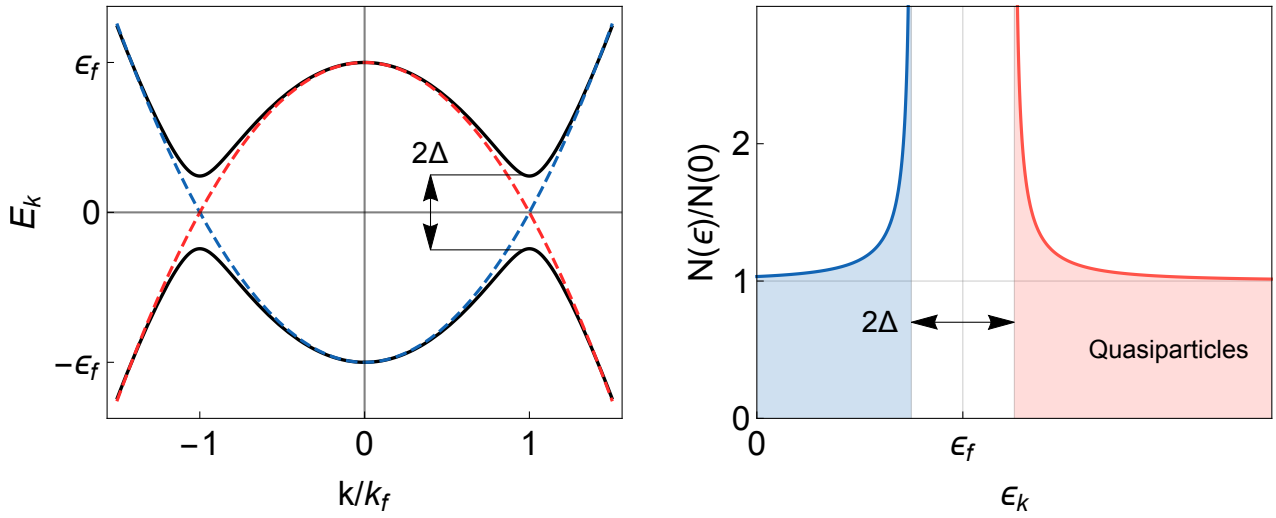

Figure 2.10: Left: Energy spectrum of unpaired electrons for a normal conductor (dashed) with hole and electron solutions for energy. In a superconductor (solid lines), the excitations are increased by a gap energy Δ. Right: Normalized density of states for quasiparticles and quasiholes, which is split around the Fermi energy by Δ, so forming either kind of excitation requires an energy of at least Δ.

Breaking a Cooper pair and returning to un-paired electrons or quasiparticles now requires energy added to the system. This modifies the excitations of quasiparticles, which now follows [115, 224]:

$$E_k = \pm\sqrt{\xi_k^2 + \Delta^2} = \pm\sqrt{(\epsilon_k - \epsilon_f)^2 + \Delta^2} \tag{2.53}$$

The modified spectrum is shown with sold lines in Figure 2.10. From this we see that the minimum excitation energy (of a single particle with $k = k_f$) is Δ, so the energy required to generate either a quasiparticle or quasihole is Δ. In this picture, superconducting electrons are in pairs with energy e_f, so to break a Cooper pair will require two excitations, or least 2Δ.

This gap is more apparent if we solve for the density of states [115] by solving Equation

2.53 for ϵ_k and differentiating with E_k:

$$N(E) = \frac{d\epsilon}{dE} = \pm \frac{N(0)E}{\sqrt{E^2 - \Delta^2}} \tag{2.54}$$

This is the density of states relative to the Fermi energy, and assuming $\Delta \ll E_F$ we can approximate $N(0)$ as the density on the fermi surface:

$$N(0) = \frac{mk_F^2}{2\pi^2\hbar^2} \tag{2.55}$$

Figure 2.10 shows the density of states for the two excitations, in which the density has effectively been pinched around the Fermi energy, leaving an apparent energy gap where no states are available between $\epsilon_f \pm \Delta$.

Going one step further, since quasiparticles are by definition Fermions again, we can use Fermi statistics ($f(E,T) = 1/(e^{\beta E} + 1)$) to calculate the density of quasiparticles.

$$N_{\text{qp}}(T) = 4\int_0^\infty f(E)N(E)dE = 4\int_{\epsilon_f - \Delta}^\infty \frac{1}{e^{E/kT} + 1}\frac{N(0)E}{\sqrt{E^2 - \Delta^2}} \tag{2.56}$$

The factor of 4 arises from each energy level accommodating two spins, and allowing for both electron and hole excitations [57].

2.5.2 The Energy Gap

So far we haven't said anything about the behavior of the superconducting energy gap. In fact it has a significant dependence on temperature: which in turn determines nearly every other superconductor property we've discussed thus far. In Bardeen Cooper Schrieffer's theory of superconductivity [48], pairing between electrons occurs on the (coherence) length scale:

$$\xi_0 = \frac{\hbar v_F}{2\pi k T_c} \tag{2.57}$$

This interaction between electrons is mediated by phonons, so the attractive potential W is only defined below a critical phonon frequency ω_c (of order the Debye frequency Θ_D so still quite high). A critical assumption is that the coupling is weak, in other words $kT \ll \hbar\omega_c$ [224]. This sets the limit of energy ranges we should sum over. Since the Fermi distribution $f(E_k, T) = 1/(e^{\beta E_k} + 1)$ gives the excitation probability of quasiparticles, it follows that the probability of Cooper pairs is $1 - 2f(E_k)$. The superconducting energy gap can then be defined from the following self-consistent equation:

$$\Delta = \sum_k (1 - 2f(E_k)) \frac{\Delta}{2E_k} = N(0)W \int_0^{\hbar\omega_c} \frac{\Delta}{\sqrt{\xi_k^2 + \Delta^2}} \tanh \frac{1}{2}\beta\sqrt{\xi_k^2 + \Delta^2} d\xi \tag{2.58}$$

Where $N(0)$ is the maximum quasiparticle density (occuring at zero temperature). Figure Figure 2.11 shows the temperature dependence of the superconducting gap, which has the same fundamental behavior for superconductors with weak coupling.

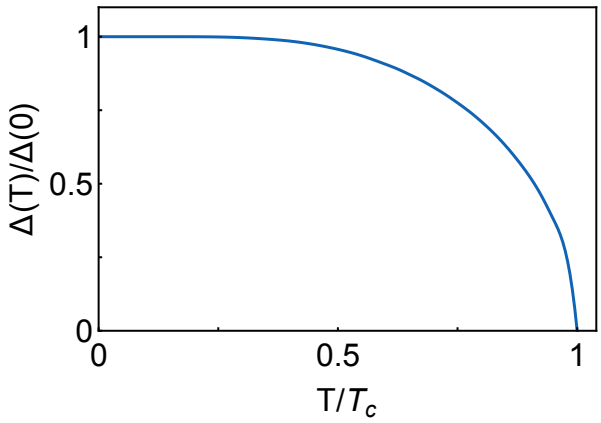

Figure 2.11: Exact solution for the temperature dependence of the superconducting energy gap plotted in reduced units with respect to T/T_c

For zero temperatures, we can ignore the tanh term completely. Thus we have a much

simpler expression for the zero-temperature gap:

$$\frac{1}{N(0)W} = \int_0^{\hbar\omega_c} \frac{d\xi}{\sqrt{\xi_k^2 + \Delta(0)^2}} \tag{2.59}$$

$$\Delta(0) \approx 2\hbar\omega_c e^{-1/N(0)W} \tag{2.60}$$

We can repeat the process for finite but low temperatures $T \ll T_c$ where the tanh term is still negligible:

$$\frac{1}{N(0)W} = \int_0^{\hbar\omega_c} \frac{d\xi}{\sqrt{\xi_k^2 + \Delta(0)^2}} \tag{2.61}$$

$$\Delta(0) \approx 2\hbar\omega_c e^{-1/N(0)W} \tag{2.62}$$

Above a critical temperature T_c, superconductivity disappears, so we must have the gap disappear also $\Delta(T_c) \to 0$ in order to return to the normal conductor model. This condition means $\epsilon(k, T_c) = \xi_k$, and we can solve a simpler integral:

$$\frac{1}{N(0)W} = \int_0^{\hbar\omega_c} \frac{1}{\xi} \tanh \frac{\xi}{2kT_c} d\xi \tag{2.63}$$

This yields a relation for critical temperature, which can be expressed as a function of the zero-temperature gap in Equation 2.62:

$$kT_c \simeq 1.13\hbar\omega_c e^{-1/N(0)W} = 1.76\Delta(0) \tag{2.64}$$

For most superconductors $2\Delta/kT_c$ is between 3.5–4 (see Table 2.1) so this reasonable accurate. Notably for niobium this ratio tends to be a bit higher.

2.6 Mattis-Bardeen Theory of Complex conductivity

2.6.1 Conductivity Integrals

Recalling Section 2.1.1, conductivity can be expressed as a real conductivity σ_n for normal electrons and a complex conductivity σ_s for superconducting Cooper pairs. Earlier we saw that at finite temperatures T below T_c, there are a nonzero density of quasiparticles in equilbrium. However since quasiparticles can also be generated dynamically (from incident radiation $hf > 2\Delta$, or left behind by a vortex) the non-equilibrium quasiparticle density is actually higher [144]. This suggests that we can't neglect contributions from these un-paired particles, and the current in a superconductor is carried partially by Cooper pair supercurrent and partially by a normal current. This means we should divide the superconductor conductivity into real and imaginary parts:

$$\sigma = \sigma_1 - i\sigma_2 \tag{2.65}$$

The full effects of the quasiparticle energy density along with the localized interactions of the electron pairs on length scale ξ are taken into account by the Mattis-Bardeen integrals for conductivity σ_1 and σ_2 [145, 167]. Relative to the normal conductivity σ_n, the conductivities are defined by the following chonky integrals:

$$\sigma_1 = \frac{1}{\omega}\int_{\Delta}^{\infty} d\epsilon \frac{(\epsilon+\omega)\epsilon + \Delta^2}{\sqrt{(\epsilon+\omega)^2 - \Delta^2}\sqrt{\Delta^2 - \epsilon^2}} \left(\tanh\frac{\epsilon+\omega}{2\tau} - \tanh\frac{\epsilon}{2\tau}\right) \tag{2.66}$$

$$\sigma_2 = \frac{1}{\omega}\int_{\Delta-\omega}^{\infty} d\epsilon \frac{(\epsilon+\omega)\epsilon + \Delta^2}{\sqrt{(\epsilon+\omega)^2 - \Delta^2}\sqrt{\Delta^2 - \epsilon^2}} \left(\tanh\frac{\epsilon+\omega}{2\tau}\right) \tag{2.67}$$

Where $\tau = \hbar/kT$ and ω is written in energy units. With the gap solved for a particular temperature (by integrating Equation 2.58), the conductivities can be solved for a given frequency and temperature (also by integrating). Interestingly, this conductivity does not

depend on any material properties [224]. The two conductivities are plotted relative to the normal conductivity σ_n in Figure 2.12. Notably, higher frequencies cause σ_2 to drop much more dramatically than σ_1, however their temperature dependence is approximately the same.

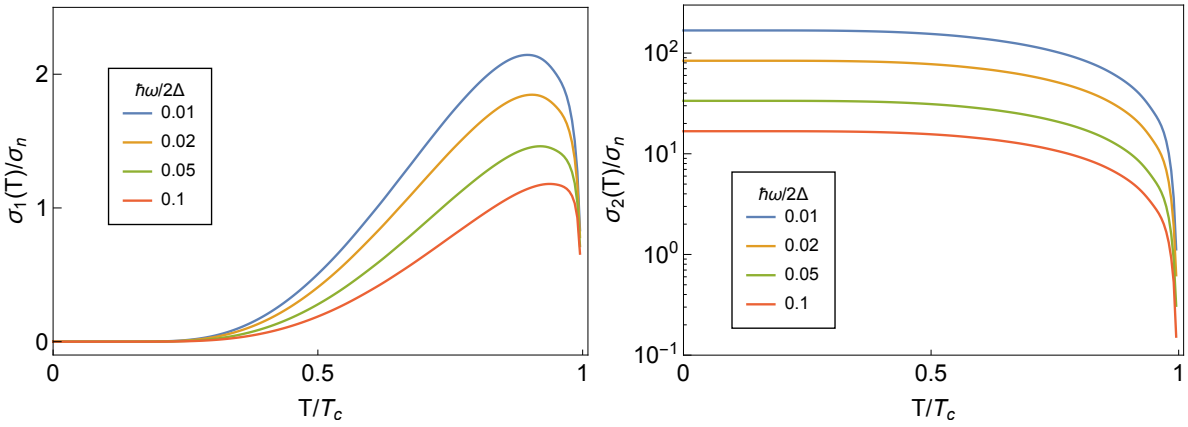

Figure 2.12: Numerically solved complex conductivities, solved for various frequencies (for niobium this range roughly corresponds to 10-100 GHz). Real conductivity σ_1 from quasiparticles is shown on the left, and complex conductivity σ_2 from Cooper pairs is shown on the right.

2.6.2 Resistance and Reactance

Now that we split the conductivity into $\sigma_s = \sigma_1 - i\sigma_2$, we should also divide the superconductor conductivity and impedance into real and imaginary parts:

$$Z_s = R_s + iX_s \tag{2.68}$$

We actually already kind of derived the reactance earlier in Equation 2.17 when we found the kinetic inductance, where we determined $X_s = \omega L_{\mathrm{T}}$. Knowing that for practical temperatures, $\sigma_2 \gg \sigma_1$ (see Figure 2.12), so we can take a similar approach to extract R_s. With the same rectangular strip as before (in which the current flows in the surface), we can write

the resistance in terms of resistivity $\rho = 1/\sigma$:

$$R_s I^2 = \mathrm{Re} \int_A \rho J^2 = \mathrm{Re} \int_A \frac{J^2}{\sigma_1 - i\sigma_2} = \frac{\sigma_1}{\sigma_1^2 + \sigma_2^2} \int_A J^2 \tag{2.69}$$

Remember, we already did a similar integral to find the sheet inductance in 2.17. This lets us write $\int_A J^2 = L_T I^2 / \mu_0 \lambda_L^2$. Thus, the resistance is:

$$R_s = \frac{\sigma_1}{\sigma_1^2 + \sigma_2^2} \frac{L_T}{\mu_0 \lambda_L^2} \approx \frac{\sigma_1}{\sigma_2} \frac{L_T}{\mu_0 \lambda_L^2} \tag{2.70}$$

Here we have assumed $\sigma_2 \gg \sigma_1$. Importantly, from this we can determine how "good" our superconductor is, by comparing the ideal reactance to the resistance (or the quality factor):

$$Q = \frac{X_s}{R_S} = \frac{\sigma_2}{\sigma_1} \omega \mu_0 \lambda_L^2 \tag{2.71}$$

Thus, by integrating $\sigma_2(T)/\sigma_1(T)$ we can estimate superconductor losses as as function of temperature.

CHAPTER 3

HIGH-FREQUENCY CIRCUIT COMPONENTS

As the frequency of an signal increases, new complexities are continuously introduced to designing electrical systems that interact with it. A conductive material of practically any shape can hold a constant voltage and conduct a fixed current, so the only primary concern for continuous currents is dissipation in the form of resistance. As a result, the basic passive electrical components (capacitors, inductors and resistors) behave exactly as expected at low frequencies regardless of their physical shape, making it quite simple to design complex circuits. However as the wavelength of the signal begins to shrink at high frequencies, the reactive components of impedance become more and more pronounced. While this means the same effect can be accomplished with smaller components (both physically and electrically) it also means that previously unimportant effects whose absence we took for granted when designing low circuits become increasingly pronounced: capacitors gain parasitic inductance, inductors gain parasitic capacitance, wire shapes begin to affect signal transmission, new forms of energy dissipation become relevant, and signals are increasingly happy to radiate and become increasingly less confined to the metal designed to carry them.

As a consequence, after frequency passes radio-waves and reaches into the microwave range ($\sim 1-10$ GHz) the reactive impedance from shaped metal on a substrate is sufficient for most applications, and we no longer have to rely on discrete capacitors and inductors! This means microwave circuits frequently look quite different from low-frequency counterparts, since filters, shunts, splitters, couplers etc. can be printed directly on the circuitboard. [1] For microwaves, planar structures have low enough parasitic effects that we can also still use geometry to define discrete capacitors or inductors and expect them to behave correctly. In practice this method of trying to form systems out of discrete microwave components

1. It is a fun exercise to visually locate the high-frequency microwave signal in a commercial circuitboard containing elements with multiple frequencies (such as one from an oscilloscope).

is frequently done by physicists designing complex circuits in order to simplify the task of modeling their theoretical behavior.

Relying to this low-frequency component-first approach proves increasingly less tractable as the circuit frequency increases into the millimeter-wave range (30-300 GHz). In this domain microfabricated planar geometries have such pronounced parasitic reactance that capacitance and inductance begin to blur together, and even simple geometry results in unexpected resonances. Empowered by the modern computational power of electromagnetic simulations, we can embrace this fluid nature of high-frequency signals and arrive at more organic-looking circuit designs. Rather than deal with parasitic reactances and attempt to refine the design until it matches the circuit we want, we can instead take a phenomenological approach and model the design behavior with very abstract components.

In this chapter, we will discuss basic components (inductors and capacitors), waveguides, resonators, and nonlinear resonators in context of high-frequency devices. While this list seems short, we can get a lot of mileage out of just these few components!

3.1 Circuit Components at High Frequencies

Passive circuits, which can store energy, move it around and behave in a linear fashion, can be fundamentally broken up into a network of inductors, capacitors and resistors. Consider the impedance of these components in an ideal case:

$$Z = R + iX = \begin{cases} R \\ 1/(i\omega C) \\ i\omega L \end{cases} \tag{3.1}$$

From the dependence on ω we see that capacitor reactance drops at higher frequencies, making them more transparent; meanwhile inductor reactance increases, making small inductive

effects more pronounced. Real devices however are inherently a combination of ideal components, since the device geometry itself contributes to the overall reactance. These (typically unwanted) effects result in small amounts of extra parasitic inductance and capacitance, and are particularly noticeable when considering the behavior deviations of real inductors or capacitors from ideal components, especially when reactance of the ideal component is small.

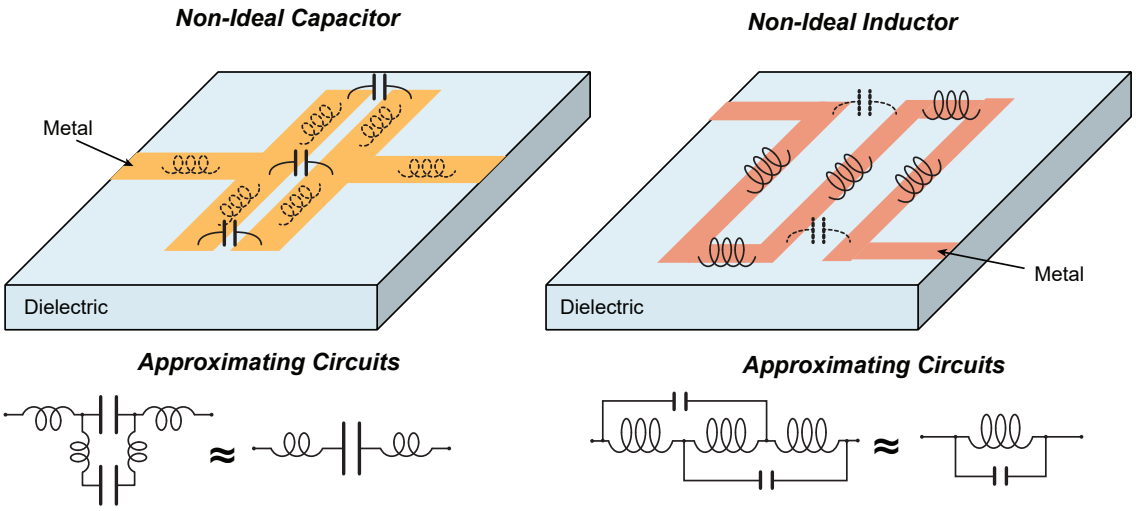

Figure 3.1: Illustration of a realistic non-ideal capacitor (left) and inductor (right). Parasitic reactances (inductive and capacitive respectively) are shown with dashed lines. Approximate circuits modelling the parasitics are shown below. Since the parasitics should be small we should be able to take some shortcuts.

Consider a simple capacitor such as the one shown on the left of Figure 3.1 consisting of two long plates on a dielectric substrate. Now considering the nonzero self inductance (or parasitic inductance) of any wire section[2], we subdivide the overall capacitance into smaller capacitors, and assign each length of wire between them a tiny inductance dL, resulting in

2. From Chapter 1, we found that kinetic inductance alone can quickly generate appreciable inductance in just one square of material, making the motivation for this kind of analysis clear.

the approximating circuit shown in Figure 3.1. The total impedance becomes:

$$Z_{\text{C,eff}} = 2i\omega dL + \left(i\omega \frac{C}{2} + \frac{1}{2i\omega dL - 2i/\omega C} \right)^{-1} \tag{3.2}$$

This is a much nastier expression, and even has zeros– meaning that at a particular frequency the capacitor will resonate! However since we are relatively sure the parasitic effects are small, we can approximate $dL \ll C$, which simplifies the expression to

$$Z_{\text{C,eff}} \simeq \frac{1}{i\omega C} + \frac{5}{2} i\omega dL \tag{3.3}$$

From which we see that the capacitor has gained an inductive correction term to its total reactance (this term is reduced but still there if we don't consider the leads).

We can follow a similar approach to consider the parasitic capacitance of an inductor, such as the one shown on the right of Figure 3.1, and subdivide the inductor with parasitic capacitors dC inserted. Making a symmetry argument[3] we can approximate the impedance by examining the circuit of a single unit cell, which gives:

$$Z_{\text{L,eff}} = \left(i\omega dC + \frac{1}{i\omega L} \right)^{-1} = \frac{i\omega L}{1 - LdC\omega^2} \tag{3.4}$$

Again we see that the effective impedance is transformed from the ideal case, and we have a self-resonance (this time a bit simpler to see). For a much more detailed analysis of the parasitic inductance in a meander inductor see Ref [158]. From this we conclude that any real circuit element will have a cutoff frequency ω_c, which can be loosely approximated by decomposing component geometry into ideal inductors and capacitors.

As we observed, the parasitic effects become more pronounced if the inherent inductance is high (eg. from kinetic inductance) or the inherent capacitance is high (eg. a high dielectric

3. Otherwise the approximate circuit becomes much more difficult to solve

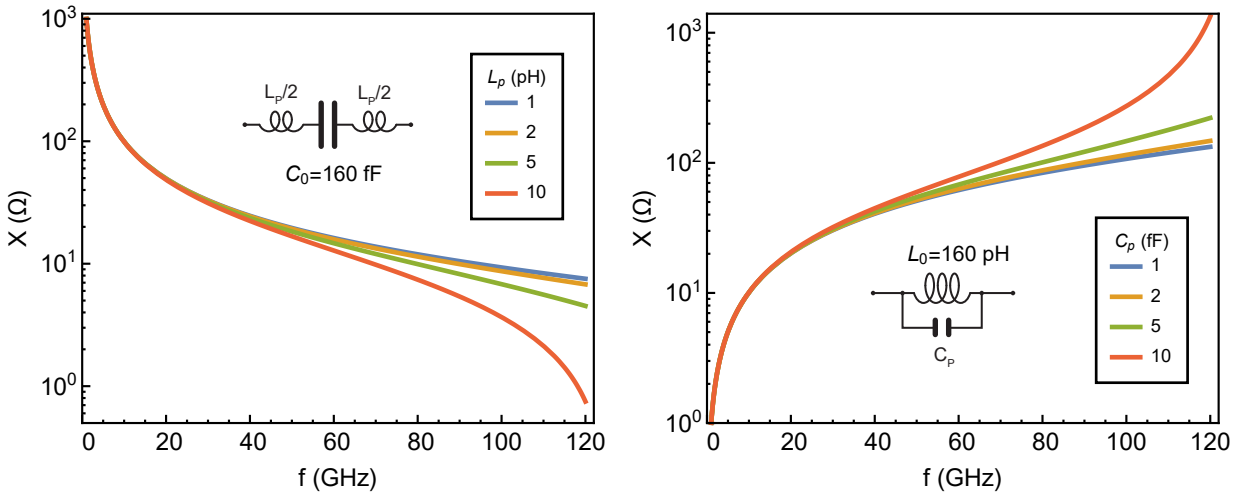

Figure 3.2: Reactance plotted as a function of frequency for a realistic capacitor (left) and inductor (right). Even up to microwave frequencies (10 GHz), the parasitics have little effect and can safely be ignored, but by 100 GHz their effects become much more significant.

constant substrate), but also should become more pronounced as frequency increases, or in other words as we approach the cutoff frequency of the device. In Figure 3.2, we plot the effects of added parasitic elements on the reactance of an capacitor and inductor with realistic values[4]. While the parasitic effects are functionally negligible at low frequency, and even don't make much of a difference at microwave frequencies, by millimeter-wave frequencies we begin to see more pronounced deviations from the ideal behavior. This means that for circuit design at 100 GHz, we need to use smaller inductors and capacitors, which are in turn now more sensitive to parasitic corrections!

3.2 Waveguides

One of the most important things we would like to do with high frequency signals is move them around. Turns out, this is exactly what electromagnetic waves like to do when left alone in a vacuum, and we could shape microwave and millimeter-wave signals with dielectric lenses

4. In practice usually one encounters higher capacitances and lower inductances (unless there is kinetic inductance), but for illustrative purposes we chose the values to be symmetric.

[244] in a very similar way to how we work with optics. Millimeter and especially microwave wavelengths are quite large however, so the size and scale of experiments would have to be quite large with this free-space approach[5] Instead, a much more practical method for moving signals around is to confine them with waveguides.

3.2.1 Hollow Waveguides

Similar to how optical fibers can transmit light through total internal reflection, one simple way to confine a signal is enclosing it in a tube with reflective walls. While in optics the reflections are generated by a difference in dielectric constants inside and outside the fiber[6], in our case, the reflective walls are typically conductors and the middle of the tube is typically air or vacuum, or another medium with minimal loss.

Let's remind ourselves of how waveguides work. Inside the dielectric (or vacuum) of the waveguide, the profile of an electromagnetic field oscillating at frequency ω will be determined by Maxwell's equations:

$$\nabla^2 \vec{H} = -\frac{\omega^2}{\mu\epsilon}\vec{H} \tag{3.5}$$

The shape of the field inside the waveguide is entirely determined by the boundary conditions on the waveguide walls, or in other words the shape of the boundaries (walls) of the waveguide. The two simplest waveguide shapes, rectangular and cylindrical are shown in Figure 3.3. Let's consider the rectangular waveguide (since we will be using these a lot). To make life simpler, we will skip ahead and only consider modes propagating in the $\hat{z}$ direction, and specifically those that have a electric field transverse to the propagation direction (TE). In other words, consider the z-component of a magnetic field with the separable form

5. Imagine how big a messy optics table from a typical laser experiment would be if you scaled it up until each lens was 100 cm in diameter!

6. Actually dielectric waveguides (essentially big optical fibers) are a thing for millimeter-waves, they just tend to be relatively lossy so aren't very popular unless you really want your waveguide to be flexible.

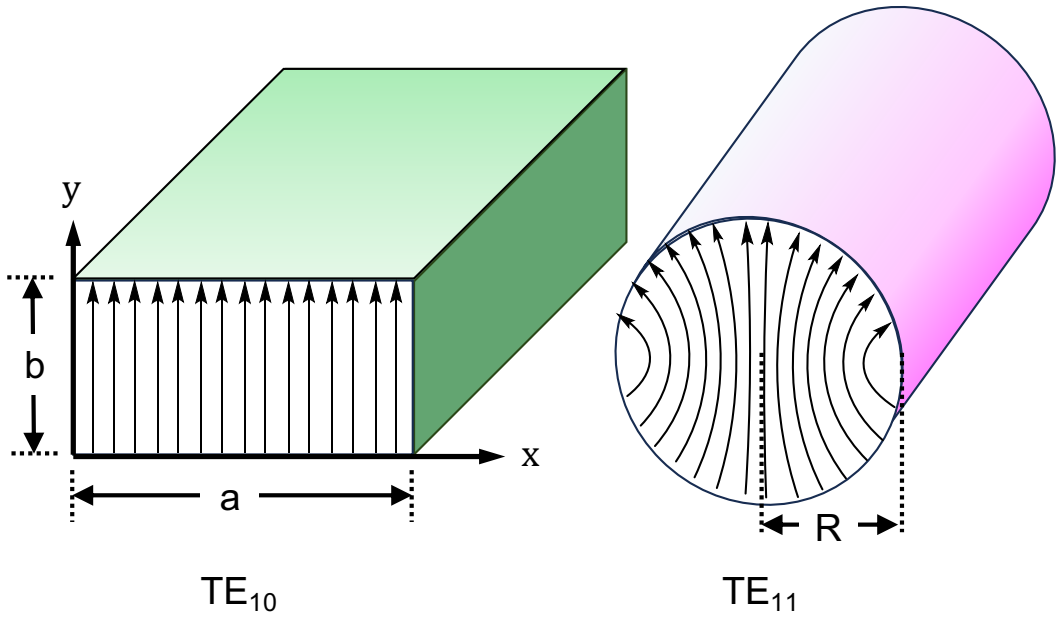

Figure 3.3: Rectangular waveguide (left) showing the electric fields of its lowest frequency TE_{10} mode, and cylindrical waveguide (right) showing its lowest frequency TE_{11} mode.

$H_z = X(x)Y(y)e^{-ik_z z}$ which satisfies Equation 3.5 in the z-direction by prapagating through the waveguide with wavevector k_z. The mode profile $X(x)Y(y)$ must satisfy:

$$\frac{\partial_x^2 X}{X} + \frac{\partial_y^2 Y}{Y} = k_z^2 - \frac{\omega^2}{\mu\epsilon} \tag{3.6}$$

This is satisfied if X and Y are sinusoidal functions, which based on the boundaries of the rectangular waveguide tells us the mode must have the form

$$H_Z = H_0 \sin k_x x \sin k_y y e^{-ik_z z} \tag{3.7}$$

With $k_x = m\pi/a$, $k_y = n\pi/b$ for integers m,n. By taking the curl of H we can also obtain the electric field of the TE modes (which nobody usually bothers writing out):

$$\vec{E}(x,y) = \hat{x}\frac{E_0 k_y}{k_x^2 + k_y^2}\cos k_x x \sin k_y y - \hat{y}\frac{E_0 k_x}{k_x^2 + k_y^2}\sin k_x x \cos k_y y \tag{3.8}$$

Satisfying Equation 3.6 yields the cutoff relations of the waveguide:

$$k_z = \sqrt{\mu\epsilon\omega^2 - k_x^2 - k_y^2} = \sqrt{\mu\epsilon\omega^2 - \left(\frac{m\pi}{a}\right)^2 - \left(\frac{n\pi}{b}\right)^2} \tag{3.9}$$

For the wave to propagate and not get exponentially attenuated, we must have a real-positive k_z, which means that the frequency must be at minimum:

$$\omega^2 > \frac{1}{\mu\epsilon}\left(\frac{\pi}{a}\right)^2 = \omega_{c,10}^2 \qquad f_c(TE_{10}, \text{vac}) = \frac{c}{2a} \tag{3.10}$$

Hence the first cutoff frequency for the lowest frequency mode of a rectangular waveguide: the $m = 1$,$n = 0$ or TE_{10} mode (Notice we defined a as the larger dimension). We've listed the cutoff frequencies for common sizes of rectangular waveguide in Table 3.1. For a while, this will be the only mode, until we hit either the TE_{20} or TE_{11} frequencies:

$$\omega^2 < \frac{1}{\mu\epsilon}\left(\frac{\pi}{2a}\right)^2, \quad \frac{1}{\mu\epsilon}\left(\frac{\pi}{b}\right)^2 \tag{3.11}$$

And now we understand why common rectangular waveguides have one dimension which is half the other! In Figure 3.4 we plot the dispersion relation (Equation 3.9) for a WR10 waveguide (a=2.54 mm), which shows the frequency range where there is a single mode in the waveguide.

For a cylindrical waveguide (shown on the right of Figure 3.3) the process is fairly similar, except for using polar coordinates. Notably, a closer look at Figure 3.3 shows how similar the field distributions are between the rectangular and cylindrical case: it is simple to imagine a rectangular waveguide with rounded corners has a field distribution somewhat in between the two. This modes for a cylindrical waveguide are Bessel function solutions [78]:

$$E_z(r, \phi) = J_n\left(\frac{u_{n,m}r}{R}\right)\cos n\phi \tag{3.12}$$

Band	Designation	Dimensions $b \times a$ (in)	TE_{10} f_c (GHz)	Nominal Range (GHz)
C	WR187	1.872×0.872	3.14	3.94 - 5.99
X	WR90	0.900×0.400	6.54	8.2 - 12.5
Ku	WR62	0.622×0.311	9.46	11.9 - 18.0
K	WR42	0.420×0.170	14.0	17.6 - 26.7
Ka	WR28	0.280×0.140	21.0	26.43 - 40.0
U	WR19	0.188×0.094	31.3	39.2 - 59.6
V	WR15	0.148×0.074	39.7	50 - 75
E	WR12	0.122×0.061	48.2	60 - 90
W	WR10	0.100×0.050	58.8	75 - 110
F	WR8	0.080×0.040	73.5	90 - 140

Table 3.1: Standard rectangular waveguide dimensions and frequency ranges, highlighting the origin of microwave band designations. Historically the K band was split up into K-under and K-above due to the strong water resonance at 22.24 GHz.

Which in turn yields a dispersion relation similar to the rectangular case but with Bessel function zeroes $u_{n,m}$:

$$k_z = \sqrt{\mu\epsilon\omega^2 - \frac{u_{n,m}^2}{R^2}} \tag{3.13}$$

In Figure 3.4 we plot the dispersion relation (Equation 3.9) for a cylindrical waveguide with radius R=2 mm, which behaves fairly similar to the rectangular waveuguide, except the higher modes are a bit more evenly spaced, meaning the single-moded frequency range is reduced. Although circular waveguides are much easier to machine, historically this has made rectangular waveguides more popular.

Now let's highlight a neat property of hollow waveguides. Consider a waveguide (eg rectangular or cylindrical) operating below its cutoff frquency. From the generalized form of Equations 3.9 and 3.13 rewritten in terms of the cutoff frequency ω_c:

$$k_z c = \sqrt{\omega^2 - \omega_c^2}$$

We see that below the cutoff frquency, the wavevector ceases to be real. This means the

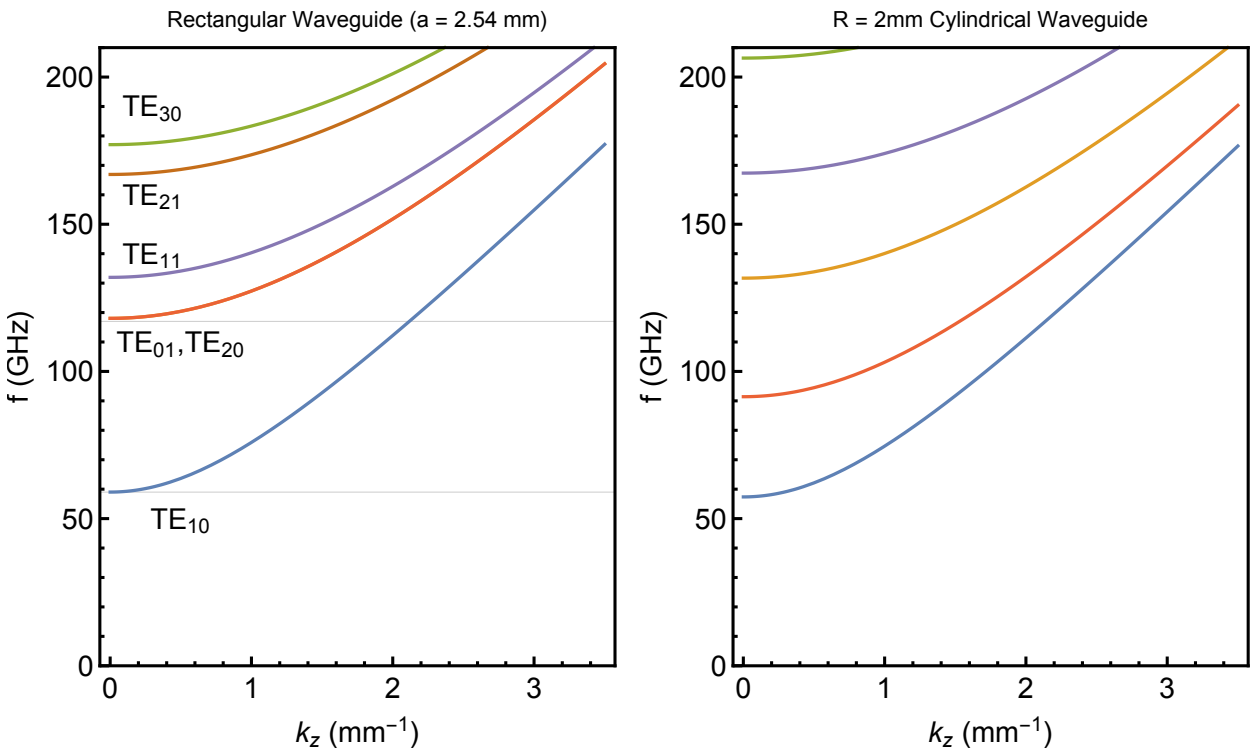

Figure 3.4: Dispersion relation of a WR10 rectangular waveguide (left) showing the cutoff frequency, and the single-moded frequency range (59–117 GHz). The dispersion relation for a cylindrical wavguide is shown on the right. The lowest frequency mode behaves fairly similar to the rectangular waveguide, however since the structure is symmetric the next highest mode cutoff is much closer.

propagating wave has a evanescent solution that decays as it propagates:

$$E(z) \sim E_0 e^{-\mathrm{Im}[k_z]z} = E_0 e^{-z\sqrt{\omega_c^2-\omega^2}/c} \tag{3.14}$$

We plot the transmission through a rectangular and cylindrical waveguide in Figure 3.5, from which we see that even a few millimeters of waveguide results in very significant attenuation! This is particularly great because it means we can rest easy and not worry about any low frequency signals propagating in our system- so any images, harmonics, noise etc. from our millimeter-wave components can safely be ignored in a waveguide since it will functionally disappear after any reasonable length of propagation.

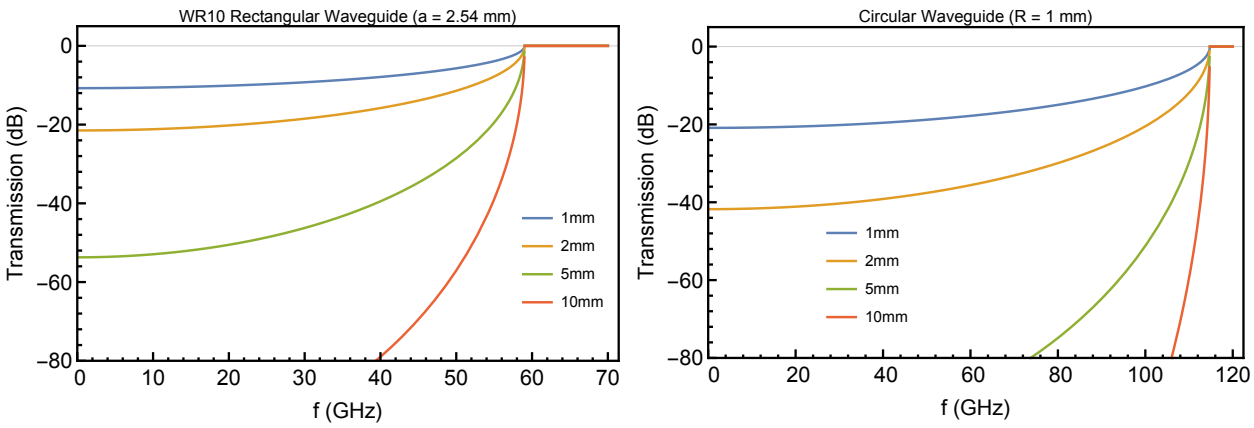

Figure 3.5: Transmission through a rectangular waveguide (left) and circular waveguide (right) below their cutoff frequencies, plotted for different lengths of waveguide section, showing the drastic attenuation resulting from the evanescent propagation.

As another benefit of this property, consider propagation through the loss-less[7] circular waveguide (very easy to machine) with radius 1 mm shown on in Figure 3.5. Suppose we create a cavity at say 70 GHz and surround it with some of these lossless evanescent waveguides: by just drilling a few holes we have very easily suppressed leakage out of the cavity! This is precisely the approach we used to realize millimeter-wave resonant cavities [217] with convenient optical access for transduction experiments [119].

Rectangular Waveguide Properties

Since we make good use of the rectangular waveguide later on, let's quickly write down some of their properties. The group velocity (the speed at which the energy propagates) can be derived from the k vector:

$$v_{\mathrm{g}} = \frac{\partial \omega}{\partial k} = c\sqrt{1 - \frac{\omega_c^2}{\omega^2}} = c\sqrt{1 - \left(\frac{c}{2af}\right)^2} \tag{3.15}$$

7. Metal waveguides already have very low loss, but superconducting waveguides even more so!

This is what we should use for calculating the group delay. Notably unlike a coaxial cable, the group velocity is not linear with frequency, so the phase a signal builds up after travelling for some distance will not be a linear function of frequency! This makes calibrating to a waveguide of known length a bit of a headache.

What about the impedance (the other important transmission line property)? Since this is a ratio of voltage to current, turns out this is not so simple to calculate for a rectangular waveguide, since we now need to integrate the fields. Waveguides have a separately defined quantity called the wave impedance [184] which in free space is

$$Z_0^w = \sqrt{\frac{\mu_0}{\epsilon_0}} \tag{3.16}$$

And for the TE modes (the one we use) of our waveguide gets modified to

$$Z_{TE}^w = \frac{Z_0^w}{\sqrt{1 - \omega_c^2/\omega^2}} \tag{3.17}$$

Above the cutoff frequency this impedance is real (in other words energy propagates). As a rule of thumb, throughout the operating band of a rectangular waveguide the wave impedance is around $\sim 500\ \Omega$. Note that this is not quite the same quantity as electrical impedance though (hence we label it Z^w).

3.3 Coaxial Lines at High Frequency

Unless you are working near or above 100 GHz hollow waveguides are quite large. Additionally, since they need walls made of conductors, waveguides are typically made up of rigid and bulky metal pieces that need to be welded or tightly secured together. Luckily high-frequency signals can also be carried by coaxial cables, which re the standard method for making microwave connections due to their (literal) flexibility.

The addition of a center conductor to the hollow cylindrical waveguide greatly simplifies

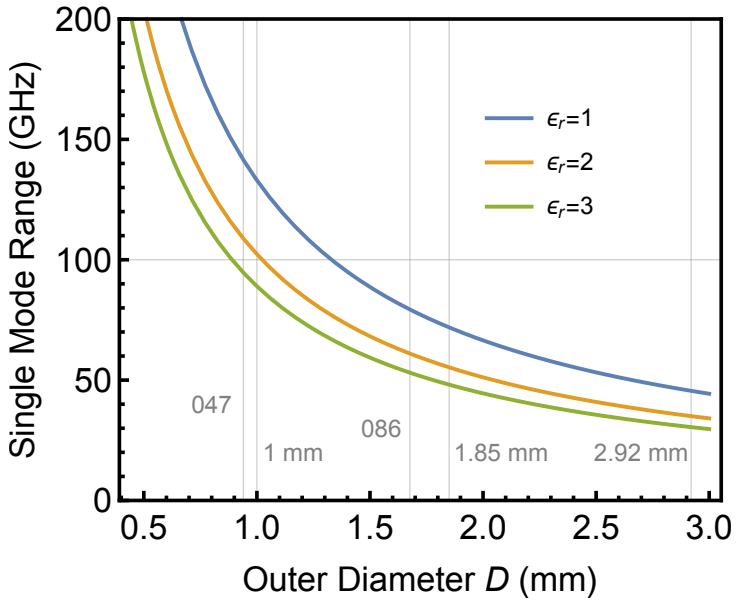

Figure 3.6: Single-moded operation range of a coaxial cable with a 50 Ω impedance, plotted as a function of outer diameter D for different dielectric constants (For example Teflon has $\epsilon_r = 2.02$). Standard coax connector and cable dimensions are marked in gray.

the boundary conditions, and instead of TE or TM allows the geometry to support a TEM wave, which has a cutoff frequency of zero! Meaning the TEM mode propagates even at DC. For a coaxial cable (inner and outer diameters d and D) filled with a dielectric[8] with permittivity ϵ_r, the the cutoff frequency of the next mode (TE_{11}) is approximately [184]:

$$f_c \approx \frac{2}{\pi(D+d)\sqrt{\mu\epsilon}} = \frac{2c}{\pi(D+d)\sqrt{\mu_r\epsilon_r}} \tag{3.18}$$

This determines the operating range of a coaxial cable or connector. The ratio of the diameters are usually constrained to give a fixed impedance (commonly 50 Ω) which is determined by [184]:

$$Z_{\text{coax}} = \frac{1}{2\pi}\sqrt{\frac{\mu}{\epsilon}}\ln\left(\frac{D}{d}\right) \tag{3.19}$$

For teflon, which is a standard dielectric with a low $\epsilon_r = 2.02$, we obtain approximately 50 Ω

8. because otherwise how would the center conductor stay in the middle and not fall down?

when $D \approx 3d$. In Figure 3.6 we plot the the cutoff frequency of the higher mode (or the single-moded operating range) for cables adjusted to 50 Ω for different dielectric constants. From this we see that for single-mode operation above 100 GHz or so, we need the outer diameter D to be below 1.2 mm for vacuum or 0.9 mm for teflon. These dimensions are small enough that the connector parts are pretty tiny (also fragile and expensive).

Figure 3.7: Photograph of a partially assembled 1 mm coaxial connector. The outer diameter of the dielectric (white ring) is less than 1 mm, so these parts are very fragile!

Another issue we haven't considered yet is loss. Unfortunately for us, the attenuation in a coaxial cable increases with frequency. The attenuation from conductive losses [184] scales with the root of frequency:

$$\mathrm{Im}[k_z] = \frac{1}{2\pi Z_0}\sqrt{\frac{\omega\mu}{2\sigma}}\left(\frac{1}{d}+\frac{1}{D}\right) \tag{3.20}$$

This is not great news for us, especially since the frequency cutoff makes us select the tiny 1 mm cables, which also increase attenuation. This means that in the absolute best case[9] cables attenuate at least 3 dB/ft [7]. Now waveguides aren't looking so bad!

9. Silver-plated copper cables with foamed PTFE dielectric seems to be the best we've seen on the market. Foamed PTFE means less garbage dielectric and more nice clean vacuum in the cable.

3.4 Planar Transmission Lines

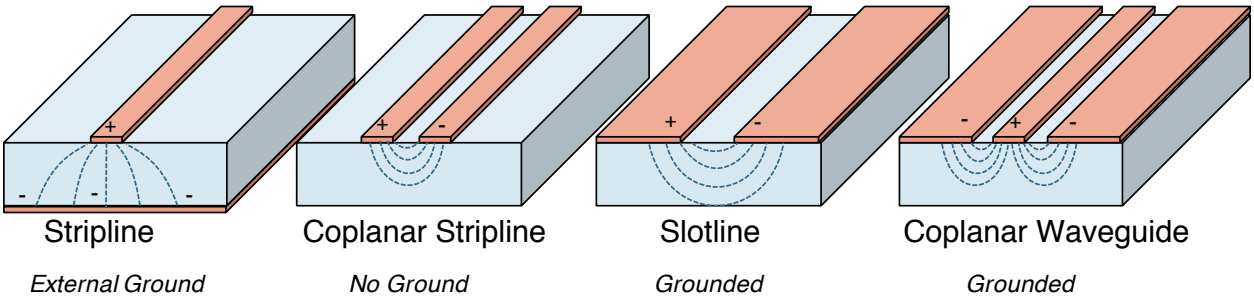

Figure 3.8: Illustration of common types of planar transmission lines, with electric fields labelled with dashed lines. Other than the stripline, all the surface transmission lines shown here can be operated with a backside or external ground.

For microstructures on chips or circuit boards, cables or waveguides aren't practical for moving signals around. Instead we can use two dimensional metal structures patterned on a flat dielectric surface: this could be anything from crystalline materials (sapphire or silicon), teflon, glass fiber dielectrics or even cheese [133]! There are several metal geometries that will do the trick: we show a few popular types in Figure 3.8. Like the coaxial cable, all of these designs work by having two primary conductors carrying differential voltage and current, meaning they can support a TEM mode (illustrated in Figure 3.9). This can either be a positive conductor and a large ground plane (as in a stripline or the coplanar waveguide) or two identical conductors (or ground planes) as in the case of the coplanar stripline or slotline.

For any of these geometries, the two conductors are patterned in a single step from a thin metal film deposited[10] on the surface of the dielectric. Most commonly we use either high-conductivity metals or superconductors along with low-loss dielectrics, so these structures can be modelled as a lossless transmission line [184] shown in Figure 3.9. Neglecting dissipation (hence no resistors) the two conductors are modelled as inductor per unit length L_0 with

10. The metal can be taped to the surface in case of the cheese dielectric [133]

capacitive coupling per unit length C_0 between them.

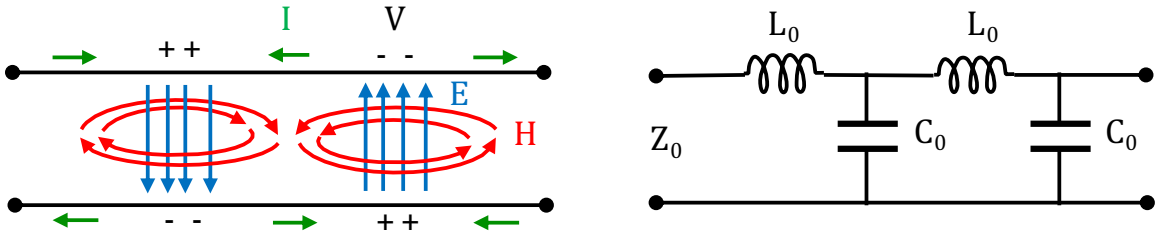

Figure 3.9: Illustration of a transmission line showing the electric and magnetic fields of its supported TEM mode (left) along with charges and currents, and the lossless model (right) of the transmission line made up of ideal components.

In this structure a signal will have a propagation constant $\beta = \omega\sqrt{L_0 C_0}$, so a wave will propagate with velocity

$$v_p = \frac{1}{\sqrt{L_0 C_0}} \tag{3.21}$$

It can be useful to think about the transmission line as a dielectric medium with an effective dielectric constant ϵ_{eff} in which the wave propagates at $v_p = c/\sqrt{\epsilon_{\text{eff}}}$. The characteristic impedance of this transmission line is simply a ratio of the inductance and capacitance (so it is independent of how long it is) given by

$$Z_0 = \sqrt{\frac{L_0}{C_0}} \tag{3.22}$$

This impedance determines the transmission characteristics, including how signals reflect off of discontinuities.

In our case these transmission lines are frequently fabricated with superconductors, which have kinetic inductance (described in Chapter 1). Previously we examined how this modifies the inductance of a stripline. Keeping things general, we will conveniently write out the inductance and capacitance for planar waveguide geometries including kinetic inductance and realistic effects such as finite substrate thickness. Defining the conductor width as w

and the gap between conductors as s, the inductance and capacitance for a coplanar stripline (see Figure 3.8) patterned on a substrate with finite thickness h will be [259]

$$C_0 = \epsilon_0 \frac{K(2k_0')}{K(k_0)} + \frac{\epsilon_r - 1}{2}\frac{K(k_2')}{K(k_2)} \qquad \text{where} \quad k_2 = \frac{\tanh\left(\frac{\pi s}{4h}\right)}{\tanh\left(\frac{\pi(s+2w)}{4h}\right)} \tag{3.23}$$

$$L_0 = \mu_0 \frac{K(k_0)}{K(k_0')} + \mu_0 \frac{2\lambda_L^2}{ts(s+2w)k_0' K(k_0')}\left(s\log\left(\frac{4w(s+2w)}{d(s+w)}\right) + (s+2w)\log\left(\frac{4sw}{t(s+w)}\right)\right) \tag{3.24}$$

The first terms for both expressions are the normal versions K is the elliptic integral of the first kind, $k_0 = s/(s+2w)$ and $k' = \sqrt{1-k^2}$. λ_L is the London length (see Chapter 1) and t is the superconducting metal film thickness. Similarly the inductance and capacitance for a coplanar waveguide (CPW) with center pin width w, gap s and substrate height h will be [246]:

$$C_0 = \epsilon_0 \frac{K(k_0)}{K(k_0')} + \frac{\epsilon_r - 1}{2}\frac{K(k_2)}{K(k_2')} \qquad \text{where} \quad k_2 = \frac{\sinh\left(\frac{\pi w}{4h}\right)}{\sinh\left(\frac{\pi(w+2s)}{4h}\right)} \tag{3.25}$$

$$L_0 = \mu 0 \frac{K(k_0)}{K(k_0')} + \mu_0 \frac{\mu_0\lambda_L^2}{2wtk_0^2 K(k_0)^2}\left(-\log\left(\frac{t}{4w}\right) - k_0\log\left(\frac{t}{4(w+2s)}\right) + \frac{2(w+s)}{w+2s}\log\left(\frac{s}{w+s}\right)\right) \tag{3.26}$$

Note that for a CPW, k_0 is now $w/(w+2s)$! These expressions have been derived in detail using conformal mapping techniques[121][11] in Refs [246, 259]. Here I've just simplified the math into unified long but useable functions useful for modelling these two common types of waveguide in one go. Practically speaking though, especially if kinetic inductance is included I've found that this analytic approach only goes so far in approximating waveguides- it is more accurate to simulate the structure, and even more accurate to measure it directly!

11. Using these conformal mapping techniques you can actually find the capacitance or inductance of virtually any waveguide geometry, usually as some combination of elliptic integrals.

3.5 Resonators

For the waveguides and transmission lines we discussed earlier, the signal propagates as $e^{ik_z z}$ or $e^{i\beta l}$ (depending on the notation used). Now consider what happens if the waveguide arrives at a discontinuity: if we arrive at a short (a wire connecting the two ends of a transmission line, or a conducting wall closing off a hollow waveguide) then the electric field (or voltage) at the discontinuity must be zero, so we will have a wave reflecting with negative polarity. Likewise if the waveguide simply ends (an open) no current can flow at the discontinuity so a wave will reflect with the same polarity[12]. Now consider a length d of waveguide or transmission line terminated on both ends by a discontinuity. You can either think of this situation as a signal bouncing back and forth between them, or by modifying the electric field solutions with the newly applied boundary conditions[13]: either way we obtain the same result. At particular frequencies satisfying $k(\omega)d = n\pi$ we will have a standing wave solution for the electric field, which resonates inside our volume. Just like that we've formed a resonator! Let's go over them in excruciating detail.

3.5.1 Rectangular Cavity (Box) Resonator

Recall the solution from earlier (Equation 3.9) for the rectangular waveguide mode propagating on the z axis. Except now we add similar boundary conditions on the z-axis:

$$\mu\epsilon\omega^2 - \left(\frac{m\pi}{a}\right)^2 - \left(\frac{n\pi}{b}\right)^2 = k_z^2 = \left(\frac{p\pi}{d}\right)^2 \tag{3.27}$$

From this we obtain discrete frequency solutions:

$$\omega_{nmp} = c\sqrt{\left(\frac{m\pi}{a}\right)^2 + \left(\frac{n\pi}{b}\right)^2 + \left(\frac{p\pi}{d}\right)^2} \tag{3.28}$$

12. For both cases the signal will actually radiate out a little bit: think of the boundary condition as a radiating dipole and you can extract the power it emits into free space.

13. Think of this as a personality test: are you an optics person or a circuits person?

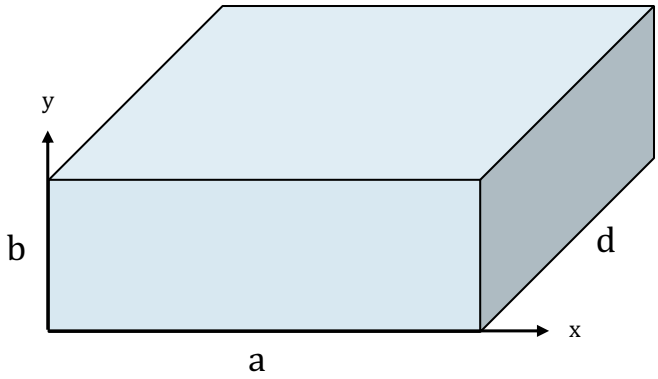

Figure 3.10: Diagram of a rectangular resonance (or box mode). Did I mention it looks like a box?

Other than the trivial solution ($\omega_{000} = 0$) the boundary conditions dictate at least two of m,n,p are nonzero, so the first few modes will be 110, 011, 101. This means the second-smallest dimension is the one that will determine the lowest mode frequency.

To get a feeling for these modes, consider a chip or substrate: a rectangular slab of dielectric, with relatively thin thickness $d < 1$ mm and horizontal dimensions a, b that are at least $\gtrsim 1$ mm (big enough to carefully handle with tweezers). From the discrete frequency solutions above, we can easily calculate the mode frequencies formed by this structure. We plot the mode frequencies for a 0.1 mm-thick[14] square ($a = b$) substrate for a few different relative dielectric constants in Figure 3.11.

Two things are immediately apparent. First we observe that if working at microwave frequencies (below 10 GHz), we don't really need to worry about these substrate modes at all for chips smaller than 5 mm! For millimeter-waves this is a different story: even for $\epsilon_r = 1$, a chip larger than 2 mm begins to have substrate modes below 100 GHz. Quartz (commonly used for millimeter-waves because of its low dielectric constant) pushes these modes down even further. If we select a high-dielectric constant substrate such a silicon or sapphire (standard low-loss substrates used for superconducting quantum devices) these

14. This is quite thin as far as wafers go!

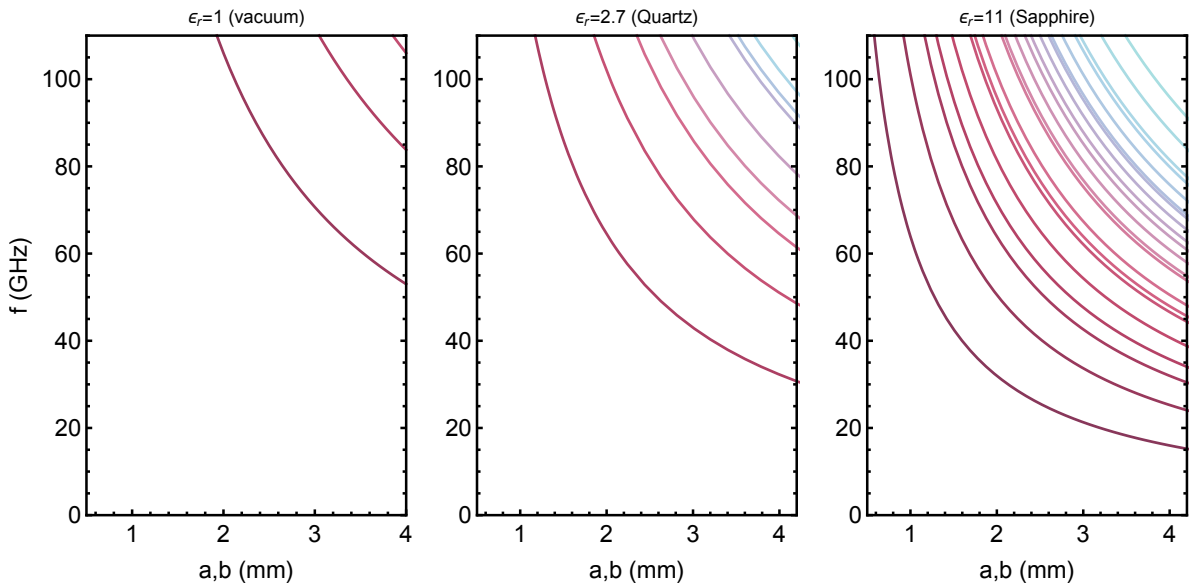

Figure 3.11: Mode frequencies of a 100 μm-thick box made from vacuum, quartz and sapphire respectively, assuming the other two dimensions are equal.

modes are quite dense in the millimeter-wave bands, and we have a box mode below 100 GHz even for a tiny 1 mm chip!

Remember though, we derived these frequencies assuming conducting boundaries on all sides of the box. So what this really means is we should be very careful when enclosing a substrate on both sides with metal, as this will almost certainly form parallel plate modes in the substrate[15]. So unless we have extremely tiny chips, we immediately rule out the popular style of packaging [77] (where a chip with a groundplane is pressed down on a metal surface) at millimeter-waves.

This is not the end of the world for millimeter-wave packaging though. Suppose we still wish to use sapphire: let's simply make the box larger (say 1 mm thick) and suspend the sapphire chip in the center. We plot the resulting modes for varying chip thicknesses in Figure 3.12: from which we see that if we keep the chip thickness small (100 μm thick) we can make a reasonably size chip (2–3 mm) and only have a few discrete modes in our

15. The electric field of these modes looks a bit like pancakes in the substrate.

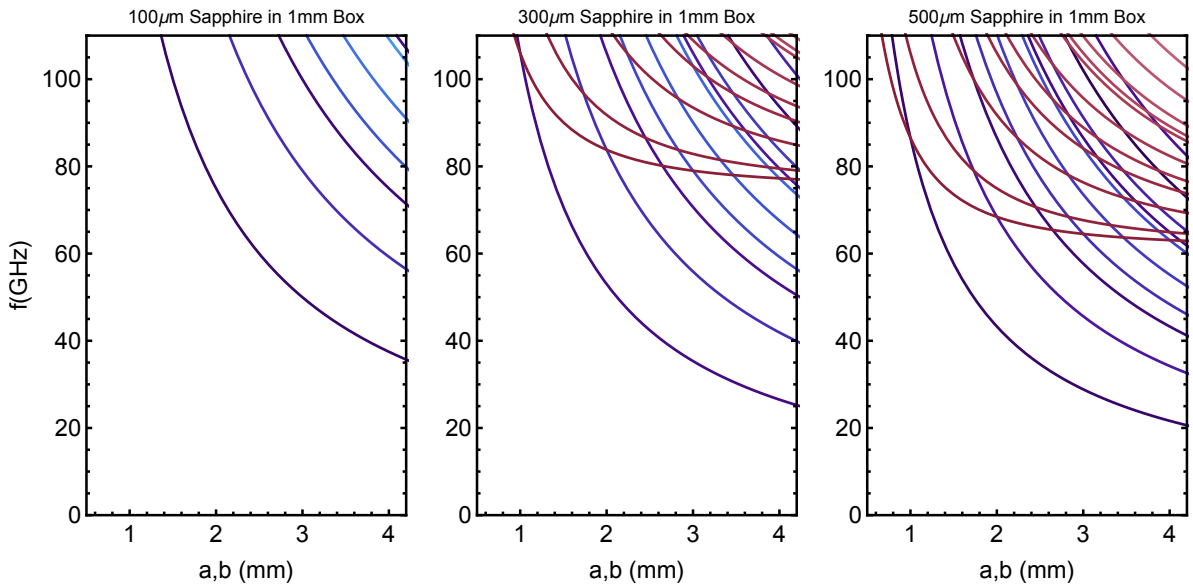

Figure 3.12: Mode frequencies of a 1 mm-thick box with varying square dimensions containing a sapphire chip of varying thickness.

frequency range! Thus we have kept the millimeter-wave box mode beast at bay.

3.5.2 Transmission Line Resonator

In a very similar way to the case of the box resonance we looked at above, we can also determine the frequencies of a length of transmission line terminated with discontinuities. If the transmission line is shorted on both ends, the voltage on either side must be zero, so $\beta l = n\pi$ for integer n. In other words the length of transmission line forms half a wavelength, so its resonant frequency will be:

$$\omega_{\lambda/2} = \frac{n\pi}{l\sqrt{L_0 C_0}} \tag{3.29}$$

Two open circuit conditions on either end result in exactly the same frequency. However a short circuit on one end and an open circuit on the other allows transmission line to support

a mode with twice the wavelength:

$$\omega_{\lambda/4} = \frac{(2n+1)\pi}{2l\sqrt{L_0 C_0}} \tag{3.30}$$

This kind of resonance is particularly useful since it requires half the space. Since one end is a voltage node and the other is a current node, it will interact capacitively on one end and inductively on the other.

3.5.3 Simple LCR Resonator

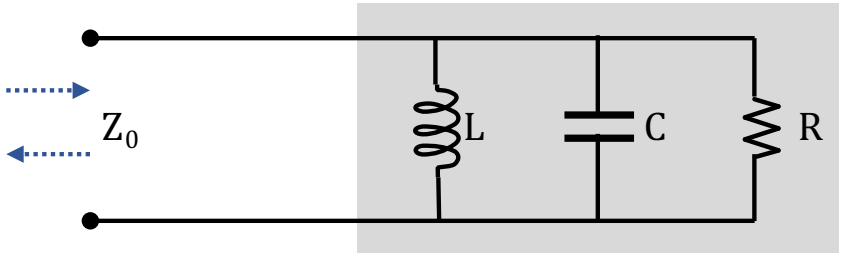

Figure 3.13: Circuit diagram of the simplest LCR resonator

A resonance can also be modelled by a set of ideal circuit components. For an arbitrary load impedance Z_L shunting an input line with impedance Z_0 we can write the reflection as an S matrix parameter [184]:

$$S_{11} = \frac{Z_L - Z_0}{Z_L + Z_0} \tag{3.31}$$

In our case, the load impedance is that of a parallel LRC resonator (Figure 3.13). We can write the effective impedance as:

$$Z_L = Z_R || Z_L || Z_C = \frac{1}{i\omega C + \frac{1}{i\omega L} + \frac{1}{R}} \simeq \frac{R}{1 + 2iQ_i\Delta} \tag{3.32}$$

Where $\Delta = (\omega - \omega_0)/\omega_0$, the resonance frequency is $\omega_0^2 = 1/LC$ and the quality factor is $Q_i = R/\omega_0 L$. From this we can evaluate the reflection, and after some algebra arrive at the

following expression:

$$S_{11}(\omega) = 1 - \frac{2Q_i}{Q_e + Q_i + 2iQ_eQ_i\Delta} \tag{3.33}$$

$$= 1 - \frac{\kappa}{\kappa+\gamma}\frac{1}{\frac{1}{2} + i\delta} \tag{3.34}$$

Where we have introduced reduced detuning $\delta = \Delta\omega/(\kappa+\gamma)$ (a concept we are borrowing from the quantum optics derivation of the resonator [239]) and quality factors and coupling rates definitions can be summarized:

$$\text{coupling Q} = \text{external Q}: \qquad Q_e = \frac{Z_0}{\omega_0 L} = \omega_0 Z_0 C = \frac{\omega_0}{\kappa} \tag{3.35}$$

$$\text{unloaded Q} = \text{internal Q}: \qquad Q_i = \frac{R}{\omega_0 L} = \omega_0 RC = \frac{\omega_0}{\gamma} \tag{3.36}$$

$$\text{loaded Q} = \text{total Q}: \qquad \frac{1}{Q_T} = \frac{1}{Q_e} + \frac{1}{Q_i} = \frac{\kappa+\gamma}{\omega_0} \tag{3.37}$$

Having abstracted its behavior using quality factors Q or linewidths κ, this simple resonator

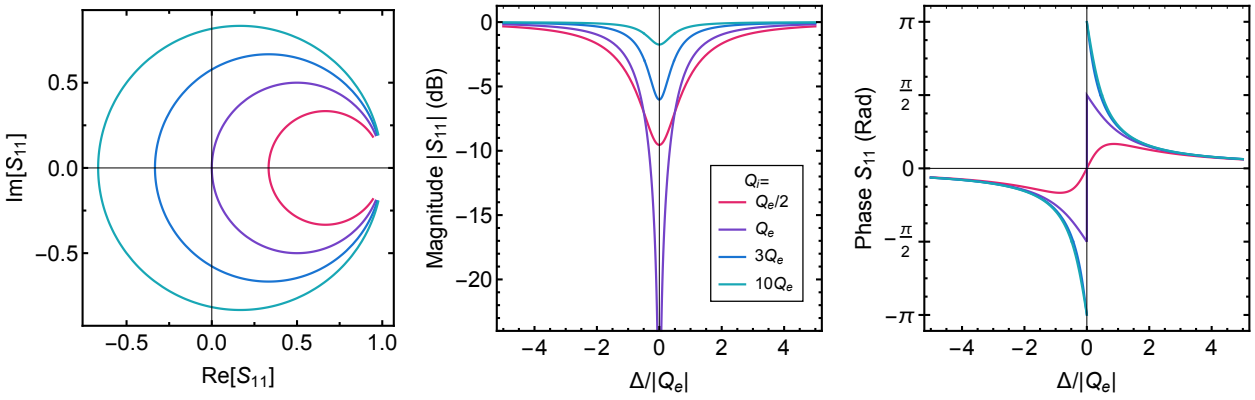

Figure 3.14: Complex reflection response of a resonator, plotted for various strengths of internal losses Q_i. The ambiguity of the magnitude response between under-coupled and over-coupled cases necessitates a full measurement of both quadratures.

model is quite powerful, since it can model the behavior of a variety of different resonators, as we will see later on. We plot the complex behavior of this function for varying values of

coupling Q_e and internal Q_i in Figure 3.14. Noting the similarity in magnitude responses between cases with $Q_e < Q_i$ and $Q_e > Q_i$ makes it necessary to measure the complex response (both real and imaginary) to get fully understand what is going on.

Frequently we also want to estimate the number of photons in the resonator given a power applied to the input port. It is important to realize that this is not just a function of applied power, but also depends on the frequency detuning [239]. We will derive this more carefully later for a nonlinear resonator, which in the current linear case simplifies to:

$$n_{\text{ph}} = \frac{P_{\text{in}}}{hf_0} \frac{\kappa}{(\kappa+\gamma)^2} \frac{1}{\frac{1}{4}+2\delta^2} \leq \frac{P_{\text{in}}}{\hbar\omega_0} \frac{4\kappa}{(\kappa+\gamma)^2} \tag{3.38}$$

3.5.4 Real Resonator Measured In Reflection

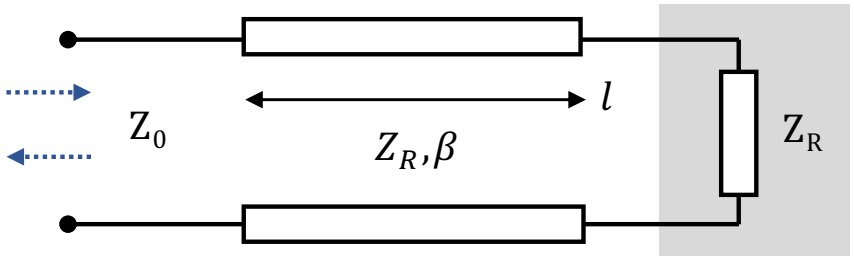

Figure 3.15: LCR resonator with multiple loss channels R_i.

In a real experimental setting resonator coupling isn't perfect: frequently the presence of the resonator modifies the impedance of the transmission line coupling to it, or even more frequently the coupling method has some electrical length (eg. consider coupling to a resonator through an evanescent waveguide we discussed earlier). Following the approach in [250], we model this realistic case with the modified circuit shown in Figure 3.15 where the resonator circuit from before is condensed into Z_R for simplicity. The coupling to the circuit is represented by a short section of transmission line, which scatters the signal before reaching the resonator, and introduces a phase $\beta l \approx \phi'$ to the signal that reaches the resonator. The modified reflection coefficient S_11 can be derived using ABCD matrices [184] (see Appendix

A for details), and approximated for small ϕ' as:

$$S_{11} \approx \frac{Z_R - Z_0 - i\phi' \left(\frac{Z_0^2}{Z_c} - Z_c \right)}{Z_R + Z_0 - i\phi' \left(\frac{Z_0^2}{Z_c} - Z_c \right)} \tag{3.39}$$

Absorbing the complex term into a complex impedance $Z_0^* = |Z_0^*|e^{i\phi}$, and therefore based on the definitions above we will have a modified complex coupling quality factor

$$Q_e^* = |Q_e^*|e^{i\phi} \tag{3.40}$$

Thus the reflection is transformed:

$$S_{11}(\omega, \phi) = 1 - \frac{2Q_i e^{-i\phi}}{Q_e + Q_i e^{-i\phi} + 2iQ_e Q_i \Delta} \tag{3.41}$$

This makes several important changes: in the complex plane, the circle that is swept by the reflection off the resonator in the complex plane is simultaneously rotated and the circle diameter increases. This means that particularly if the resonator response is being used to extract the internal and external quality factors, it is vitally important to take this rotation into account: otherwise the internal quality factor will be overestimated, since the diameter of the circle has been artificially increased by the mismatch rotation ϕ.

3.5.5 Side-coupled Resonator

Measuring a resonator in reflection makes for a simple experiment, however in practice this is a very inefficient use of resources as each resonator now requires its own transmission line[16]. A more efficient characterization method is to have one continuous transmission line with

16. Worse yet measuring in reflection requires a way to separate inputs and outputs, or a circulator for every device.

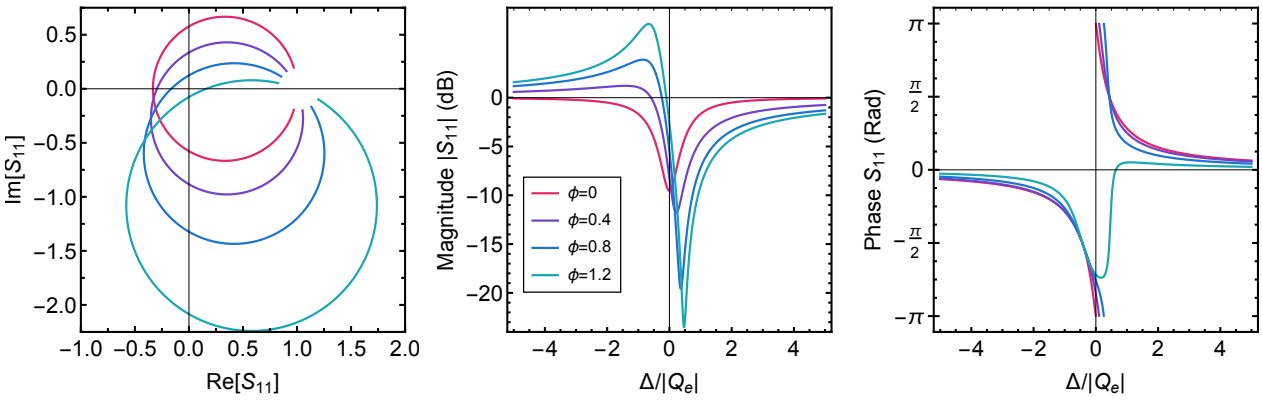

Figure 3.16: Complex asymmetric reflection response of a resonator, plotted for various strengths of asymmetry ϕ and internal losses Q_i.

many resonators coupled to it from the side (hence side-coupled)[17]. This can be realized by capacitively coupling the LCR circuit in Figure 3.13 to a continues transmission line, as shown in Figure 3.17.

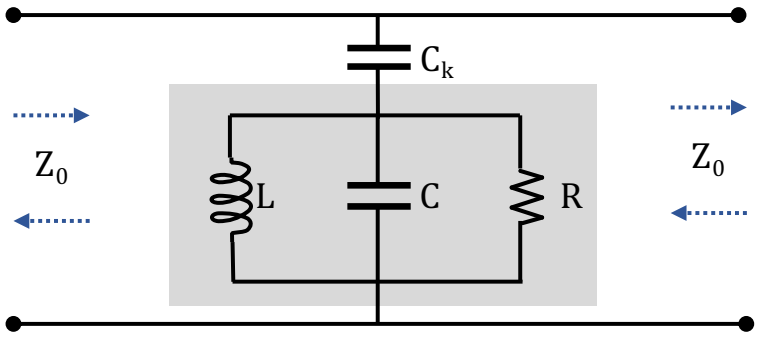

Figure 3.17: Circuit diagram for a side-coupled resonator.

We can use ABCD matrices to derive transmission and reflection parameters of this circuit by modifying the shunt impedance to $Z_L = Z_R + 1/i\omega C_k$ (see Appendix A). Notably this approach already produces an asymmetric response, with a resonanace and anti-resonance! So we would be kidding ourselves if we didn't take it into account. You can derive this asymmetry from the circuit model directly: this is done explicitly in Refs. [106, 153]. We

17. Many people like referring to these as 'hanger' measurements which I think is absolutely horrid.

will instead purposefully derive the asymmetry from a directional coupling argument to an LCR resonator so that we can conveniently take advantage of well-established physics of resonators in reflection.

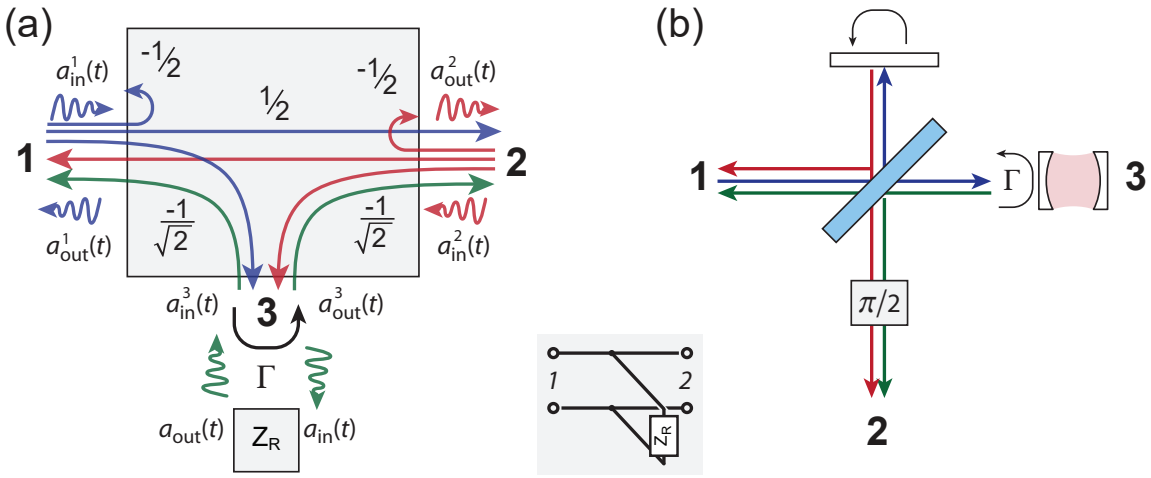

Figure 3.18: Side-coupled resonator model (a) 3 port network of an H-plane splitter coupled to a black box resonator, showing corresponding transmission coefficients with input and output fields labelled by port. The inset shows the equivalent circuit network. (b) Analogous configuration for an optical cavity, adjusted for boundary conditions.

Based on the argument that a symmetrically coupled resonator will radiate equally in both directions, we consider the the 3-port H-plane splitter. This lossless but unmatched network [184] has symmetric ports 1-2 corresponding to the transmission line, and unmatched port 3 leading to the single-port coupled resonator Z_R as shown in Fig. 3.18(a) (For optics nerds, this system can also be described by a network consisting of a 50-50 beamsplitter, perfect mirror, and $\pi/2$ phase shifter as shown in Fig. 3.18(b), which yields the same key results if we are careful to use correct boundary conditions). If we place a black-box element on port 3, we can describe it's input and output fields in terms of the waveguide input and output fields:

$$\begin{pmatrix} a_{\text{in}}(t) \\ a_{\text{out}}(t) \end{pmatrix} = \frac{-1}{\sqrt{2}} \begin{pmatrix} a_1^{\text{in}}(t) + a_2^{\text{in}}(t) \\ a_1^{\text{out}}(t) + a_2^{\text{out}}(t) \end{pmatrix} \tag{3.42}$$

If we describe the black box with an arbitrary reflection term $\Gamma = a_{\text{out}}(t)/a_{\text{in}}(t)$, the scattering matrix of the system reduces to:

$$S = \frac{1}{2}\begin{pmatrix} \Gamma - 1 & \Gamma + 1 \\ \Gamma + 1 & \Gamma - 1 \end{pmatrix} \tag{3.43}$$

We can now verify that far off-resonance, for open circuit perfect reflection $\Gamma \to 1$, we recover perfect transmission. Meanwhile substituting Equation 3.34 for Γ, we recover the limit $C_k \to 0$ from Appendix A:

$$S_{21} = 1 - \frac{Q_i}{Q_e + Q_i + 2iQ_eQ_i\Delta} \tag{3.44}$$

$$= 1 - \frac{\kappa}{\kappa + \gamma}\frac{1}{1 + 2i\delta} \tag{3.45}$$

This response function is plotted for various values of internal quality factors Q_i in Figure 3.19. As a rule of thumb, we can look for circles in the complex response to track down resonances, and having found one, see that the larger Q_i is, the deeper the resonance will be. This makes side-coupled resonators a great tool for quickly extracting information about the internal loss of a resonator!

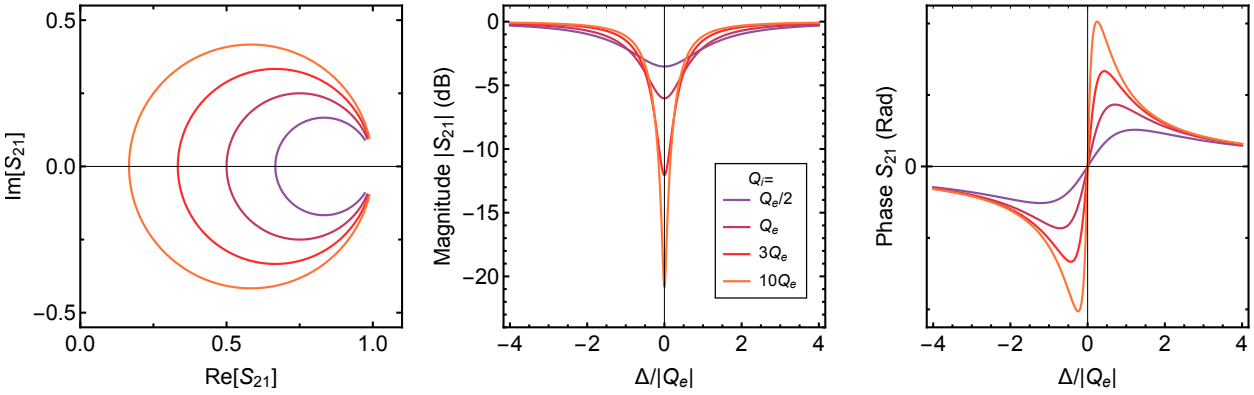

Figure 3.19: Complex transmission response of a side-coupled resonator, plotted for various strengths of internal losses Q_i relative to the coupling Q_e.

Frequently we also want to estimate the number of photons in side-coupled resonator given a power travelling through the transmission line. As with the reflectively-coupled resonator, this is not just a function of applied power, but also depends on the frequency detuning [239]. We will derive this more carefully later for a nonlinear resonator, but for now it simplifies to:

$$n_{\text{ph}} = \frac{P_{\text{in}}}{hf_0}\frac{\kappa}{(\kappa+\gamma)^2}\frac{1}{\frac{1}{2}+2\delta^2} \leq \frac{P_{\text{in}}}{\hbar\omega_0}\frac{2\kappa}{(\kappa+\gamma)^2} \tag{3.46}$$

Note the maximum value is exactly half that of the reflection case![18]

3.5.6 Asymmetry in Side-coupled Resonators

As with real experimental implementations of resonators measured in reflection, real side-coupled resonators also show signs of asymmetry. While this might sound like a minor issue and just a meaningless correction to think about, consider the following: if we use the simple model from before to fit our asymmetric resonator (see the profiles in Figure 3.20), we might end up wildly overestimating the quality factor [186]! For this reason it is vitally important in this day and age to measure both quadratures of transmission (or reflection) to claim anything about the resonator with any amount of certainty.[19]

Now naively one might say - didn't we already figure out how asymmetry works for reflection? Can't we just substitute the modified reflection case into the reflection term on the beamsplitter? But because that derivation made assumptions about the input impedance Z_0, **this would give the wrong result**. Instead, we can define this asymmetry by modifying the beamsplitter behavior (in practice this is more frequently the physical source of the asymmetry).

18. Many parametric amplifiers take advantage of this to squeeze extra performance out of a pump tone.

19. Even last year I've seen papers coming out that only measure magnitude! We really aren't in the stone age of RF measurement anymore...

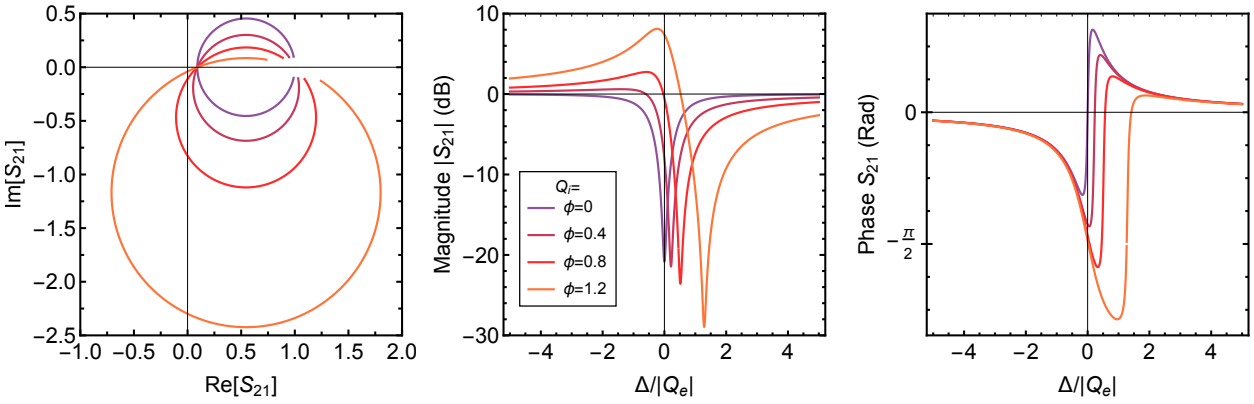

Figure 3.20: Complex transmission response of an asymmetric side-coupled resonator, plotted for various values of asymmetry ϕ (for negative values the effect is simply flipped). Here we have fixed $Q_i = 10Q_e$: notice however that the depth of the transmission dip increases as asymmetry grows making it easy to overestimate Q_i without taking into account the asymmetric response.

Far off resonance, an impedance mismatch on output port 2 results in nonzero reflection $|r| = \sin\phi$ and transmission $|t| = \cos\phi$ less than unity. To account for this while preserving the unitarity of the S matrix, we apply transformations of the form $e^{i\phi}$ to each path of the 3 port network. This gives:

$$S = \frac{1}{2}\begin{pmatrix} \Gamma e^{i\phi} - e^{i\phi} & \Gamma e^{i\phi} + e^{-i\phi} \\ \Gamma e^{i\phi} + e^{-i\phi} & \Gamma e^{i\phi} - e^{i\phi} \end{pmatrix} \tag{3.47}$$

Which in turn yields $S_{21} = (\Gamma e^{i\phi} + e^{-i\phi})/2$. Mapping the linear reflection to the modified 3 port network, we obtain the result used in the main text, which agrees with Refs. [106].

$$S_{21} = 1 - \frac{\kappa}{\kappa+\gamma}\frac{e^{i\phi}}{\cos\phi}\frac{1}{1+2i\delta} \tag{3.48}$$

A similar derivation in Ref. [153] arrives at a very similar result, however the astute researcher will notice that the **denominators** of the two respective expressions differ very

slightly: ours and Ref. [106] has a $\cos\phi$ while the Ref. [153] has a $e^{i\phi}$ term [20]. Keep in mind, this difference is splitting hairs on the end of hairs: if you plot both models on top of each other they line up for nearly any parameter combination (Note: This is not to be confused with the ϕ-rotation method, which rotates but doesn't rescale the circle resulting in a very wrong result). My conclusion from this is that for any practical parameters or measurement precision this doesn't matter[21], however maybe future context may settle the debate.

3.6 Sources of Loss

Throughout this thesis we battle against sources of dissipation in an attempt to maximize device coherence (in other words maximizing Q_i or reducing the intrinsic resonant linewidth or decay rate γ). Ultimately the final quantity we care most about is the single-photon loss $Q_i(\bar{n} = 0)$, since this either closely models the loss of a linear or nonlinear system in the quantum regime. However even with a single linear system there is still a good deal we can learn about how the loss behaves, which can help us determine where the dissipation comes from, where it goes, and thus hopefully how to get rid of it.

3.6.1 Multiple Loss Sources

Returning to the simple resonator model, consider what happens if we separate loss channels, presumably to model the effects of different kinds of loss or decoherence. The lower the resistance in the LCR circuit, the faster the energy decays, so Q_i is lower. For a system with several types of losses, we should expect the worst kind to dominate. In our model we can describe this by adding several resistors in parallel, as shown in Figure 3.21. Thus, for

20. This is a great opportunity to convince yourself that $\mathrm{Re}[Q_e^{-1}] \neq \mathrm{Re}[Q_e]^{-1}$ which I think might be where a lot of the confusion comes from.

21. What is $\mathcal{O}(\phi^2)$ between friends?

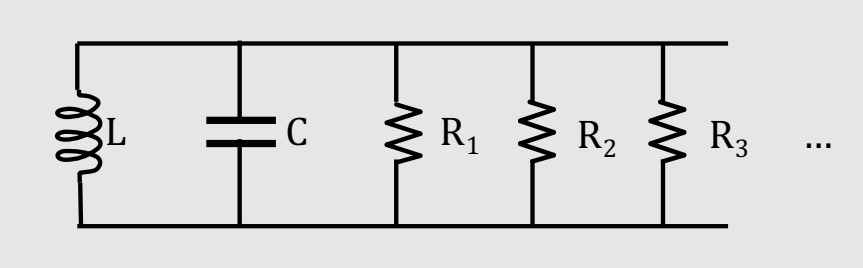

Figure 3.21: LCR resonator with multiple loss channels R_i.

several kinds of loss R_i the overall loss R will be:

$$R^{-1} = R_1^{-1} + R_2^{-1} + R_3^{-1} + ... \tag{3.49}$$

This tells us how internal quality factors add together:

$$Q_i^{-1} = Q_{i,1}^{-1} + Q_{i,2}^{-1} + Q_{i,2}^{-1} + ... = \sum_x Q_x^{-1} \tag{3.50}$$

In this case it may be more useful to think about linewidths γ:

$$\gamma = \sum_x \gamma_x \tag{3.51}$$

Now suppose the source of the loss is from different dielectric materials in a capacitor, so $Q_{i,x} \propto C_x R_x$. For a resonant system where the energy is stored in several different capacitors, the total loss can be written as a weighted sum of the individual losses:

$$Q_i = \sum_x \frac{p_x}{R_x} \tag{3.52}$$

Where the weights are determined by the energy participation ratio $p_x = C_x/C_{\text{tot}} = \mathcal{E}_x/\mathcal{E}_{\text{tot}}$. This was a quick sketch of the process, but this method can be applied to decompose losses in even more complicated quantum systems [155, 241].

3.6.2 Disentangling Losses

Having just determined how loss channels in a linear system add together, we can use this as a tool to help figure out the individual sources of loss (see Ref. [49] for a relatively thorough modern example of how this is done for tantalum microwave resonators). As it turns out, superconducting devices can have a significant number of different contributing loss channels.

Any high-frequency device has to contend with some undesirable leakage into its environment [33, 235]. While we use superconductors to minimize resistivity, involving normal metals will introduce resistive loss [19]. Luckily this can be minimized by ensuring all relevant currents are carried by superconductors[22]. Furthermore any devices that include dielectric materials (which when considering metal oxides are virtually impossible to avoid) we also must consider classical dielectric loss [165, 176]. Since the latter is proportional to the capacitance or electric field, dielectric loss is particularly straightforward to calculate using the participation ratio [54, 241] method outlined above given known loss tangents $Q^{-}1 = \tan\delta$ of bulk materials.

The losses described so far are typically constant, so to determine their relative composition would require rigorous experiments that measure loss while adjusting the relative participation. Some forms of loss however have far simpler control knobs. Recalling from Chapter 2 we remember that the superconducting wires will contribute loss from quasiparticles [39, 145], which given the superconductor properties, we can calculate as a function of temperature T from the complex conductivity integrals [145]:

$$Q_{qp}(T) = \alpha \frac{\sigma_2(T)}{\sigma_1(T)} \tag{3.53}$$

Thus we have a loss component which is a strong function of temperature! Luckily for temperatures below $\sim 0.1T_c$, $Q_{qp} > 10^{12}$ so typically doesn't pose a significant problem to

22. Similarly to electric field participation, one can calculate the inductive energy participation ratio in resistive regions of the circuit.

our superconducting devices (for now). That said, we can use the thermal dependence of Q to extract superconducting properties!

We also place particular emphasis on two level system (TLS) loss [179] from coupled two-level systems (usually also found in dielectrics). Each of these two level systems, regardless of the exact mechanism, can absorb a photon: collectively this process results in average dissipation. When exposed to sufficiently high energies, the individual two level systems saturate and do not absorb additional power: this can come from the thermal heating, external electric or magnetic fields [130], or (conveniently) from the electric fields of the resonator itself. To summarize:

$$\frac{1}{Q_{\mathrm{TLS}}} = \frac{\tanh\left(\frac{\hbar\omega}{2kT}\right)}{Q_{\mathrm{TLS},0}\sqrt{1+\left|\frac{E}{E_{\mathrm{crit}}}\right|^2}} \tag{3.54}$$

Here $Q_{\mathrm{TLS},0} = F\delta_{\mathrm{TLS}}$ contains an effective density of two level systems, and E_{crit} is the average saturating electric field. In practice, we make the assumption that external fields are negligible, and consider that a portion of the resonator electric field contributes to saturating the two level systems, allowing us to write the electric field term $E^2 = \alpha\bar{n}^\beta$ as an empirical function of the resonator photon number $\bar{n}$ [242]. Thus we have a relatively simple model of the power-dependent loss:

$$\frac{1}{Q_{\mathrm{TLS}}(\bar{n},T)} = \frac{\tanh\left(\frac{\hbar\omega}{2kT}\right)}{Q_{\mathrm{TLS},0}\sqrt{1+\left(\frac{\bar{n}}{n_{\mathrm{c}}}\right)^\beta}} \tag{3.55}$$

At low temperatures, the tanh term can be neglected, further simplifying the expression.

Putting everything together, we establish that the internal quality factor of our devices

can be split up into power-dependent and thermal components:

$$\frac{1}{Q_i(\bar{n},T)} = \frac{1}{Q_{\text{other}}} + \frac{1}{Q_{\text{TLS}}(\bar{n},T)} + \frac{1}{Q_{qp}(T)} \tag{3.56}$$

Provided the resonator electric fields are strong enough (frequently true for planar devices, but three-dimensional structures [6] tend to have more dilute electric fields which do not saturate TLSs with reasonable applied power), the power and temperature dependence of Q_i can help determine superconductor properties (eg. T_c) as well as the proportion of loss coming from two level systems. Furthermore, if we understand (or simulate) the energy distribution in the structure, the power-independent remainder of the loss Q_{other} can be decomposed into proportional components using the participation method:

$$Q_{\text{other}} = \sum_i \frac{p_i}{Q_i} \tag{3.57}$$

This makes linear systems a valuable tool for measuring sources of decoherence! Now we can whittle away at sources of loss with a slightly better understanding of where it came from.

CHAPTER 4

NONLINEAR QUANTUM SYSTEMS

Having read this far, you might say- hold on, wasn't all this supposed to be about quantum-mechanical behavior? Indeed. The very same LC resonators we have been discussing thus far can be modelled as ideal quantum harmonic oscillators[1], which can be described with the Hamiltonian:

$$\mathcal{H} = \frac{Q^2}{2C} + \frac{\Phi^2}{2L} \tag{4.1}$$

With conjugate operators Q and Φ (so that $\Phi = L\dot{Q}$) which satisfy the commutation relation $[Q, \Phi] = i\hbar$. For more details on where this comes from and working with quantum circuits in general, see Refs. [170, 187, 189] Similar to the mechanical quantum harmonic oscillator, we can transform this with ladder operators:

$$\mathcal{H} = \hbar\omega_0 \left(a^\dagger a + \frac{1}{2} \right) \tag{4.2}$$

Where the collapse operator a satisfies $[a, a^\dagger] = 1$. This harmonic system has evenly spaced energy levels, as shown in Figure 4.1, making them indistinguishable from each other. An applied signal at the resonance frequency ω_0 will continue to excite the harmonic oscillator indefinitely.

The key to observing quantum phenomena is introducing photon interactions, or in other words nonlinearity. Consider the modified LC circuit on the right of Figure 4.1, which replaces the inductor with a nonlinear current-dependent inductance:

$$L(I) = L \left(1 + \frac{I^2}{I_*^2} \right) \tag{4.3}$$

The resulting Hamiltonian will gain a quartic term [99, 170, 187] and can be expressed in

1. Traditionally every thesis must include an introduction to the harmonic oscillator

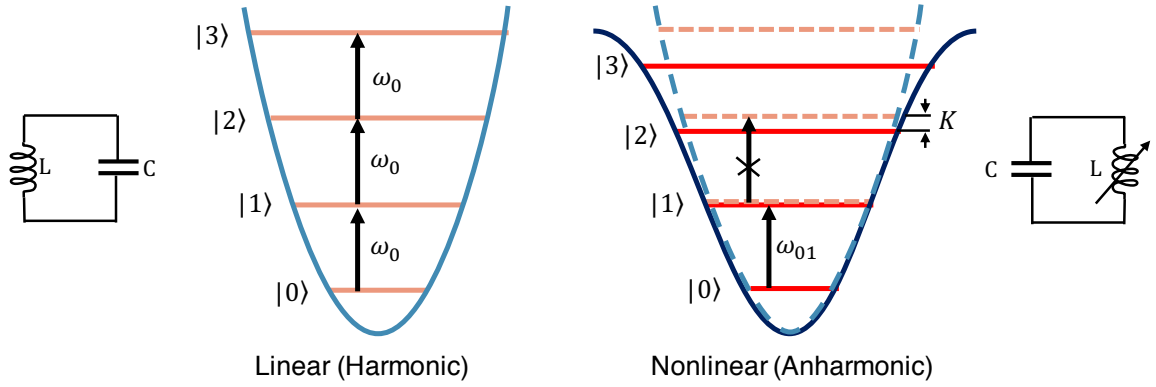

Figure 4.1: Energy levels of a harmonic oscillator (left) which are all evenly spaced, and an anharmonic oscillator (right) whose potential isolates the energy levels making them individually addressable.

terms of an anharmonicity, or Kerr nonlinearity K [104, 239]:

$$\mathcal{H} = \frac{Q^2}{2C} + \frac{\Phi^2}{2L} + \frac{\Phi^4}{4I_*^2 L^3} \tag{4.4}$$

$$\simeq \hbar\omega_0 \left(a^\dagger a + \frac{1}{2}\right) + \hbar\frac{K}{2} a^\dagger a^\dagger a a \tag{4.5}$$

$$= \hbar(\omega_0 + K a^\dagger a)\left(a^\dagger a + \frac{1}{2}\right) \tag{4.6}$$

We suggestively absorbed the scale of the nonlinearity into K and written out the spectrum to show that the frequency now depends on the energy level of the system $n = a^\dagger a$:

$$\omega_n = \omega_0 + Kn \tag{4.7}$$

For this nonlinear system, each energy level transition has its own individually addressable frequency separated by the anharmonicity K, allowing us to observe and manipulate individual quantum states!

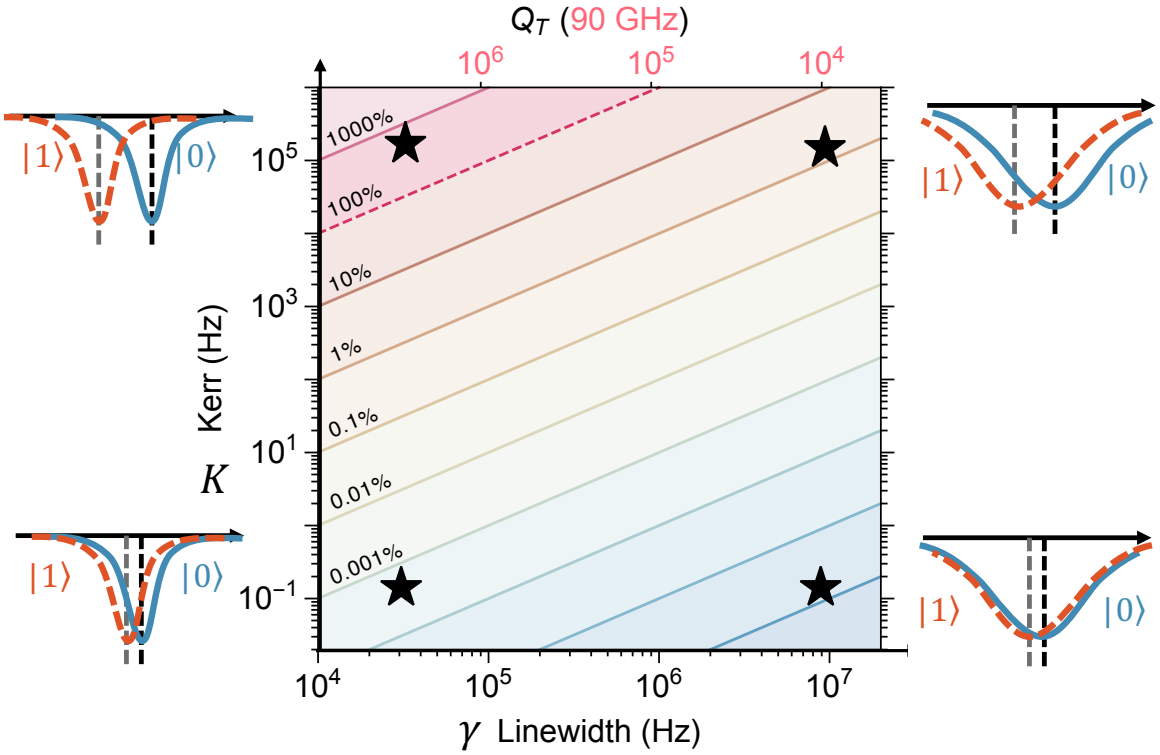

Figure 4.2: Nonlinearity strength relative to linewidth (or loss). For large linewidths or weak nonlinearity, individual photon levels are indistinguishable from one another and the system is governed by classical behavior, and many photons are required to observe the nonlinear effects. Only when the anharmonicity exceeds the linewidth can we address energy levels individually.

Unfortunately real systems have finite linewidth κ, which blur together the delicately separated transition frequencies of the anharmonic system. We illustrate this effect for several cases of linewidth compared to nonlinearity in Figure 4.2. For systems with either weak nonlinearity or wide linewidth the individual transitions are externally indistinguishable, and the levels are not isolated, so an applied signal will continue to excite the system into higher energy levels, like in the harmonic case in Figure 4.1.

The nonlinearity is still present however, so enough photons in the system will engage the nonlinearity enough that the frequency will shift enough for us to observe the effect. If

we ensure that the anharmonicity is greater than the transition linewidth however, as is the case in the top left of Figure 4.2, we can be sure that the first two energy levels are isolated, enabling us to treat the anharmonic oscillator as a two-level system. Throughout this thesis, our goal will be realizing this condition ($K > \kappa$) at millimeter-wave frequencies.

4.2 Kerr Nonlinear Resonator

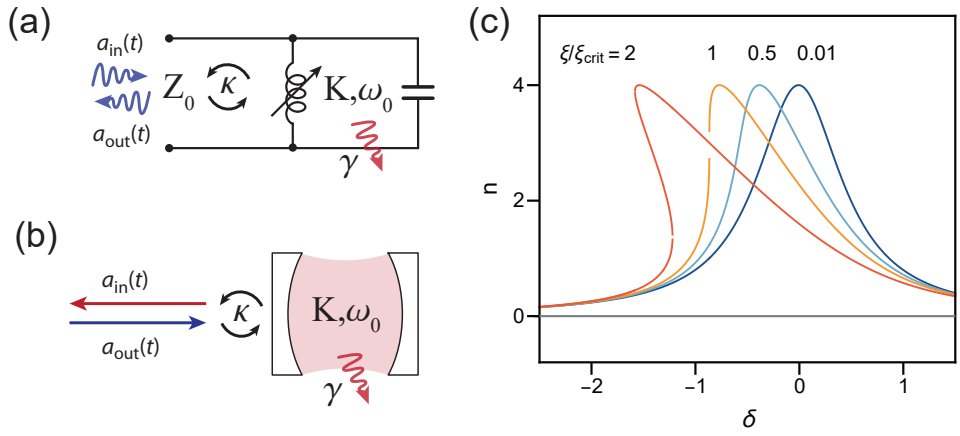

Figure 4.3: a) Circuit model of a nonlinear resonator described by coupling κ, loss γ, frequency ω_0 and nonlinearity K, with input and output fields. b) Optics model of the same system where the resonator is a Kerr-nonlinear optical cavity. c) Normalized intracavity photon number as a function of reduced pump frequency δ for reduced drive strengths ξ, where $\xi_{\text{crit}} = -1/\sqrt{27}$.

Now that we have a hamiltonian written out for a nonlinear resonator, we can use input output formalism [239] to derive its complex response. Consider the models of the system illustrated in Figure 4.3, similar to the direct-coupled resonator measured in reflection discussed earlier in Chapter 2, with an added nonlinear term K. Given that we will observe the nonlinear resonator by letting it interact with photons, instead of finding solutions for the degrees of freedom Q, Φ it is much more useful to determine the dynamics of the operator $a(t)$. The input and output fields[2] $a_{\text{in}}(\omega, t)$ oscillate with frequency ω, and based on the

2. Similar to the operators in a harmonic oscillator, these can be thought of as ladder operators that create propagating photons.

coupling strengths we write out the quantum Langevin equation of motion for $a(t)$ in the resonator:

$$\partial_t a = i[\mathcal{H}, a] - \frac{\kappa + \gamma}{2} a - \sqrt{\kappa} a_{\text{in}} \tag{4.8}$$

Performing the commutator algebra with Equation 4.6 and assuming the resonator field oscillates at the driven frequency ω, we arrive at

$$i(\omega - \omega_0)a + \frac{\kappa + \gamma}{2} a - iK a^\dagger a a^\dagger = \sqrt{\kappa} a_{\text{in}} \tag{4.9}$$

Now we make the key assumption that for weak nonlinearity or large enough excitations, the resonator field can be linearized, allowing us to approximate the operators with a complex-valued function [239], which simplifies the above into something we can solve for a:

$$i(\omega - \omega_0)a + \frac{\kappa + \gamma}{2} a - iK|a|^2 a^* = \sqrt{\kappa} a_{\text{in}} \tag{4.10}$$

We have been careful to use the microwave convention for Fourier transforms, and $n_{\text{ph}} = |a|^2$ corresponds to the average number of photons in the resonator. Multiplying by the complex conjugate, we obtain an equation governing the normalized number of photons in the resonator n.

$$\left(\frac{1}{4} + \delta^2\right) n - 2\delta\xi n^2 + \xi^2 n^3 = 1 \tag{4.11}$$

Where similar to Ref. [61], n, ξ and δ are defined as:

$$n = \frac{|a|^2}{|a_1^{\text{in}}|^2} \frac{(\kappa + \gamma)^2}{\kappa} \tag{4.12}$$

$$\xi = \frac{|a_1^{\text{in}}|^2 \kappa K}{(\kappa + \gamma)^3} \tag{4.13}$$

$$\delta = \frac{\omega - \omega_0}{\kappa + \gamma} \tag{4.14}$$

We plot n as a function of δ for varying drive strengths ξ in Fig. 4.3c, finding that n

reaches a maximum value of 4. At the critical value $\xi_{\text{crit}} = -1/\sqrt{27}$, Eq. 4.17 has 3 real solutions, leading to the onset of bifurcation. Based on the resonator boundary conditions[3] $a_{\text{out}} = a_{\text{in}} + a\sqrt{\kappa}$ [194] and Eq. 4.10, the reflection coefficient Γ will be [61, 239, 264]

$$\Gamma = 1 - \frac{\kappa}{\kappa + \gamma} \frac{1}{\frac{1}{2} + i(\delta - \xi n)} \tag{4.15}$$

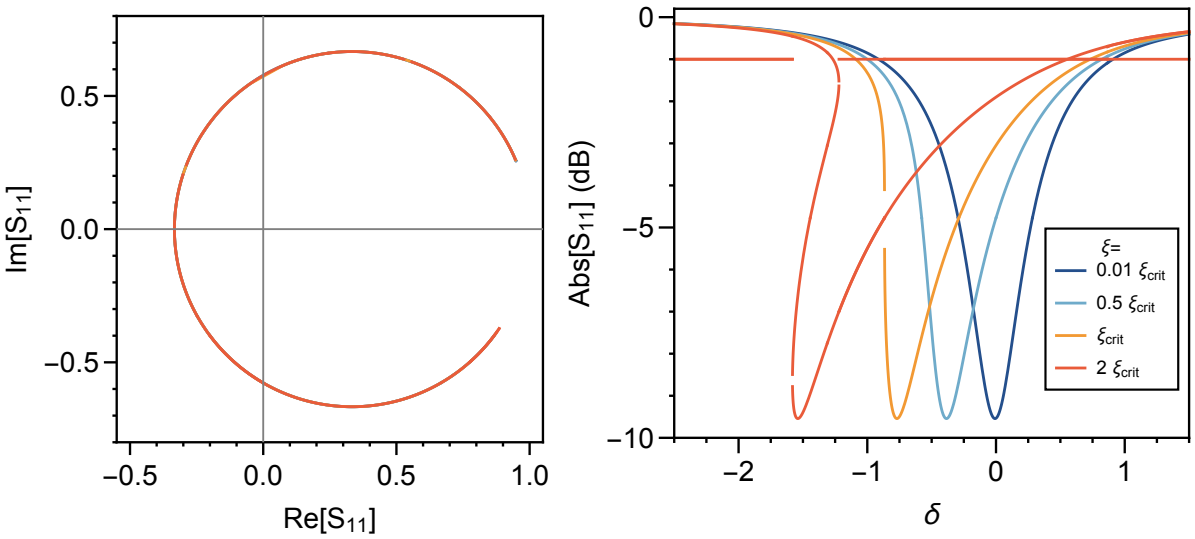

Figure 4.4: Nonlinear reflection response of a Kerr nonlinear resonator in the complex plane (left) and in terms of magnitude (right). We chose $\kappa = 2\gamma$ to make the magnitude response apparent. Notably the circle profile in the complex plane is unaffected by the nonlinearity.

We plot the nonlinear reflection response (Equation 4.15) for increasing drive strengths ξ in Figure 4.4. Above the critical drive strength $\xi_{\text{crit}} = -1/\sqrt{27}$ the transmission bifurcates (similar to the photon number in Figure 4.3): in practice this means depending on the frequency sweep direction we may see hysteresis in the reflection response []. In practice, frequently the measurements are not adiabatic enough to keep the nonlinear resonator in one branch and the response instead probabilistically switches between the solutions[4] Notably

3. Note: we used microwave conventions here. If you want this to match up with Ref. [239] flip the sign.

4. If the measurement is averaged, this can look like a curved response which deviates from the model, however really it is just averaged probability.

the shape of the response in the complex plane does not change- this property is useful for determining whether the nonlinear effect is actually caused by a Kerr nonlinearity, or if it is instead a hysteretic heating effect of some sort.

4.2.1 Side-coupled Nonlinear Resonator

Having already developed a method for translating the physics from a resonator in reflection, it is straightforward to repeat the process above for a side-coupled resonator. Since we separated the effects of nonlinearity and asymmetry, our model can take into account both effects. For our devices we find that this helps explain experimental data, however I will state that the asymmetric effects are based on an approximation, and could likely use additional scrutiny.

One convention difference between the side-coupled case and the direct-coupled reflection case is a factor of two on the input power. This results in a slightly different linearized equation of motion:

$$i(\omega-\omega_0)a+\frac{\kappa+\gamma}{2}a-iK|a|^2a^*=\sqrt{\kappa}a_{\text{in}}=-\sqrt{\frac{\kappa}{2}}a_1^{\text{in}} \tag{4.16}$$

$$\left(\frac{1}{4}+\delta^2\right)n-2\delta\xi n^2+\xi^2n^3=\frac{1}{2} \tag{4.17}$$

Notably this results in a photon number a factor of two smaller than in the reflection case:

$$n_{\text{ph}}(\xi\to 0)=\frac{P_{\text{in}}}{hf_0}\frac{\kappa}{(\kappa+\gamma)^2}\frac{1}{\frac{1}{2}+2\delta^2}\leq\frac{P_{\text{in}}}{\hbar\omega_0}\frac{2\kappa}{(\kappa+\gamma)^2} \tag{4.18}$$

Mapping Eq. 4.15 to the modified 3 port network, we obtain the full nonlinear asymmetric model for a side-coupled resonator, which in respective limits agrees with Refs. [106] and [219].

$$S_{21}=1-\frac{\kappa}{\kappa+\gamma}\frac{e^{i\phi}}{\cos\phi}\frac{1}{1+2i(\delta-\xi n)} \tag{4.19}$$

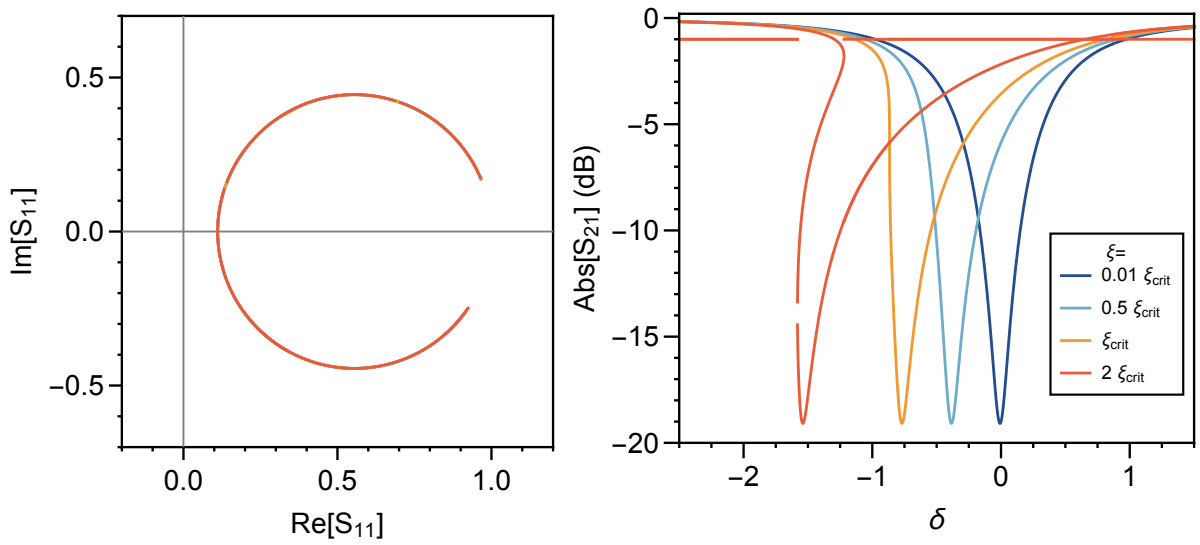

Figure 4.5: Nonlinear transmission response of a side-coupled Kerr nonlinear resonator in the complex plane (left) and in terms of magnitude (right). Here κ is chosen to be 8γ.

We plot the nonlinear transmission response of a side-coupled Kerr nonlinear resonator (Equation 4.19) with no asymmetry ($\phi = 0$) in Figure 4.5 and with significant asymmetry ($\phi = \pm 0.5$) in Figure 4.6. Notice from Equation 4.19 that the asymmetry will only change the response in the complex plane, and will not nominally affect the frequency response (ie the resonator will still trace out the same circle in the complex plane regardless of power). This effect is particularly nice, since it means that regardless of the strange deformations[5] occuring in frequency space, we can still fit the polar response of the resonator and extract quality factors!

However since the Kerr nonlinearity changes the frequency dependence of n, this translates into a pronounced distortion of the regular frequency response. In practice this can also result in strange bifurcation effects, where the hysteretic jump is negative (which seem counterintuitive from a adiabatic sweep argument). Highly asymmetric and nonlinear devices aren't very common, so we look forward to applying this approach to future measured devices to help either experimentally corroborate or further refine the model.

5. You can try fitting a model with three solutions to experimental data, but I guarantee you it will not be a fun time!

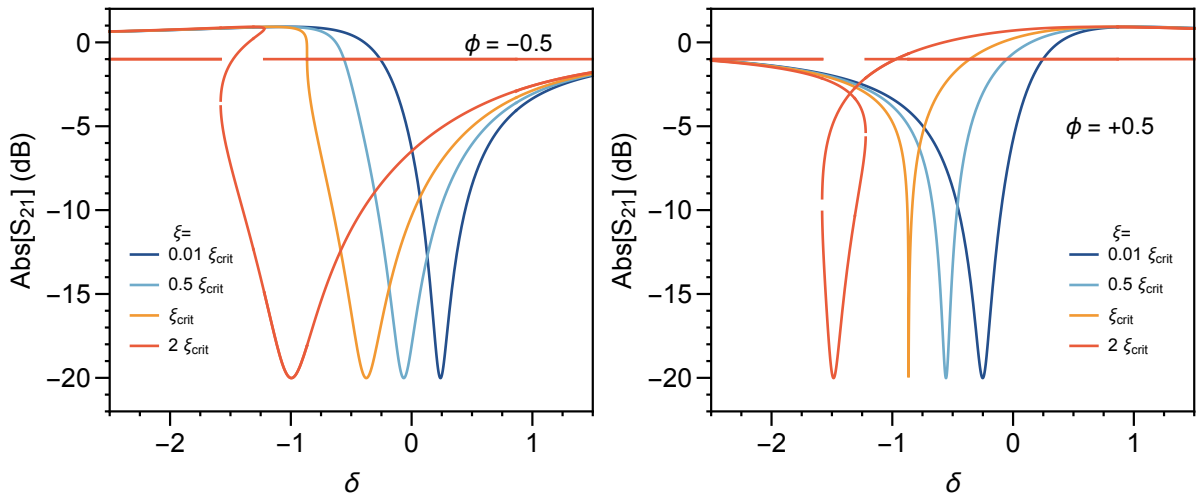

Figure 4.6: Nonlinear transmission response of an asymmetric side-coupled Kerr nonlinear resonator shown for significant asymmetries $\phi = -0.5$ (left) and $\phi = 0.5$. Here κ is also chosen to be 8γ.

4.3 Josephson Junctions

So far we have been a bit vague about the source of the Kerr nonlinearity in our circuit. Conveniently, if we look back to the expression for nonlinear kinetic inductance in Chapter 1, we see that this matches the general form expression we used! However a much more commonly used source of nonlinearity for superconducting quantum devices is the Josephson junction [**?**]

A conventional (SIS)[6] Josephson junction is formed by the interface of two superconductors separated by a thin insulating barrier: for the Cooper pairs to tunnel through the barrier, the thickness must be significantly less than the coherence length of the superconductors on either side. These barriers are typically between 1-2 nm, typically formed by thermally oxidizing a layer of metal (such as aluminum), however they can also be explicitly deposited (such as in the case of AlN barriers [109, 166] Recalling our brief examination of the boundary conditions between two adjacent superconductors from Equation 2.52, we can

6. Superconductor - Insulator - Superconductor

write the voltage and current through the Josephson junction in terms of the superconducting phase difference φ[7]:

$$I_J(t) = \frac{dQ}{dt} = I_c \sin(\varphi(t)) \tag{4.20}$$

$$V_J(t) = \frac{\Phi_0}{2\pi}\frac{d\varphi(t)}{dt} \tag{4.21}$$

Notably, the current will always be below the critical value I_c. Knowing the voltage and current, we can also evaluate the inductance:

$$L = V_J \left(\frac{dI_J}{dt}\right)^{-1} = \frac{L_J}{\cos(\varphi)} \tag{4.22}$$

Where we have collected constants into the Josephson inductance $L_J = \Phi_0/2\pi I_c$, related to the critical current. For our applications, we will be only be running tiny currents through the junction, so we can expand for small φ:

$$L \simeq L_J \left(1 + \frac{\varphi^2}{2} + \mathcal{O}(\varphi^4)\right) \tag{4.23}$$

$$I \simeq I_c\varphi + \mathcal{O}(\varphi^3) \tag{4.24}$$

And thus we recover the generic form of nonlinear inductance we used earlier:

$$L(I) = L_J \left(1 + \frac{I^2}{2I_c^2} + \mathcal{O}(I^4)\right) \tag{4.25}$$

This tells us that for small currents, the Josephson junction generates a Kerr nonlinearity!

7. The circuit phase used in the LC quantization is related to the superconducting phase by $\Phi = \hbar\varphi/2e$ [187], which can be worked out from the voltage expression.

4.3.1 Critical Current from Superconducting Properties

The Josephson junction critical current I_c determines nearly everything about its behavior: knowing the critical current we can directly determine the inductance L_J as well as the Josephson energy $E_J = \Phi_0^2/(2\pi)^2 L_J = \Phi_0 I_c/2\pi$. Importantly, the critical current can be derived from material properties of the junction. For a Josephson junction with two identical superconductors on either side of the tunnel barrier[8], the critical current [10] can be expressed in terms of its normal resistance R_n and the superconducting energy gap $\Delta(T)$:

$$I_c = \frac{\pi}{2e}\frac{1}{R_n}\Delta(T)\tanh\left(\frac{\Delta(T)}{2kT}\right) \tag{4.26}$$

At extremely low temperatures $\Delta \to \Delta_0 = 1.7\text{–}2kT_c$[9], and the tanh term will approach unity, which simplifies the relationship to

$$I_c R_n = \frac{\pi\Delta_0}{2e} \tag{4.27}$$

This $I_c R_n$ product can be directly measured from the cryogenic current voltage relations of a single Josephson junction (For practical experiments it helps if R_n is low and I_c is high), but since it depends solely on bulk superconductor properties, it should be fairly constant for junctions made with the same materials. This a very convenient property for us, since having calibrated the $I_c R_n$ product, we can predict low-temperature junction properties from the

8. For two different superconductors, the density primarily scales with the smaller energy gap of the two along with a correction term [10]

9. Note that since we sometimes use niobium in our junctions, we should make an effort to avoid the $1.76kT_c$ simplification.

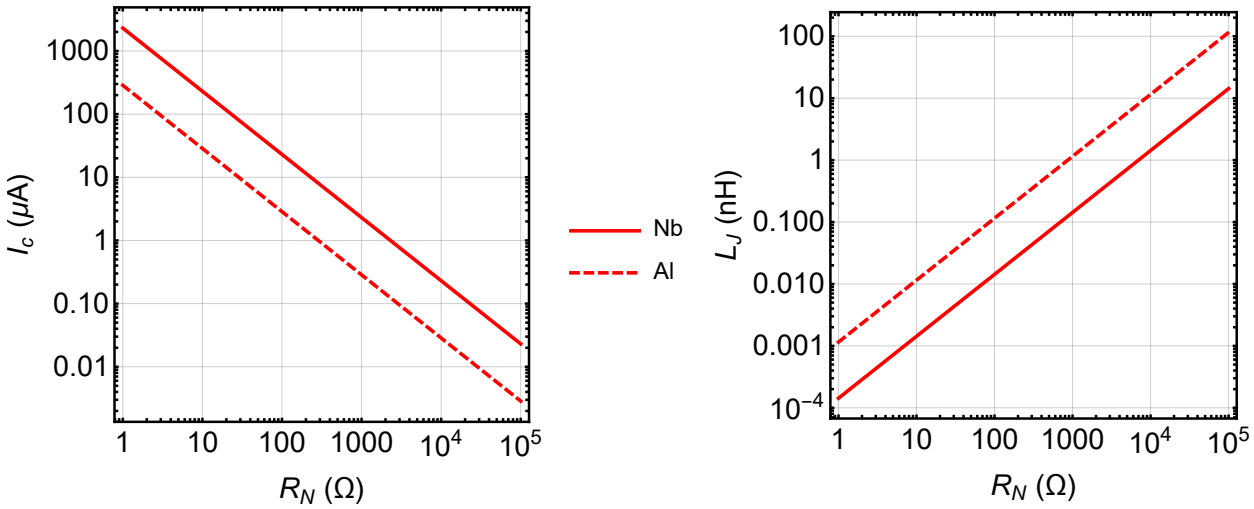

Figure 4.7: Josephson junction critical current and Josephson inductance as a function of room temperature resistance for a junction made with niobium and aluminum. Note the nearly order of magnitude difference between the two materials.

room-temperature junction resistance R_n.

$$L_J = \frac{\hbar}{\pi \Delta_0} R_n \tag{4.28}$$

$$E_J = \frac{\hbar}{(2e)^2} \frac{\pi \Delta_0}{R_n} \tag{4.29}$$

In Figure 4.7, we plot these parameters as a function of the junction room temperature resistance for two common junction superconducting materials: aluminum ($T_c = 1.2$ K, $\Delta = 1.76 k T_c$) and niobium ($T_c = 9.2$ K, $\Delta = 1.85 k T_c$). This allows us to quickly estimate of the junction properties.

The critical current of the junction I_c can also be expressed in terms of a critical current density $J_c = I_c/A$, which will be independent of the specific junction area A and should be constant for a given junction fabrication process. Usually written in kA/cm^2, J_c is useful for comparing different junction processes[10], and junctions made with different gap energies.

10. As a handy conversion for practical sizes, 10 kA/cm^2 = 1μA/μm^2. Because who in their right mind makes centimeter-sized junctions?

The critical current density will depend exponentially on the barrier thickness d [224]:

$$J_c \simeq \frac{n_s \hbar e}{m} \frac{\kappa}{\sinh(2\kappa d)} \tag{4.30}$$

This expression can be derived from the Landau Ginzburg equations applied on the barrier, in this case κ is the characteristic decay length of the order parameter in the insulator (similar to ξ).

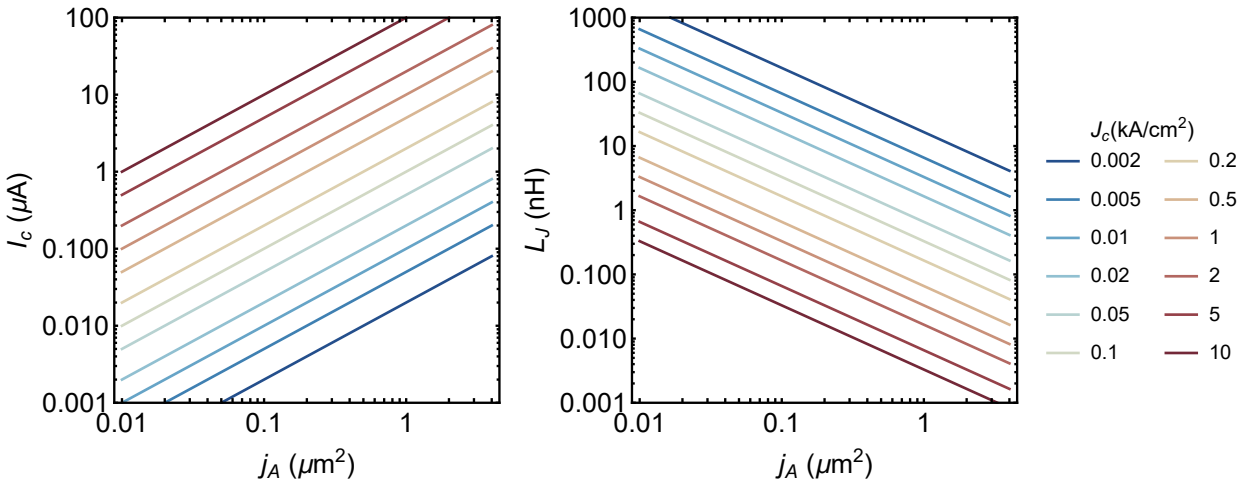

Figure 4.8: Josephson junction critical current and Josephson inductance as a function of junction area for varying junction critical current density J_c. Since the superconducting properties of the junction are encoded in J_c, these relationships are agnostic to junction materials. For the junctions used in microwave qubits, typically $L_J \sim$ 1–10nH.

4.3.2 Junction Capacitance

Since the SIS Josephson junction is formed with two superconducting leads separated by a dielectric, it will also have some natural[11] capacitance between the two leads, as illustrated in Figure 4.9. As we learned before, no capacitance is too small to think about, so we will

11. Not to be confused with organic or farm-raised capacitance.

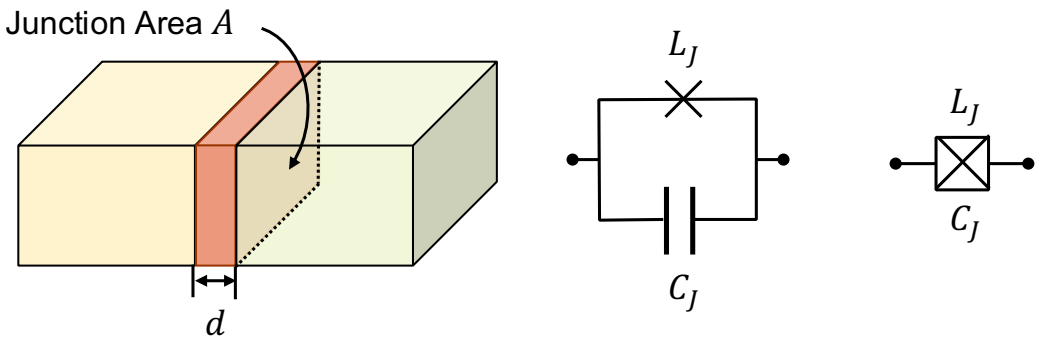

Figure 4.9: Illustration of a generic Josephson junction, highlighting the parallel plate capacitance formed on the barrier. For this reason, junctions are modelled with a capacitance C_J, sometimes abbreviated with a single circuit element.

model the Josephson capacitance as a parallel plate capacitor:

$$C_J = \frac{\epsilon_0 \epsilon_r A}{d} \tag{4.31}$$

For Josephson junctions using aluminum oxide as a barrier (assuming amorphous aluminum oxide, or alumina, $\epsilon_r = 9.8$), the barrier thickness is typically between 1-2 nm. This results in a specific junction capacitance between 44–86 fF/μm^2, depending on the exact barrier thickness[12]. Notably this capacitance scales inversely to the barrier thickness, while the tunneling inductance will scale **exponentially** with the barrier thickness [10, 162, 224]. Since the inductance is exponentially sensitive to barrier thickness, we can usually treat this independently from the capacitance.

As a particularly interesting result of this natural junction capacitance, consider the self-resonant frequency of the junction itself (this will show up as a resonant frequency of the

12. What do you do if you barrier is lumpy like cheese? Probably throw a bit more capacitance in there just to be safe.

circuit even if the junction is shunted by other impedances).

$$\omega_p = \frac{1}{\sqrt{L_J C_J}} = \left(\sqrt{\frac{\Phi_0}{2\pi J_c} \frac{\epsilon_0 \epsilon_r}{d}} \right)^{-1} \tag{4.32}$$

This plasma frequency of the junction is entirely independent of junction area! So given

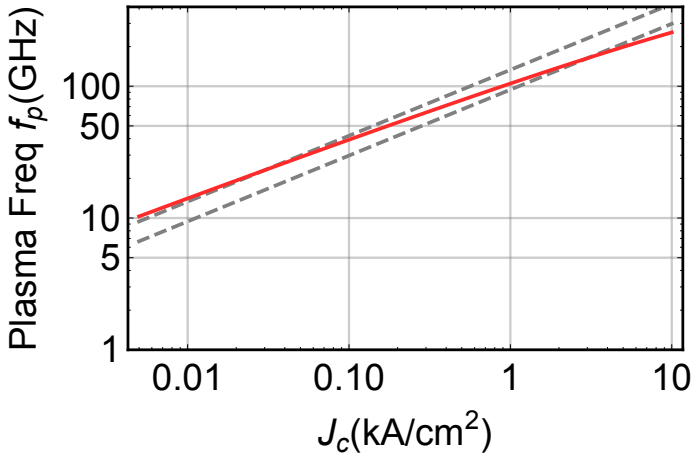

Figure 4.10: Josephson junction plasma frequency, plotted as a function of critical current density J_c, highlighting the need for high-J_c junctions for millimeter-wave applications. Here we assumed that a 1 nm and 2 nm barrier result in 3 kA and 30 A/cm^2 respectively: neglecting the effects of barrier thickness on capacitance (grey dashed lines corresponding to fixed 1 and 2 nm) still gives a pretty good estimate.

a barrier thickness (or having measured specific capacitance and specific inductance) we can determine the self-oscillation frequency for a junction, regardless of its size! This is particularly important for designing circuits at high frequencies, since the lowest resonant frequency of a circuit containing at least one Josephson junction will be determined by the junction plasma frequency[13]. From Figure 4.10 in which we plot the plasma frequency for varying critical current densities, we see that ω_p primarily scales with J_c, meaning that for a high-frequency circuit around 100 GHz we will require junctions with a critical current density above 1 kA/cm^2! Thus high-J_c and high-frequency can be used interchangeably

13. Okay I suppose I can't stop you from using higher modes of a circuit to do what you want, but I promise you it won't be a fun time if you're working above the plasma frequency.

when describing junctions (since one follows from the other).

4.4 Josephson Nonlinear Oscillator

In this section we will conveniently skip over several decades of research, however Refs. [111, 116] can be a useful resource for details about how these superconducting quantum systems work and how they were developed.

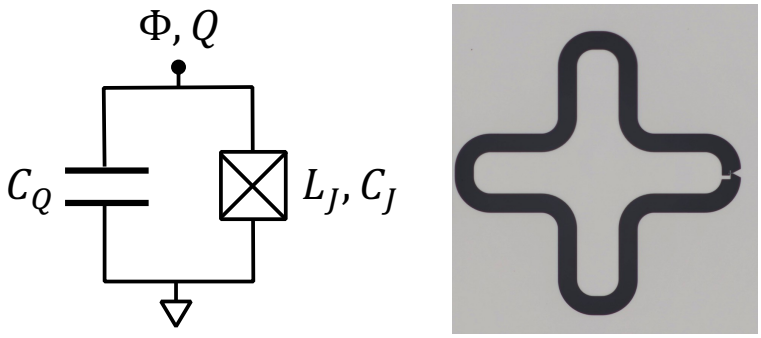

Figure 4.11: Circuit diagram for a simple anharmonic (nonlinear) circuit consisting of a Josephson junction shunted by a capacitance. A microscope image of one of such circuits is shown on the right, where the shunting capacitor has a large cross shape. This particular circuit has a frequency of 19 GHz, and a nonlinearity around 200 MHz.

For resolving quantum phenomena in our superconducting circuits, we would like a resonant system with a high level of nonlinearity or anharmonicity. Due to the low critical current densities they can achieve (which sets the nonlinearity scale, as we showed in Equation 4.25), Josephson junctions are perfect for this task! So to make a single-photon-resolved nonlinearity you will need to obtain the following components:

1. A Josephson junction

That's it! (Check out Refs. [141, 268]). Remember that the inherent capacitance associated with the junction means that the junction itself is an extremely nonlinear resonator already. However most people are unhappy with just the junction resonance itself (it is frequently

above 30 GHz, as you can see in Figure 4.10) and would also like a bit more control over the transition frequency. To do this we can simply adjust the resonant frequency by enhancing the Josephson junction's natural capacitance with a shunt capacitor, as shown in Figure 4.11: and now we have a nonlinear circuit with a controllable frequency and nonlinearity[14]. It is easy to see why this qubit design[113, 178] is so popular!

Josephson junctions in particular make this situation interesting, since their critical current density is directly tied to their inductance by fundamental constants (since $L_J I_c = (\Phi_0/2\pi)^2$). Let's quickly get a feel for what this means for the nonlinearity relative to the frequency[15]. Consider the circuit shown in Figure 4.11. For a small current, we can approximate the resonant frequency as perturbation of an LC oscillator, with total capacitance $C_\Sigma = C_J + C_Q$:

$$\omega = \frac{1}{\sqrt{L_J C_\Sigma}} \tag{4.33}$$

Expanding for small currents, allows us to reform a linearized frequency in terms of $\omega_0^2 = 1/L_J C_\Sigma$:

$$\omega\big|_{I\to 0} \simeq \frac{1}{\sqrt{L_J(1+\frac{I^2}{2I_c^2})C_\Sigma}} \approx \omega_0 - \frac{\omega_0 I^2}{8I_c^2} \tag{4.34}$$

We can rewrite this in terms of inductive energy $E = \frac{1}{2}L_J I^2$ and transform everything into inductance. Making a further assumption that the nonlinearity scale is much smaller than the resonant frequency, or $E_n \approx n\hbar\omega_0$ allows us to obtain a frequency shift per photon:

$$\omega_n \approx \omega_0 - \frac{\omega_0}{8}\left(\frac{2\pi}{\Phi_0}\right)^2 L_J^2 \frac{2E}{L_J} \tag{4.35}$$

$$= \omega_0 - L_J\omega_0^2\frac{e^2}{2\hbar}n \tag{4.36}$$

$$= \omega_0 - \frac{e^2}{2\hbar C_\Sigma}n \tag{4.37}$$

14. As Elzar would say — "Bam!"

15. We're going to play fast and loose with some approximations

Thus the frequency shift per photon, or the Kerr nonlinearity strength for this circuit will be approximately

$$\frac{K}{2\pi} = -\frac{e^2}{2hC_\Sigma} \equiv -\frac{E_C}{h} \tag{4.38}$$

As it turns out the anharmonicity is defined by the charging energy E_C of the total circuit capacitance.

4.4.1 Cooper Pair Box in the Transmon Limit

In deriving the nonlinear behavior of the capacitively-shunted Josephson junction, we took a pretty quick approach. To rectify this let's take a look at the circuit behavior with less approximations. Written out explicitly the Hamiltonian of the circuit in Figure 4.11 will be [113, 116, 187]:

$$\mathcal{H} = \frac{Q^2}{2C_\Sigma} - E_J \cos\varphi \tag{4.39}$$

Using the charging energy $E_C = e^2/2C_\Sigma$ and defining the number of Cooper pairs on the capacitor $n = Q/2e$ simplifies the hamiltonian to the form:

$$\mathcal{H} = 4E_c n^2 - E_J \cos\varphi \tag{4.40}$$

Realistically however, we need a method to interact with this nonlinear system. This is

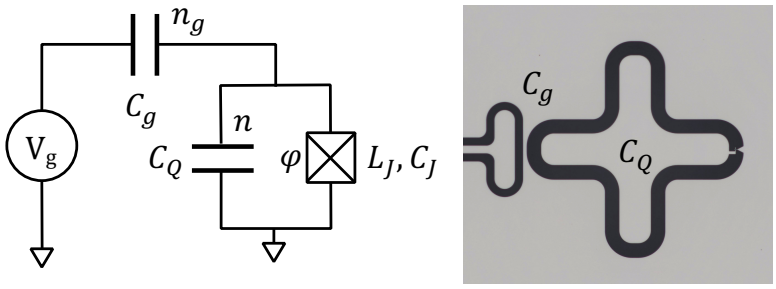

Figure 4.12: Circuit diagram for a cooper pair box, where the capacitively shunted Josephson junction is capacitively coupled to a charging voltage.

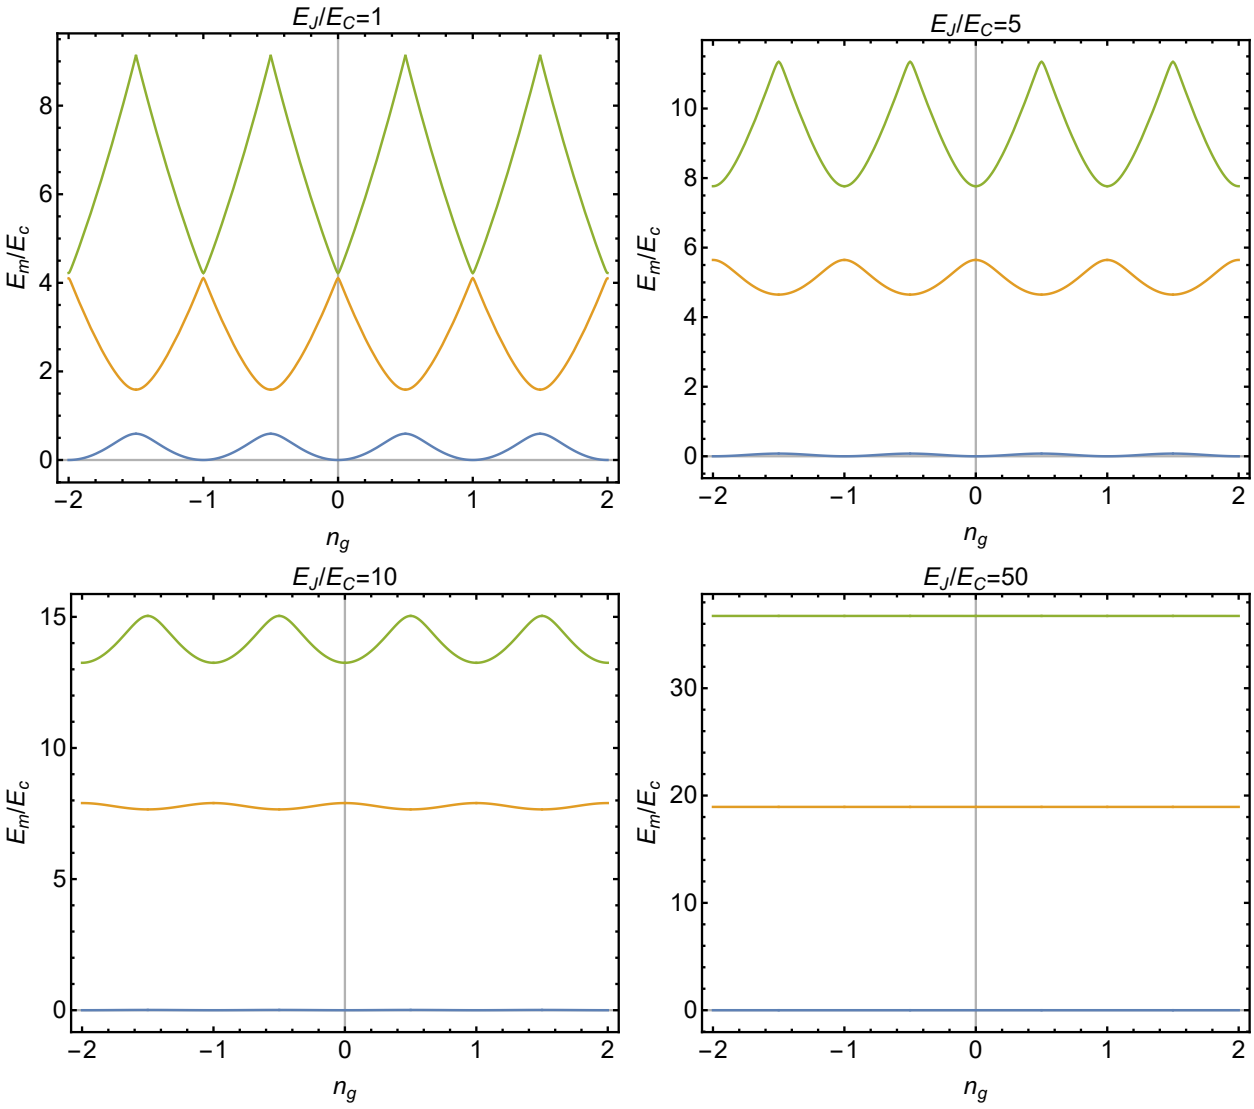

Figure 4.13: Energy levels in the transmon limit. For sufficiently large $E_J \gg E_c$ the energy levels have nearly no dependence on the gate charge n_g making the qubit.

practically accomplished by capacitively coupling a gate voltage V_g to the junction. This modifies the energy spectrum with a gate charge[16] $n_g = C_g V_g/2$. The hamiltonian now becomes:

$$\mathcal{H} = 4E_c(n - n_g)^2 - E_J \cos\varphi \tag{4.41}$$

16. Technically we couple to the qubit through a resonator, so this gate charge will be further modified to include the resonator charge.

This can be solved in the φ basis, obtaining energy solutions[113] in terms of Mathieu functions $a_n(x)$:

$$E_m(n_g) = E_c a_{2(n_g+k_{m,n_g})}(-E_J/2E_C) \tag{4.42}$$

Here $k(m, n_g)$, which is plotted in Figure 4.14 is a sorting function to make the energy

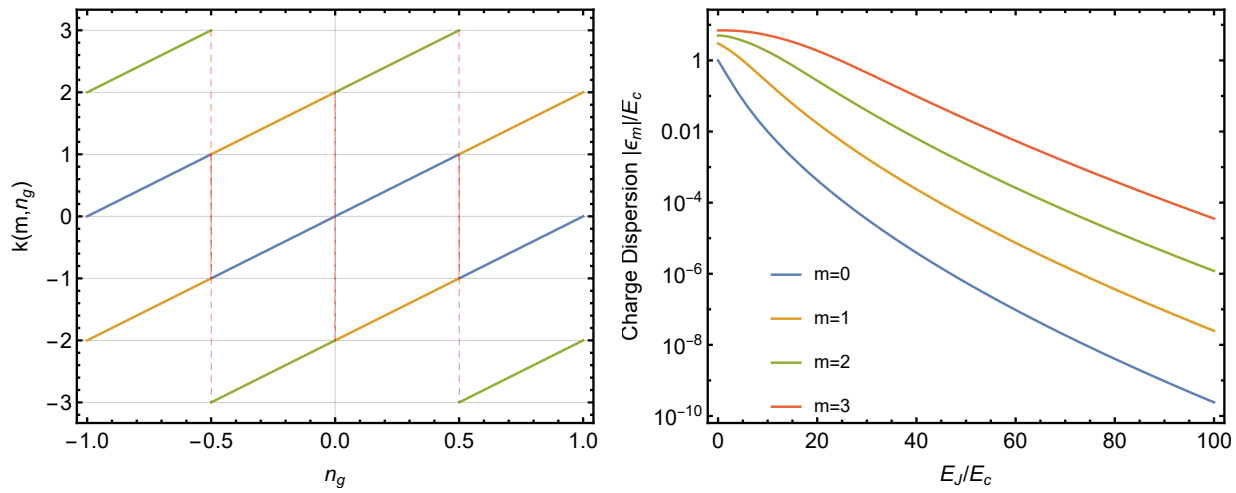

Figure 4.14: Left: sorting function $k(m, n_g)$ used to organize the energy solutions of the transmon. Right: variation of each energy level ϵ_m with respect to the charging energy.

solutions line up. We plot the energy spectrum for varying E_J and E_C in Figure 4.13. As we increase the ratio of E_J to E_C above ~ 20 the energy levels become increasingly flat: this is also explicitly plotted in Figure 4.14. This means that the system will be increasingly less sensitive to charge fluctuations, which would otherwise cause dephasing of the qubit state!

In this limit, the cosine term of the Hamiltonian can be expanded out to fourth order [113], and the full energy spectrum will be:

$$E_n \simeq -E_J + \sqrt{8E_J E_C}\left(m + \frac{1}{2}\right) - \frac{E_C}{12}(6m^2 + 6m + 3) \tag{4.43}$$

If we consider the difference between these energy levels, we recover the anharmonic frequen-

cies from our approximation earlier which are separated by E_C:

$$hf_{m\to m+1} = \sqrt{8E_J E_C} - E_C(m+1) \tag{4.44}$$

$$f_{m\to m+1} = \frac{1}{2\pi\sqrt{L_J C_\Sigma}} - \frac{E_C}{h}(m+1) \tag{4.45}$$

Here we rewrote the frequencies in terms of the expected resonance frequency for convenience.

4.4.2 Energy Participation

Now that we are more confident about how this circuit behaves, consider the energy stored in the junction[17]. In terms of electric field, each respective capacitance in the circuit stores an energy proportional to $E_x = \frac{1}{2}C_x V^2$. So the proportion of energy stored by the junction capacitance will be:

$$p_J = \frac{\mathcal{E}_J}{\mathcal{E}_{\text{Tot}}} = \frac{C_J}{C_\Sigma} \tag{4.46}$$

Interestingly enough we can write this out in terms of qubit frequency relative to the junction plasma frequency:

$$p_J = \frac{C_J L_J}{C_\Sigma L_J} = \left(\frac{\omega_q}{\omega_p}\right)^2 \tag{4.47}$$

4.5 Experiments with a Two Level System

So we have a sufficiently anharmonic oscillator to be a qubit! This means we can treat the first two levels as an isolated system, which in our case are separated by $\hbar\omega_q = \sqrt{8E_J E_C} - E_C$. Luckily this system has been studied very extensively (eg. see Refs. []) We can model this as a with a spin Hamiltonian, expressed in the two level basis in terms of Pauli matrices:

$$\mathcal{H}_q \simeq \hbar\omega_q \frac{\sigma_z}{2} \tag{4.48}$$

17. The inductive energy gets complicated if you include parasitic inductance, so let's not talk about that for now.

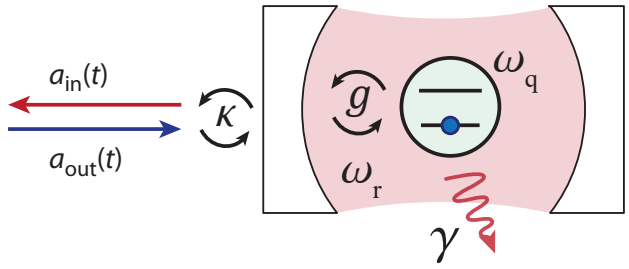

Figure 4.15: Illustration of the interacting cavity (or resonator) and two level system (atom or qubit) characterized by the Jaynes Cummings hamiltonian

4.5.1 Qubit Resonator Interactions

Now instead of measuring the qubit directly, we can use concepts from cavity quantum electrodynamics and instead couple our artificial two-level system to a cavity (a resonator in our case) and take advantage of the dispersive interactions between them [239]. Using the rotating wave approximation (neglecting the counter-rotating terms that oscillate at $\omega_q + \omega_r$ etc.) [239], this coupled system can be modelled by the Jaynes-Cummings Hamiltonian:

$$\mathcal{H}_{JC} = \mathcal{H}_0 + \mathcal{H}_{\text{int}} = \hbar\omega_r \left(a^\dagger a + \frac{1}{2}\right) + \hbar\omega_q \frac{\sigma_z}{2} + \hbar g(\sigma_+ a + \sigma_- a^\dagger) \tag{4.49}$$

In the limit of weak coupling, or $g \ll \Delta = \omega_q - \omega_r$, we can transform to the interacting frame ($\exp[\frac{i}{\hbar\Delta}\mathcal{H}_{\text{int}}]$) and describe the system with the dispersive Hamiltonian[]

$$\mathcal{H}_{\text{d}} \approx \hbar(\omega_r + \chi\sigma_z)\left(a^\dagger a + \frac{1}{2}\right) + \hbar\omega_q \frac{\sigma_z}{2} \tag{4.50}$$

From this we observe that the readout resonator will shift by 2χ depending on the state of the qubit. This property is especially useful for reading out the qubit state! Since the qubit and resonator interact, the resonator state will also impact the qubit, especially if we

regroup the Hamiltonian in a suggestive way:

$$\mathcal{H}_{\mathrm{d}} \approx \hbar\omega_r \left(a^\dagger a + \frac{1}{2}\right) + \frac{\hbar}{2}\left(\omega_q + 2\chi a^\dagger a + \chi\right)\sigma_z \tag{4.51}$$

If the anharmonic nature of the qubit is taken into account (ie the next highest levels, which are separated by K), the dispersive shift will really be:

$$\chi = -\frac{g^2}{\Delta}\left(\frac{1}{1+\Delta/K}\right) \tag{4.52}$$

With this state-dependent dispersive shift we can use weak measurements of the readout resonator to non-destructively measure the qubit state [113, 200].

Microwave signals applied on the system allow us to control the qubit state [168], adding an interaction similar to the electric dipole interactions used to control a spin.

$$\mathcal{H} \approx \hbar(\omega_r + \chi\sigma_z)\left(a^\dagger a + \frac{1}{2}\right) + \frac{\hbar\omega_q}{2}\sigma_z + \hbar\Omega_0 2\sigma_x \tag{4.53}$$

If we weak drive the system $\Omega_0 \ll K$ with a low detuning $\Delta_\Omega = \omega - \omega_q$ (ie close to the ground state transition $\Delta_\Omega \ll K$) we can simplify the system in the drive frame [168] and only consider the qubit

$$\mathcal{H}_\Omega \approx \frac{\hbar\Delta_\Omega}{2}\sigma_z + \hbar\frac{\hbar\Omega_0}{2}\sigma_x \tag{4.54}$$

This produces new dressed eigenstates with energies

$$E_\pm = -\frac{\hbar\Omega_0}{2} \pm \frac{\hbar}{2}\sqrt{\Omega_0^2 + \Delta_\Omega^2} \tag{4.55}$$

On resonance $\Delta_\Omega \to 0$, this effect is well described by an oscillation[18] between the ground and excited states of the qubit at the Rabi frequency Ω_0. This process gives us coherent

18. Also sometimes called Rabi flopping.

control over the qubit state.

4.5.2 Rabi Oscillations

As it is simpler to deal with static hamiltonians, experimental Rabi oscillations can be described by evolving the qubit driven by a continuous weak applied drive signal, as described above. In this way, we can approximate the final state of the qubit after a pulse of a particular length t by the driven qubit state at time t.

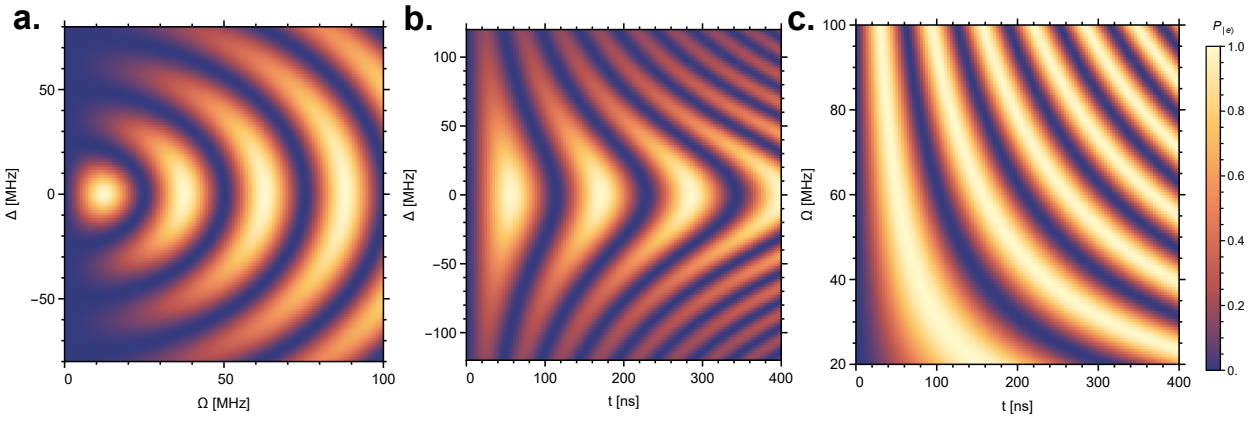

Figure 4.16: Excited state population after being driven by a signal with amplitude Ω detuned from the qubit by Δ for some time t. a) Amplitude-freqency Rabi oscillations increase in frequency off-resonance while also changing in contrast as a function of amplitude, resulting in a WiFi pattern. b) The most popular chevron pattern appears from time-frequency Rabi oscillations, for which contrast is a fixed function of detuning. c) At $\Delta = 0$, the oscillation rate can be expressed as a function of pulse area Ωt, so we observe constant population fringes for curves of constant Ωt.

For a continuous signal detuned from the qubit transition by Δ, the excited state probability of the qubit after starting in its ground state would be given by [26]

$$P_e(t) = \frac{\Omega^2}{\Omega^2 + \Delta^2} \sin^2\left(\frac{\Omega_T}{2} t\right) \tag{4.56}$$

We have dropped a few subscripts for simplicity and used the effective Rabi rate $\Omega_T^2 = \Omega^2 + \Delta^2$. From this expression we observe that on-resonance ($\Delta = 0$) the behavior is oscillatory

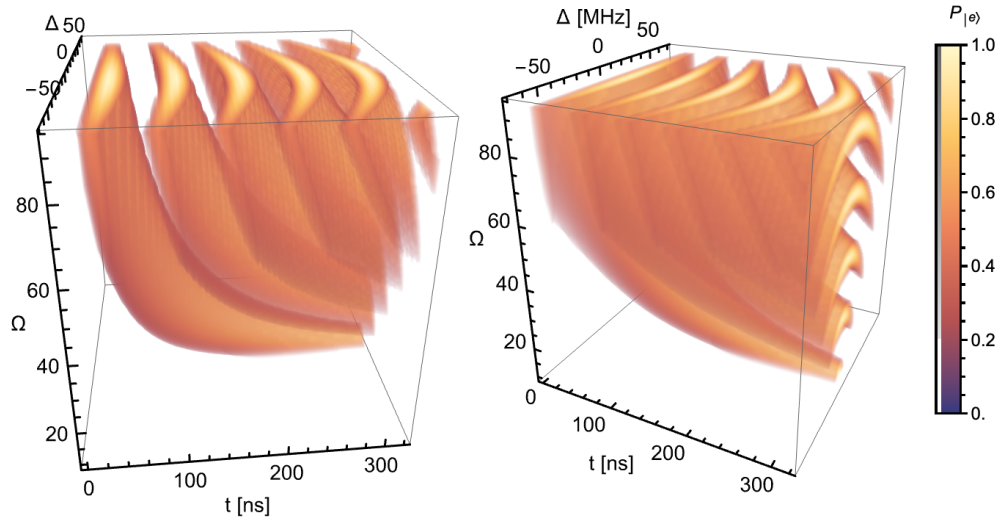

Figure 4.17: Rabi oscillations, shown by plotting the excited state population as a function of drive amplitude Ω (in MHz), detuning from the qubit Δ (in MHz), and time t (in ns). Section cuts from respective axes yield the plots in Figure 4.16.

with respect to the pulse area Ωt. Off-resonance, the oscillations become a function of pulse amplitude Ω and time t, while still achieving maximum population when $\Omega_T t = m\pi$ for odd integer m. We summarize the behavior of the oscillations (Equation 4.56) in Figure 4.16 and Figure 4.17, which shows different fringes with respect to each variable. While the frequency (or detuning) dependence varies whether amplitude or time is being swept, either will result in oscillatory behavior.

Since a square pulse isn't perfectly sharp in practice [19] and has quite a wide Fourier transform, we frequently instead use more rounded pulses which have narrower bandwidth. Thus in experiments we instead explore Rabi oscillation dynamics by applying fixed-length Gaussian pulses near the qubit frequency. These have the following pulse envelope function:

$$\Omega(t) = \Omega_0 \exp\left[-\frac{t^2}{2\sigma^2}\right], \qquad -n\sigma < t < n\sigma \tag{4.57}$$

For practical purposes, the pulse length is finite, which is achieved by truncating the Gaussian

19. Real pulses have curves!

envelope at $\pm n\sigma$. The finite Gaussian nature of the pulse we use complicates the qubit evolution however, since the Rabi evolution rate is non-uniform during the pulse [67] and the truncated pulse edges introduce non-adiabaticity in the qubit state evolution [24]. For a truncated Gaussian pulse described in Eq. 4.57, the excited state population can be modelled by [26, 67]:

$$P_e = \frac{\Omega^2 e^{-n^2}}{\Omega^2 e^{-n^2} + \Delta^2} \sin^2\left(\frac{\sigma}{2}\int_{-n}^{n} \partial\tau \sqrt{\Omega^2 e^{-\tau^2} + \Delta^2}\right) \tag{4.58}$$

On resonance ($\Delta = 0$) we recover the familiar expression $P_e = \sin^2(\pi\sigma\Omega)$. However off resonance the expected power-broadening of the Rabi oscillations in frequency space is reduced, as shown in Figure 4.18.

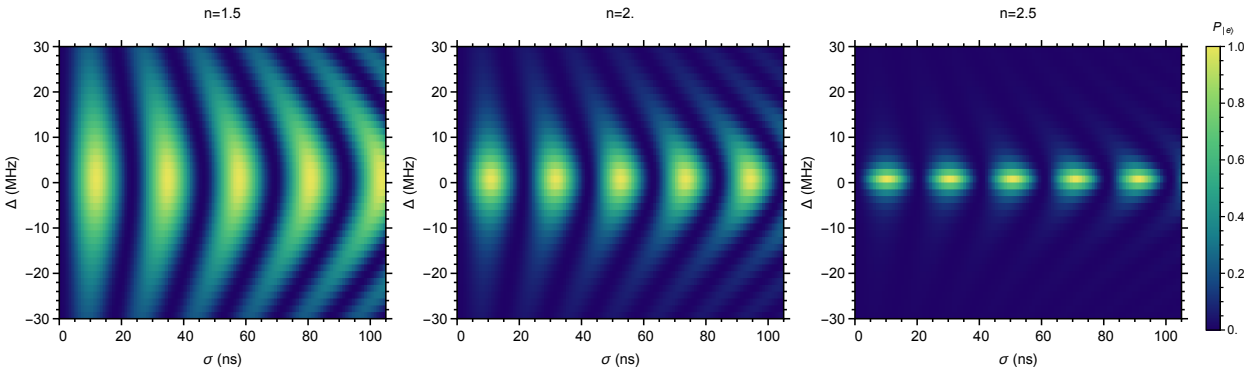

Figure 4.18: Excited state population following a finite-length Gaussian pulse terminated at $\pm n\sigma$, shown for different values of n. For smaller n the pulse shape is closer to a square pulse, and the fringes exhibit significant power broadening. As more of the Gaussian envelope is used the fringe bandwidth is reduced, limiting the power broadening effect. Here the pulse lengths σ have been adjusted to yield similar oscillation rates at $\Delta = 0$.

For small n the behavior is similar to that of a square pulse [26], however as we include more of the Gaussian profile by increasing the cutoff length n, the bandwidth of the oscillations decreases while the frequency scaling of the oscillations remains constant. While this model neglects bandwidth effects on the pulse envelope reaching the qubit, the behavior shown in Figure 4.18 qualitatively matches experimental behavior (see Chapter 8).

4.6 Dissipation for a Two-Level System

Whether from incident noise or from energy decay, real quantum systems couple to their environment to some degree, resulting in decoherence. To describe the quantum dynamics of a qubit taking into account sources of loss we use collapse operators C_k. The density matrix of the system ρ will be governed by the Liouville equation:

$$\frac{\partial \rho}{\partial t} = -\frac{i}{\hbar}[H, \rho] + \sum_k \left(2C_k \rho C_k^\dagger - \rho C_k^\dagger C_k - C_k^\dagger C_k \rho \right) \tag{4.59}$$

Let's first consider energy relaxation. Modelling the qubit as a two-level system (and neglecting the physical temperature of the coupled bath), the collapse operator will be

$$C_1 = \sqrt{\gamma_1}\sigma_- = \begin{pmatrix} 0 & 0 \\ \sqrt{\gamma_1} & 0 \end{pmatrix} \tag{4.60}$$

In terms of the density matrix elements ρ_{ij}, the equation of motion for the two level system becomes:

$$\frac{\partial \rho}{\partial t} = \begin{pmatrix} -\gamma_1 \rho_{11} & -i\omega_q - \frac{\gamma_1}{2}\rho_{12} \\ i\omega_q - \frac{\gamma_1}{2}\rho_{21} & \gamma_1 \rho_{11} \end{pmatrix} \tag{4.61}$$

Suppose we begin in the excited state $\rho = |1\rangle\langle 1|$. Then the two level system will exponentially decay into its ground state over time, with $\rho_{11} = e^{-\gamma_1 t}$. This characteristic decay rate γ_1 corresponds the energy relaxation time $T_1 = 1/\gamma_1$. Now suppose we begin in a superposition state: the off-diagonal terms would be nonzero, but decay at a rate of $\gamma_1/2$. Thus relaxation also results in decay of phase information, or in other words dephasing!

Consider the addition of a pure dephasing operator:

$$C_2 = \sqrt{\frac{\gamma_\phi}{2}}\sigma_z = \sqrt{\frac{\gamma_\phi}{2}} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \tag{4.62}$$

The equation of motion gains a term of the form $\frac{\gamma_\phi}{2}(\sigma_z\rho\sigma_z - \rho)$, which results in modified off-diagonal terms:

$$\frac{\partial\rho}{\partial t} = \begin{pmatrix} -\gamma_1\rho_{11} & -i\omega_q - (\frac{\gamma_1}{2} + \gamma_\phi)\rho_{12} \\ i\omega_q - (\frac{\gamma_1}{2} + \gamma_\phi)\rho_{21} & \gamma_1\rho_{11} \end{pmatrix} \tag{4.63}$$

The off diagonal terms decay at a rate of $\gamma_\phi + \gamma_1/2$. Thus we can describe the dephasing of the two-level system quantum state in terms of a pure dephasing rate $\gamma_\phi = 1/T_\phi$:

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi} \tag{4.64}$$

As a result, the dephasing rate of the two-level system is limited by its relaxation rate, or $T_2 \leq 2T_1$. We can experimentally extract the pure dephasing rate T_ϕ by measuring both T_1 as well as T_2! As a final note, suppose the qubit frequency ω_q isn't constant: the jitter introduced in the off-diagonal terms appears as additional dephasing when averaged.

CHAPTER 5

MILLIMETER-WAVE FOUR-WAVE MIXING

Recent interest in next-generation communication devices [29, 171] has led to important advances in sensitive millimeter-wave measurement technology around 100 GHz. Realizing quantum systems at these frequencies however requires both the demonstration of low-loss components — device materials with low absorption rates [34, 41, 265] and resonators with long photon lifetimes [62, 71, 85, 114, 118, 205] — and most importantly, elements providing nonlinear interactions, which for circuit quantum optics can be realized with four-wave mixing Kerr terms in the Hamiltonian. One approach commonly used at microwave frequencies relies on aluminum Josephson junctions [32], which yield necessary four-wave mixing at low powers. However to avoid breaking Cooper pairs with high-frequency photons, devices at millimeter-wave frequencies are limited to materials with higher superconducting critical temperatures (T_c). Higher T_c junctions have been implemented as high-frequency mixers for millimeter-wave detection [103, 151, 230], and ongoing efforts are improving losses for quantum applications [79, 175].

Kinetic inductance (KI) offers a promising alternative source of Kerr nonlinearity arising from the inertia of Cooper pairs in a superconductor, gaining recent interest for microwave quantum applications [196, 204], and has also been successfully used for millimeter-wave detection [83, 172]. Niobium Nitride (NbN) is an ideal material for KI, as it has a high intrinsic sheet inductance [22, 93, 101], a relatively high T_c between 14-18 K [22, 93, 101, 210] making it suitable for high-frequency applications [265], and has good microwave loss properties [169]. Among deposition methods, atomic layer deposition (ALD) offers conformal growth of NbN [210] and promising avenues for realizing repeatable high KI devices on a wafer-scale [204].

In this chapter, we explore kinetic inductance as a nonlinear element for quantum devices at millimeter-wave frequencies using high KI resonators in the W-Band (75-110GHz)

fabricated from thin films of NbN deposited via ALD. We study potential loss mechanisms at powers ranging down to single photon occupations. Using the power dependent frequency shift, we characterize the nonlinearity arising from KI, the strength of which varies with wire width and material properties. With two tone spectroscopy, we observe degenerate four-wave mixing near single photon powers. These measurements demonstrate the necessary core components for millimeter-wave circuit quantum optics, paving the way for a new generation of high-frequency high-temperature experiments.

5.1 High-KI Niobium Nitride Films

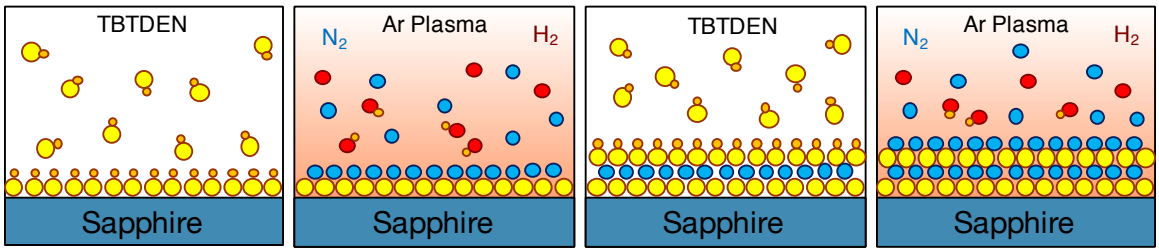

Figure 5.1: Illustration of the plasma-enhanced atomic layer deposition process (PEALD or ALD for short). After an atomic monolayer of the precursor (TBTDEN) coats the surface, a nitrogen and hydrogen plasma reacts with the TBTDEN ligand and incorporates nitrogen into the niobium matrix.

In order to understand the quality of the NbN films grown by ALD and accurately predict resonant frequencies, we characterize material properties with DC electrical measurements. The niobium nitride devices are deposited on sapphire with a process based on Ref. [210], and etched with a fluorine based inductively coupled plasma.

DC film characterization measurements were performed in a Quantum Design Physical Property Measurement System (PPMS) with a base temperature of 1.8 K. Test structures consisting of a 1.5 mm$\times$40 μm wire were patterned on 7×7 mm chips going through the process described above along with device wafers, then wirebonded for four-wire measurements.

Finished structures were kept in a low $\sim$ 500 mTorr vacuum[1] in an effort to minimize oxide growth prior to measurement, as we observed decreases up to 1 K in critical temperatures for samples aged several days in atmosphere, likely a result of oxide growth [152] reducing the superconducting film thickness.

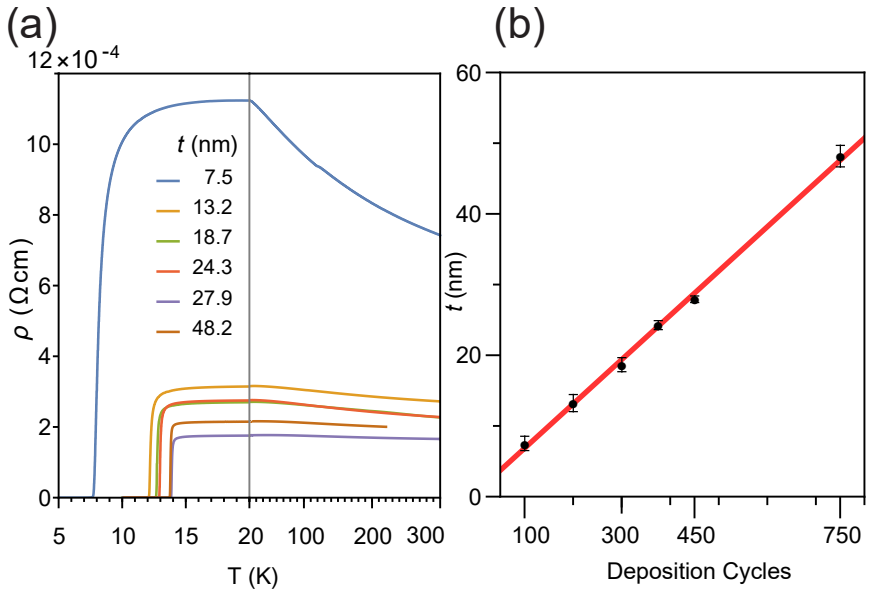

Figure 5.2: Film measurements (a) Measured resistivity as a function of temperature showing decreasing resistivity with increasing temperature above the superconducting transition characteristic for NbN. (b) Thicknesses measured by profilometry as a function of deposition cycles, with a linear fit overlaid in red. We extract a growth rate of 0.63 Å per cycle.

After cooling the samples to 10 K (3 K in the case of the 8 nm film) in ambient magnetic fields, we verified that the residual resistance of the film dropped below the instrument noise floor of around $5 \times 10^{-3}\Omega$. After thermalizing for one hour, the samples were warmed up to 20 K at a rate of 0.1 K/min, then warmed to 300 K at a rate of 1 K/min. In Fig. 5.2(a) we plot measured resistivity as a function of temperature for the films in this chapter, which we use to extract T_C, ρ_n and calculate $R_\square$ and $L_\square$ for the films. Similar to previous studies [169], resistivity decreases with temperature above the superconducting transition, which is typical for strongly disordered materials [159, 169].

1. This vacuum was pumped out by hand!

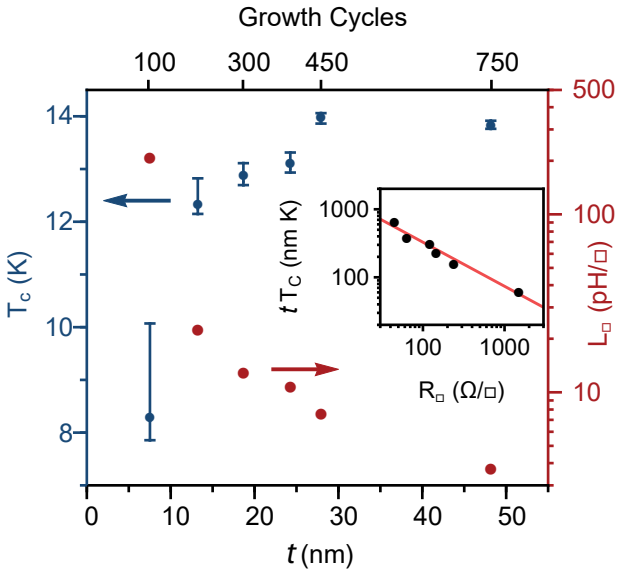

Figure 5.3: Superconducting critical temperature (T_c) and sheet inductance ($L_\square$) of deposited NbN as a function of film thickness (t). Bars denote temperatures corresponding to 90% and 10% reductions from maximum resistivity. The inset shows the dependence of tT_c on $R_\square$ with a fit (red) to $tT_c = AR_\square^{-B}$; $A = 6487 \pm 1607$,$B = 0.647 \pm 0.05$

We measure resistivity at ambient magnetic fields as a function of temperature, which we use to extract T_c for a range of film thicknesses [See Figure 5.3]. The inset also shows that our films follow a universal relation observed for thin superconducting films [93] linking T_c, film thickness t, and sheet resistance $R_\square$: we find that our results are similar to NbN deposited with other methods [22, 93]. For thicker films, T_c appears to saturate at 13.8-13.9 K which is comparable to other materials studies [93, 101, 210], while decreasing to 8.7 K for the thinnest film (t=7 nm), which can be attributed to disorder enhanced Coulomb repulsions [58, 208]. We also find that the superconducting transition width increases significantly for the thinner films, which can in turn be attributed to disorder broadened density of states [58] or reduced vortex-antivortex pairing energies at the transition [160, 169].

From the resistivity and critical temperature, we determine the sheet inductance $L_\square = \hbar R_\square/\pi\Delta_0$, where the normal sheet resistance $R_\square = \rho_n/t$ is taken as the maximum value of normal resistivity ρ_n, occurring just above T_c, and $\Delta_0 = 2.08T_c$ is the superconducting energy gap predicted by BCS theory for NbN [101, 169]. We observe a monotonic increase

in $L_\square$ for thinner films, achieving a maximum $L_\square = 212$ pH/$\square$, comparable to similar high KI films [159, 204].

5.2 Quantum Measurements at Millimeter-wave Frequencies

As we established early on, millimeter-wave systems will remain in their ground state (in the quantum regime) and will be largely unaffected by thermal photons at 1–2 K. Taking advantage of this, we flex the strengths of millimeter-wave circuits and use a custom built ^{4}He adsorption refrigerator to cool all of our millimeter-wave experiments. The adsorption refrigerator, which relies on helium-4 evaporation [183], has a base temperature of 0.86 K, and a cycle duration of approximately 3 hours[2]. We generate millimeter-wave signals (75-115 GHz) at room temperature by sending microwave signals (12-19 GHz) into a frequency multiplier. We convert the generated waveguide TE_{10} mode to a 1 mm diameter stainless steel and beryllium copper coaxial cable, which carries the signal to the 1 K stage of the fridge, thermalizing mechanically at each intermediate stage, then convert back to a WR-10 waveguide which leads to the device under test. The cables and waveguide-cable converters have a combined frequency-dependent loss ranging from 55.6 dB to 75.8 dB in the W-Band, which is dominated by the cable loss. We confirm the attenuation and incident device power at room temperature with a calibrated power meter[3] (Agilent W8486A) and a referenced measurement with a VNAx805 receiver, however cryogenic shifts in cable transmission and minute shifts in waveguide alignment likely result in small variations in transmitted power. We are able to further confirm the applied power by measuring the lowest observed bifurcation point, and find that most bifurcation powers agree with predictions, yielding a maximum combined power uncertainty of approximately ± 5 dBm, which sets the uncertainty in our

2. This cycle time is actually a complete nightmare to work around since it breaks up any long experiments you may want to do.

3. We borrowed the power meter from a CMB astronomy lab! Thank you whoever agreed to this - every single one of our millimeter-wave calibrations can be traced back to this one day of calibration.

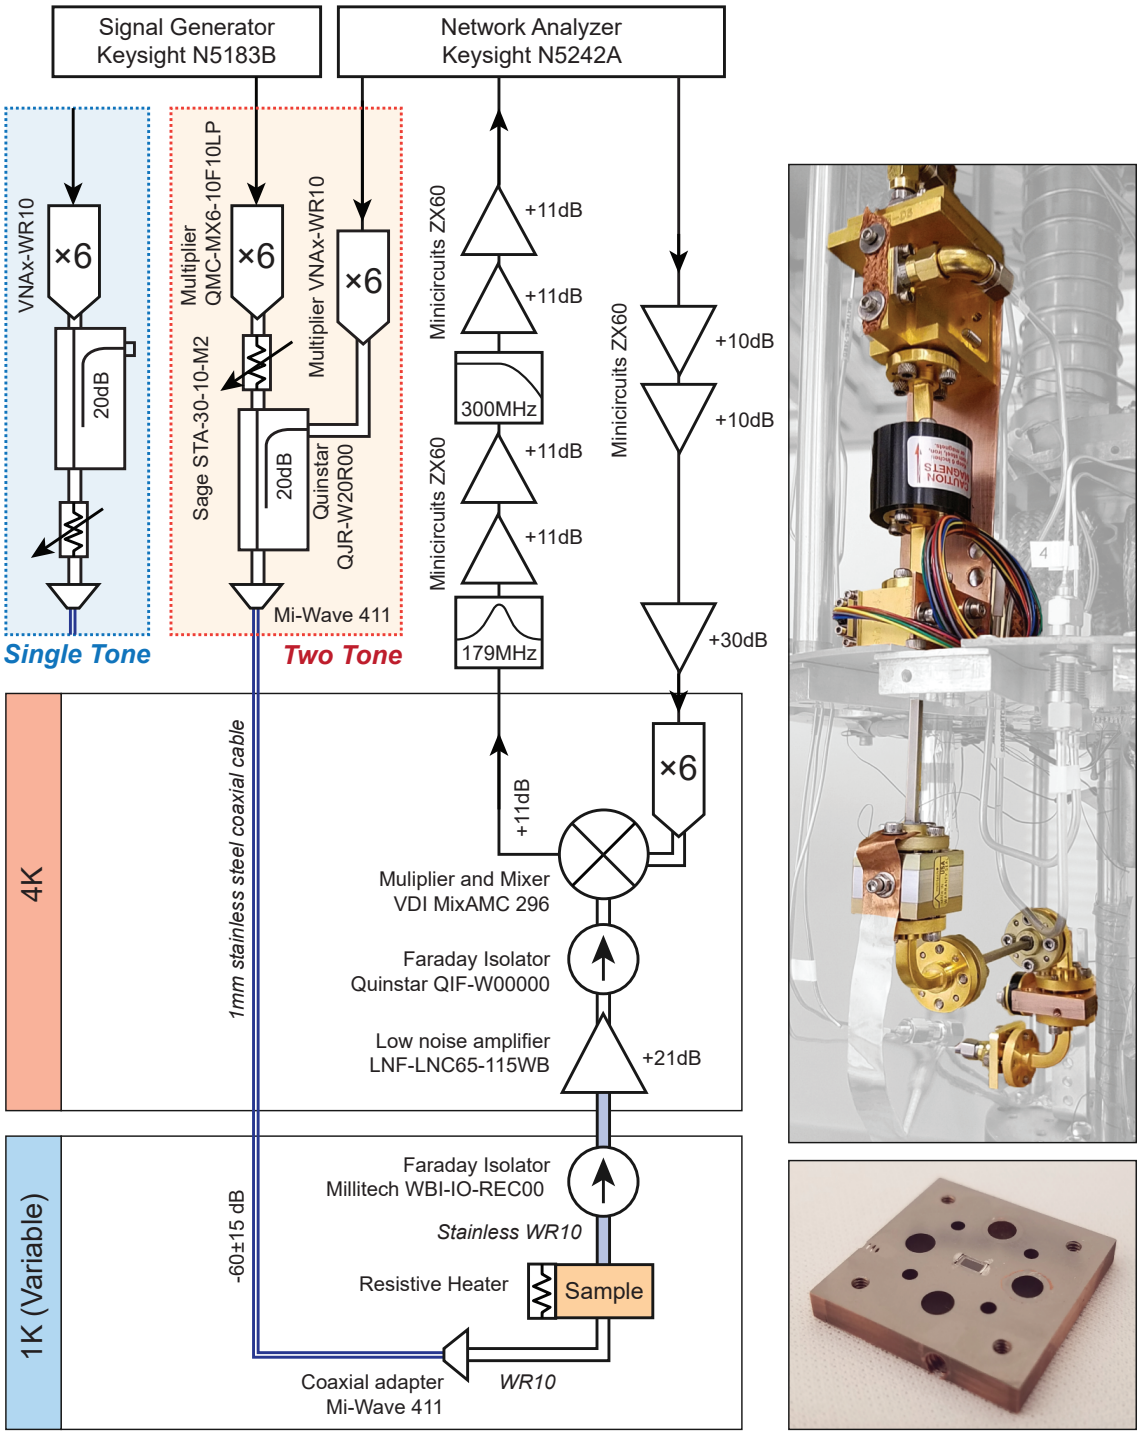

Figure 5.4: Schematic of millimeter-wave measurement setup for single and two-tone configurations. Colored tabs show temperature stages inside the ^{4}He adsorption refrigerator. A photograph on the right highlights relevant hardware inside the fridge. The bottom left shows a photograph of the sample with top waveguide removed.

photon number measurements[4].

The sample is isolated from both millimeter-wave and thermal radiation from the 4 K plate with two stainless steel waveguides 2 inches long and a faraday isolator. Using a resistive heater and a standard curve Ruthenium oxide thermometer we can perform temperature sweeps on the sample without significantly affecting the fridge stage temperatures. A low-noise amplifier ($T_N \sim 28$ K) amplifies the signal before passing through another faraday isolator, which further blocks any leaking signals. The signal then passes to a cryogenic mixer, which converts the signal to radio-wave (100-300 MHz) which we filter, amplify and measure at room temperature with a network analyzer. We control signal power by varying input attenuation and multiplier input power, confirming with room temperature calibrations as described above. For two-tone measurements, we move the signal path to the 20 dB port of the input directional coupler, and add an additional frequency multiplier fed by a reference-locked microwave signal generator. For single-tone measurements, the 20 dB port is capped with a short to minimize incident stray radiation.

5.3 Exploring loss with Millimeter-wave Loss with Resonators

We investigate properties of millimeter-wave high KI resonators in the quantum regime (at temperatures of 1 K) in a Helium-4 adsorption refrigerator. Using a frequency multiplier, cryogenic mixer and low noise amplifier, we measure the complex transmission response as shown in Figure 5.5a. Input attenuation reduces thermal noise reaching the sample, enabling transmission measurements in the single photon limit set by the thermal background. Rectangular waveguides couple the signal in and out of a 200 μm deep slot, the dimensions of which are carefully selected to shift spurious lossy resonances out of the W-band. To reduce potential conductivity losses, the waveguide and slot are coated with 200 nm of evap-

4. Much much later on we performed an AC stark shift calibration, which agreed with our calibration standard to within 3dB! So this calibration was actually pretty decent.

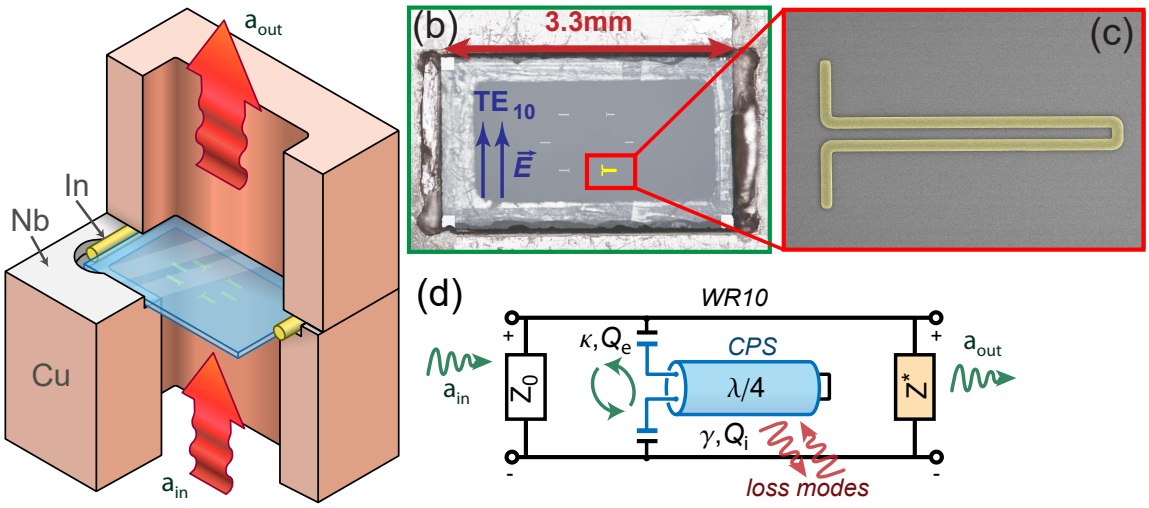

Figure 5.5: Device characterization and design. (a) Six-fold frequency multipliers convert microwave to millimeter-wave signals (green), which are demodulated with a cryogenic mixer. A cutaway shows copper WR-10 rectangular waveguides coupling the signal in and out of a Nb coated slot, into which we mount a chip patterned with 6 resonators. (b) Top down composite micrograph showing a mounted chip with the top waveguide removed. (c) Scanning Electron Micrograph of a typical resonator used in this chapter, with wire width $w = 4$ μm and film thickness $t = 27.8$ nm (NbN false colored yellow). Dipole coupling antennas extend on the left of the quarter-wave resonator. Measurements can be described with input and loss couplings Q_e and Q_i using the circuit model in (d), which takes into account the impedance mismatch between waveguide Z_0 and slot with sapphire chip Z^*.

orated Niobium. Below 9 K, this helps shield the sample from stray magnetic fields; however devices with higher T_c are not shielded from magnetic fields while cooling through their superconducting transition. We use indium to mount a chip patterned with 6 resonators in the slot, as shown in Figure 5.5b. Devices are patterned on 100 μm thick sapphire, which has low dielectric loss and minimizes spurious substrate resonances in the frequency band of interest. The planar resonator geometry shown in Figure 5.5c consists of a shorted quarter wave section of a balanced mode coplanar stripline waveguide (CPS), which couples directly to the TE_{10} waveguide through dipole radiation, which we enhance with dipole antennas. We find that this design is well described by the analytic model presented in Ref. [259] which takes into account the thin film linear kinetic inductance. For very thin or narrow wire widths, the total inductance is dominantly kinetic, making the resonators extremely

sensitive to superconducting film properties.

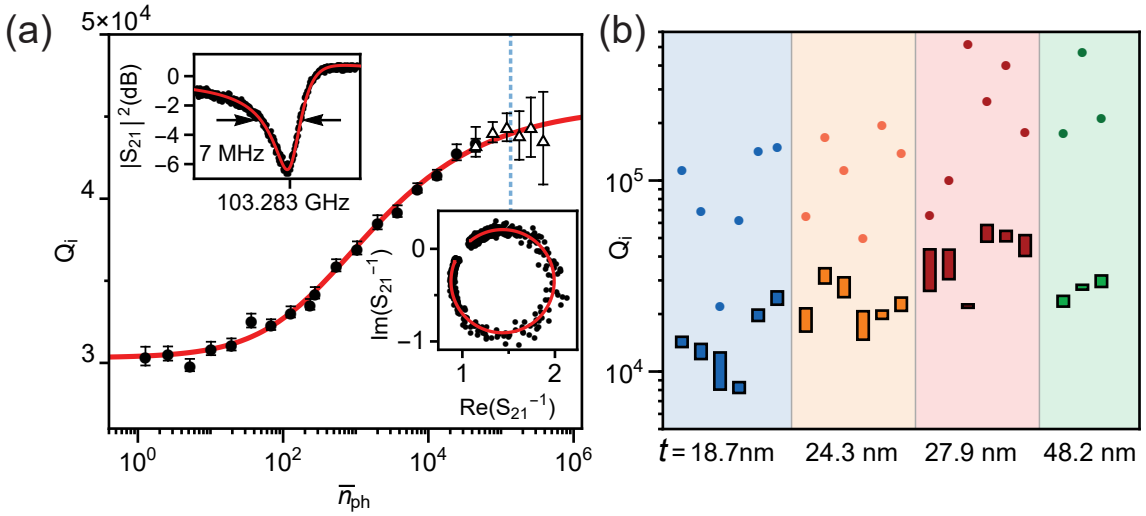

Figure 5.6: (a) Power dependence of the internal quality factor for a resonator with $Q_e^* = 2.759 \times 10^4$ patterned on a 27.9nm thick NbN film, measured at 1 K. White triangles are fits to a nonlinear response model near and above the bifurcation power (dashed line). The red line is a fit to a model including power-dependent loss from two-level systems, and power-independent loss. Insets show lineshape and fits at average photon occupation $\bar{n}_{\mathrm{ph}} \approx 1.2$. (b) Internal quality factors for resonators in this chapter, grouped by film thickness. The top and bottom of the colored bars correspond to fitted low and high power saturation values, while points correspond to two-level system induced Q_i with high power loss subtracted.

By characterizing complex transmission spectra of resonators fabricated on a range of film thicknesses, we explore loss mechanisms at millimeter-wave frequencies. The sheet inductance, thickness, and T_c measured for a given film are used to adjust the resonator design length. This spreads resonances out in frequency from 80 to 110 GHz, while varying antenna lengths allows us to adjust coupling strengths. A typical normalized transmission spectrum taken at single photon occupation powers ($\bar{n}_{\mathrm{ph}} \approx 1.2$) is shown in the inset of Figure 5.6a. On resonance, we observe a dip in magnitude, which at low powers is described well by [106]:

$$S_{21} = 1 - \frac{Q}{Q_e^*} \frac{e^{i\phi}}{1 + 2iQ\frac{\omega - \omega_0}{\omega_0}} \tag{5.1}$$

where $Q^{-1} = Q_i^{-1} + \mathrm{Re}[Q_e^{-1}]$ [106] and the coupling quality factor $Q_e = Q_e^* e^{-i\phi}$ has under-

gone a complex rotation ϕ due to an impedance mismatch [106, 153], likely induced by the sapphire chip and slot altering waveguide geometry. Plotting fitted internal quality factors with respect to photon occupation in Figure 5.6a shows that Q_i increases with power. This can be described by a power dependent saturation mechanism [195, 242], likely originating from two-level systems (TLS) in the slow-growing amorphous surface oxide [152]. At high powers, Q_i approaches a limit, indicating that another power-independent loss mechanism such as dielectric or radiative loss is dominant in this regime. For some samples, this limit is obscured by the early onset of nonlinear effects (discussed below).

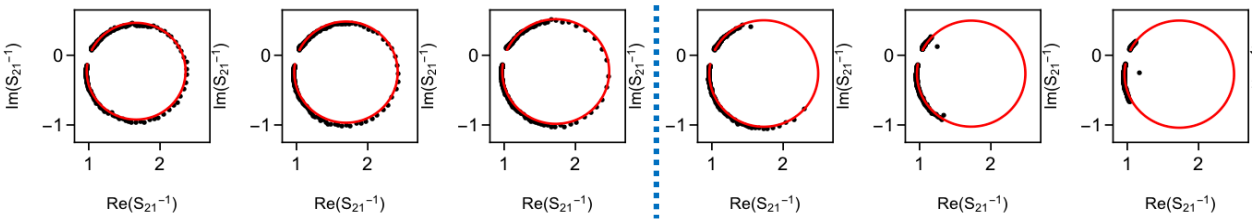

Figure 5.7: Complex transmission along with fits for a nonlinear resonator at powers near and above the bifurcation power (blue line) demonstrating how the quality factors can still be extracted from a bifurcated response, albeit with increasingly less certainty.

To study effects of film thickness on Q_i, we repeat measurements summarized in Figure 5.6a for devices varying in thickness from 19.5 nm to 48.8 nm and show the results in Figure 5.6b. We plot the low and high power limits of Q_i as well as a lower bound for the TLS induced Q_i for devices from six separate chips grouped by film thickness. For films thicker than 20 nm, we consistently find $Q_i > 10^4$, with TLS limiting Q_i around 10^5. We find a weak correlation of Q_i with film thickness, which could be explained by several additional potential sources of loss. Thinner films exhibiting higher disorder have also been connected with a nonlinear resistance [264], resulting in loss mechanisms proportional to kinetic inductance (See Appendix 5.4.1). Additionally, since devices are unshielded from ambient magnetic fields at the superconducting transition, vortices trapped in the thin films may cause additional dissipation [110, 160, 169, 174] dependent on film thickness. Resonances patterned from thinner films proved experimentally difficult to distinguish from background

fluctuations, possibly indicating low values of Q_i or frequencies outside the measurement bandwidth.

5.3.1 Thermal Losses from Complex Conductivity

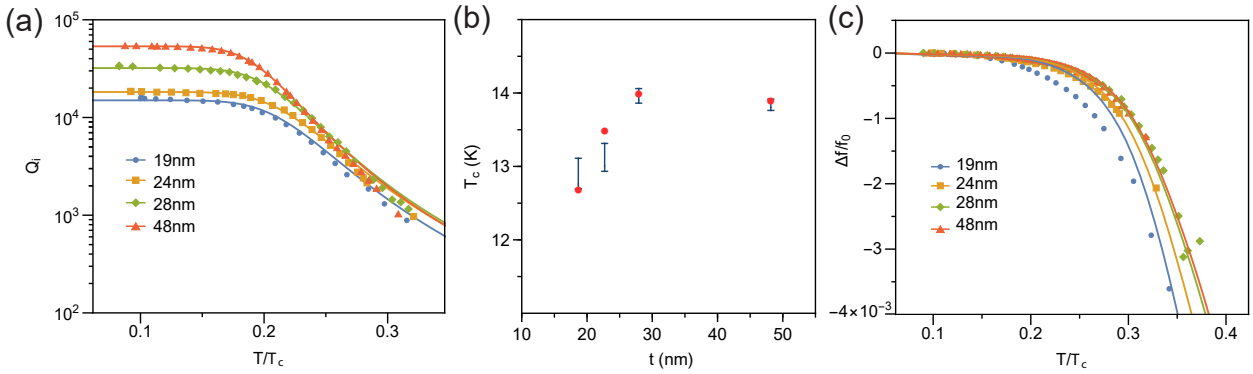

Figure 5.8: Temperature dependence of BCS conductivity. (a) High power Q_i as a function of normalized temperature for four resonators of different film thickness. Solid lines correspond to a BCS model with T_c and kinetic inductance fraction α as fit parameters. (b) Extracted T_c from fitting to BCS model (red dots), compared to T_c from DC resistivity measurements. (c) Normalized frequency shift of the same resonators as a function of temperature, with overlaid predictions from the Mattis-Bardeen equations for σ_2/σ_n with parameters taken from fits in (a).

Due to the large kinetic inductance fraction α, or magnetic field participation ratio of the thin film resonators, we expect higher sensitivity to loss from complex conductivity, which in turn is sensitive to temperature. In Figure 5.8a we show the quality factor decrease as a function of temperature for resonators with four different thicknesses, with solid lines corresponding to a model of the form

$$Q_i(T)^{-1} = Q_{i,\max}^{-1} + Q_\sigma(T)^{-1} \tag{5.2}$$

where $Q_{i,\max}$ is a temperature independent upper bound arising from other sources of loss,

and the conduction loss Q_σ is given by [189]:

$$Q_\sigma(T) = \frac{1}{\alpha}\frac{\sigma_2(T,T_c)}{\sigma_1(T,T_c)} \tag{5.3}$$

where σ_1 and σ_2 are the real and imaginary parts respectively of the complex surface impedance, calculated by numerically integrating the Mattis-Bardeen equations for σ_1/σ_n and σ_2/σ_n [145, 189, 224]. We use α and T_c as fit parameters in the model. Below 2 K ($T/T_c \sim 0.15$), Q_i saturates, which indicates that conduction loss does not limit Q_i for these devices. We note minor deviations from theory at higher temperatures, which may be a result of physical deviations from the standard curve calibrations used for the ruthenium oxide thermometer. Since these resonators were fabricated with $Q_e > 10^4$ measuring resonators at higher temperatures where Q_i is below 10^3 proved experimentally challenging. In Figure 5.8b we plot the fitted T_c values against those obtained with DC measurements and find reasonable agreement for higher thickness films, however note that the uncertainty in temperature calibration combined with the relatively low saturation values result in fitted T_c uncertainties around 0.4 K.

Bardeen-Cooper-Schrieffer theory also predicts a shift in London length as a function of temperature, which in the dirty (high disorder) limit is given by [87, 224]:

$$\frac{\lambda(T)}{\lambda(0)} = \frac{1}{\sqrt{\frac{\Delta(T)}{\Delta_0}\tanh\left(\frac{\Delta_0}{2k_bT}\right)}} \tag{5.4}$$

We can measure this by tracking the resonant frequency shift. For sufficiently large kinetic inductance fractions, or $L_k \gg L_g$, the kinetic inductance will dominate the total inductance, so the normalized frequency shift will be [87]

$$\frac{f_0(T)}{f_0(0)} = \sqrt{\frac{\Delta(T)}{\Delta_0}\tanh\left(\frac{\Delta_0}{2k_bT}\right)} \tag{5.5}$$

Figure 5.8c shows the normalized frequency shift as a function of normalized temperature and predictions from Eq. 5.5 with parameters α and T_c taken from fits to $Q_i(T)$ above. Notably, we find significant deviation from the BCS theory for lower thicknesses, which has been previously observed for high-disorder films [22, 87, 204].

5.4 Single-Tone Kerr Nonlinear Dynamics

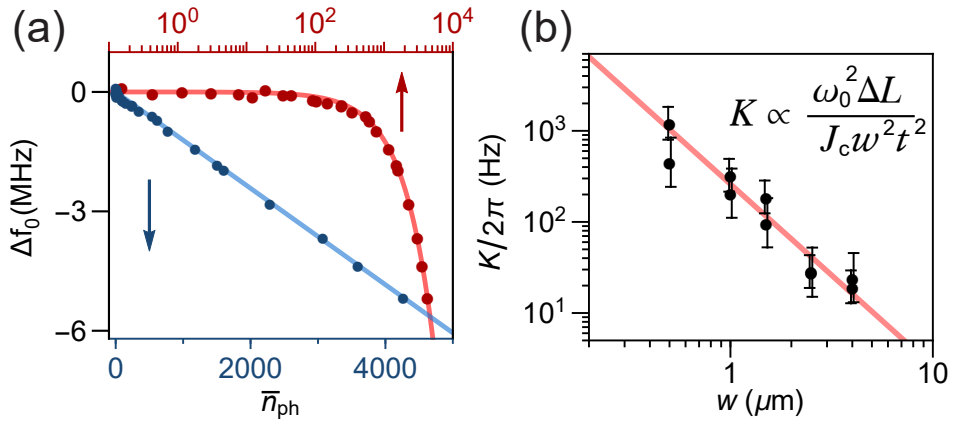

Figure 5.9: Measuring Kerr nonlinearity (a) Frequency shift versus average resonator photon number $\bar{n}_{\mathrm{ph}}$ in linear and log-scale ($\bar{n}_{\mathrm{ph}}$ accurate within a factor of ~ 10). (b) Extracted self-Kerr coefficients versus wire width w for resonators fabricated from a 29 nm thick film. Predicted w^{-2} dependence is shown in red. We find no significant impact of w on Q_i.

A key aspect of high KI resonators is their fourth-order nonlinearity: an important component for realizing quantum devices, and similar to the nonlinearity term found in Josephson junctions for low powers. Nonlinear kinetic inductance takes the general form $L = L_k + \Delta L_k I^2/I_c^2$, where L_k is the linear kinetic inductance, ΔL_k the nonlinear kinetic inductance, and I_c the critical current which sets the nonlinearity scale [108, 264]. This adds nonlinear terms of the form $\frac{\hbar}{2}K(a^\dagger a)^2$ to the Hamiltonian, with $K \propto \omega_0^2 \Delta L_k/I_c^2$, shifting the fundamental frequency ω_0 by the self-Kerr constant K for each photon added. To characterize the strength of the resonator nonlinearity, we measure the resonance frequency shift [5] as a function of photon number [138]. A linear fit for a resonator ($t = 29$ nm, $w =$

5. The point diametrically opposite $S_{21} = 1$

0.5 μm) yielding $K/2\pi = 1.21$ kHz is shown in Figure 5.9a. Although the statistical errors of the fits are small, we note this value has a multiplicative uncertainty of $10^{\pm 0.4}$: despite careful calibration, systematic variations in received power across the chip and between separate experiments limit best estimates of photon number to within a factor of $\sim$10. By writing the self-Kerr coefficient in terms of a current density $I_c = J_c wt$, we expect K to scale as w^{-2}, so in Figure 5.9b we plot the self-Kerr coefficients of devices (and error relative to each other) with respect to their wire width. These results are comparable to self-Kerr strengths of granular aluminum nanowires [138] or weakly nonlinear Josephson junctions [61].

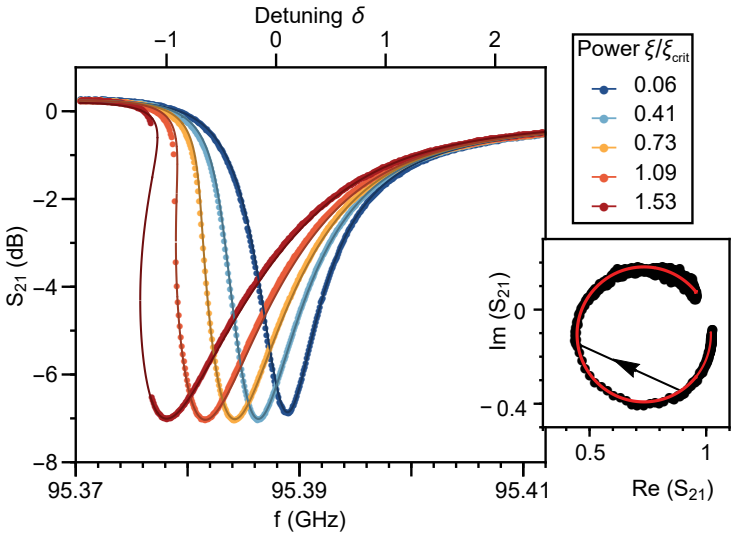

Figure 5.10: Transmission of a typical resonance at a range of powers near and above bifurcation showing good agreement with a Kerr nonlinear response. The inset shows overlaid data and fits in the complex plane just below and above the bifurcation point.

A hallmark of Kerr nonlinearity is the distortion of the transmission line-shape in frequency space at high powers, ultimately leading to a multi-valued response above the bifurcation power. Re-writing $\gamma = \omega_0/Q_i$ and $\kappa = \omega_0/\mathrm{Re}[Q_e]$, the steady-state nonlinear response takes the form derived from Refs. [219, 264] (See Chapter 4):

$$S_{21} = 1 - \frac{\kappa}{\kappa + \gamma} \frac{1 + i \tan\phi}{1 + 2i(\delta - \xi n)} \tag{5.6}$$

where the frequency detuning is written in reduced form $\delta = \frac{\omega-\omega_0}{\kappa+\gamma}$, and $n = n_{\rm ph}/\tilde{n}_{\rm in}$ is a function of frequency and reduced circulating power $\xi = K\frac{\kappa}{hf(\kappa+\gamma)^3}P_{\rm in}$. We plot steady-state transmission data taken near the bifurcation power in Figure 5.10 along with fits to Eq. 5.6, with system parameters κ, ϕ and ω_0 constrained to low power values and find the model in good agreement with measurements.

5.4.1 *Controlling nonlinearity in the presence of additional losses*

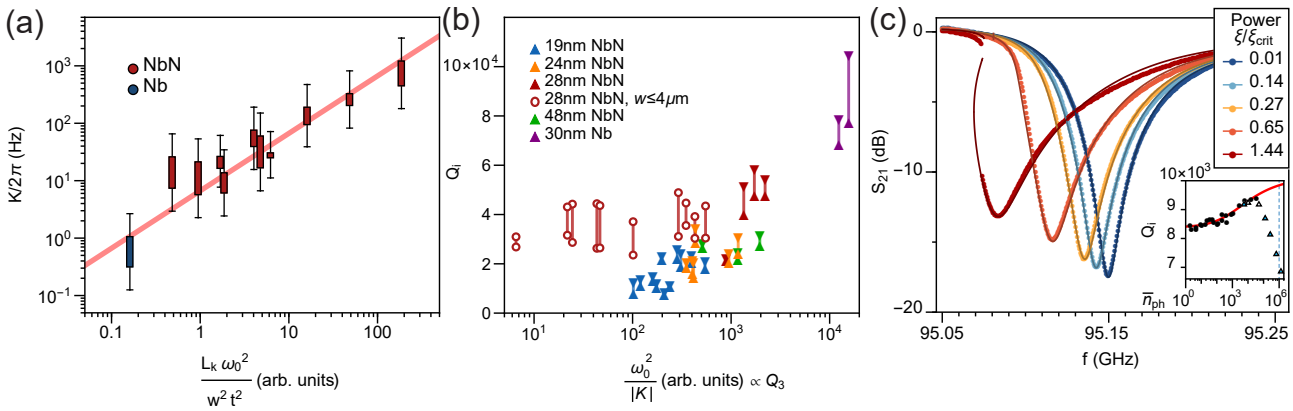

Figure 5.11: (a) Self-Kerr constant $|K_{00}|$ as a function of parameters in Eq. 5.7, with a linear fit overlaid as a solid line. Solid bars correspond to value ranges for groups of similar film thicknesses and wire widths, with error bars marking systematic uncertainty. (b) Q_i as a function of $\omega_0^2/|K_{00}|$ which corresponds to the loss Q_3 associated with kinetic inductance. Points correspond to low and high power limits of Q_i. Note that devices with varying wire width (empty circles) do not appear to be correlated with Q_3. (c) Transmission as a function of frequency for a 18.7 nm thick device at 95.15 GHz taken at increasing powers ξ, with the inset highlighting decreasing Q_i near the bifurcation power $\xi_{\rm crit}$ (dashed blue line) deviating from two level system loss model (red line). Triangles correspond to nonlinear model fits, with traces shown in main panel marked in blue.

From Ref. [264], we expect the self-Kerr nonlinearity originating from kinetic inductance of a $\lambda/4$ resonator to be

$$K = -\frac{\hbar\omega_0^2}{I_c^2}\int_0^l dx u_0^4 \Delta L \propto -\frac{\hbar\omega_0^2 L_k}{J_c^2 w^2 t^2} \tag{5.7}$$

where in our case the nonlinear kinetic inductance ΔL is constant along the transmission

line, so integrating over the fundamental mode profile u_0 yields a constant. We have also transformed the critical current I_c into a critical current density J_c, and used the assumption that the nonlinear kinetic inductance is proportional to the linear kinetic inductance [108, 219]. Figure 5.11a expands on Figure 5.9b, showing measured self-Kerr constants for all resonators in this chapter (grouped into points by film thickness and wire width) as a function the parameters in Eq. 5.7, with the solid line corresponding to a linear fit. We have also added data from identical resonators fabricated from 30nm electron-beam evaporated niobium to extend the parameter range. We find reasonable agreement with dependence on the parameters in Eq. 5.7, however note that the dependence is much less clear than that on wire width w.

Nonlinear kinetic inductance is also associated with a nonlinear resistance of the same form $R = R_0 + \Delta R I^2/I_c^2$. Based on Ref. [264], and assuming the nonlinear resistance scales with kinetic inductance, losses associated with nonlinear resistance will be

$$\gamma_3 = \frac{\omega_0}{Q_3} = \frac{3\hbar\omega_0}{8I_c^2}\int_0^l dx u_0^4 \Delta R \propto \frac{\hbar\omega_0 L_k}{J_c^2 w^2 t^2} \sim \frac{-K}{\omega_0} \tag{5.8}$$

This indicates that upper bounds on nonlinear losses should scale as $Q_3 \sim \frac{\omega_0^2}{|K|}$. In Figure 5.11b we plot low and high power limits of Q_i devices in this chapter with the addition of 30 nm Niobium devices described above, and find that for resonators with fixed wire widths $w = 4$ μm, there appears to be a potential correlation of Q_i with Q_3 indicating nonlinear resistance may be a potential loss mechanism.

In our analysis, we have also neglected to take into account higher harmonics of the resonator, which will be coupled to the fundamental mode by cross-Kerr interactions χ_{mn}, which for evenly spaced harmonics scale as [264]

$$\chi_{mn} = -\frac{3\hbar\omega_m\omega_n}{I_c^2}\int_0^l dx u_m^2 u_n^2 \Delta L \propto K \tag{5.9}$$

Given the proportionality to K, the correlation described above may also potentially be a result of cross-Kerr effects. For line-widths large enough to cover any deviations from evenly spaced higher harmonics, we anticipate see power dependent conversion processes: in particular for a Kerr nonlinear system with harmonics spaced at ω_0 and $3\omega_0$, at powers approaching the critical power we would expect increased conversion efficiency from the fundamental to third harmonic [86], which in our experiment would be observed as increased resonator loss at higher powers.

In Figure 5.11c we show the atypical transmission spectra of a 18.7 nm thick, 4 μm wide device with a particularly large line-width showing decreasing Q_i near the bifurcation power (above $n_{\mathrm{ph}} \sim 10^5$), departing from the two-level system model described in the main text. This additional power-dependent loss may be the result of the nonlinear mechanisms described above, but may also be a result of circulating currents exceeding the thin film critical current density, which is lowered by the increased London lengths of the thinner films [220, 229]. However since the loss could also simply be a result of frequency dependent dissipation, the source remains unclear.

In Figure 5.11b, we also observe that resonators achieving higher nonlinearities by reducing wire width do not appear to be affected by the nonlinear loss rate described above. We also find that these devices do not showcase obvious signs of high power loss shown in Figure 5.11c. While this may be a result of the difference in fabrication methods (see Appendix B), the thinner wires may have higher vortex critical fields [211] leading to reduced vortex formation, and thus lower loss associated with vortex dissipation. Additionally, the thinner wires at the shorted end of the quarter wave section of the resonator further shift the higher harmonics, potentially resulting in lower cross-Kerr conversion loss.

5.5 Degenerate Four-wave mixing

We further explore nonlinear dynamics by stimulating degenerate four-wave mixing with the addition of a continuous wave classical pump (see Fig. 5.4). When a high power pump tone at ω_p is on resonance with the down-shifted resonance frequency, and a low power signal at ω_s is at a frequency detuning Δ from the pump, we expect to observe the net conversion of two pump photons into a signal photon and an idler photon at their sum-average frequency [Figure 5.12a inset]. This effect is most pronounced when all frequencies are within the resonant bandwidth, and the pump power ξ approaches the bifurcation point ξ_{crit}, but is limited by the loss fraction γ/κ.

Inspired by Refs. [70, 146] we can decompose a side-coupled resonator into a linear network containing a one-sided cavity, which is very well understood in the language of input-output theory used in quantum optics [239]. This allows us to map well-modelled dynamics of a Kerr nonlinear cavity driven in reflection [61, 264] to a side-coupled resonator measured in transmission, obtaining results in agreement with Ref. [219], which uses a more direct approach. We follow a similar approach to obtain expressions for parametric conversion gain using the derived input-output relations to map the key results from Ref. [61] to the waveguide inputs and outputs. Using microwave conventions for Fourier transforms, the one-port gain of a signal detuned from the pump by $+\Delta = \frac{\omega_s - \omega_p}{\kappa+\gamma}$ is given by:

$$g_s = \frac{a_\Delta^{\text{out}}}{a_\Delta^{\text{in}}} = 1 - \frac{\kappa}{\kappa+\gamma}\frac{\frac{1}{2} - i(\delta - 2\xi n - \Delta)}{(i\Delta + \lambda_+)(i\Delta + \lambda_-)} \tag{5.10}$$

With $\lambda_\pm = \frac{1}{2} \pm \sqrt{(\xi n)^2 - (\delta - 2\xi n)^2}$. Using the 3 port network transformations above yields the normalized forward (in direction of propagation) signal gain:

$$g_s^+ = \frac{a_{2,\Delta}^{\text{out}}}{a_{1,\Delta}^{\text{in}}} = 1 - \frac{\kappa}{\kappa+\gamma}\frac{e^{i\phi}}{\cos\phi}\frac{\frac{1}{2} - i(\delta - 2\xi n - \Delta)}{2(i\Delta + \lambda_+)(i\Delta + \lambda_-)} \tag{5.11}$$

We measure the pump-signal conversion efficiency of a high-bandwidth, high-K device in

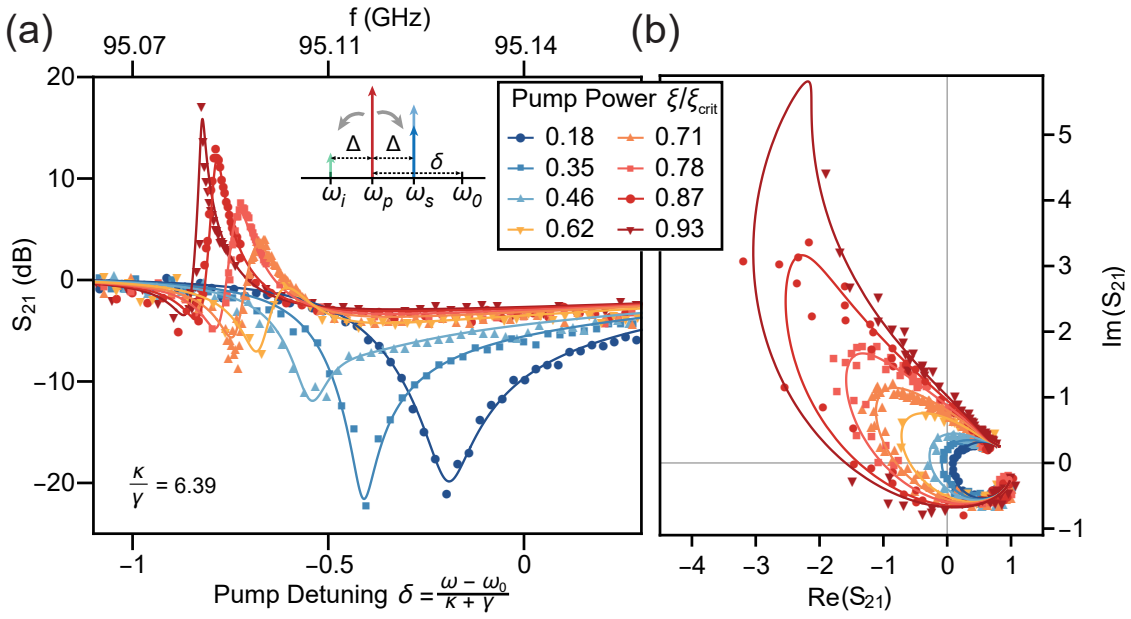

Figure 5.12: Four-wave mixing. (b-c) Parametric conversion gain with a 95.1 GHz device with the same film thickness as 5.10 as a function of reduced pump frequency δ for a fixed signal detuning Δ of +450 kHz, taken at increasing pump powers. Solid lines correspond to theoretical response. The initial forward de-amplification is better understood when the response is viewed in the complex plane (c), where we observe smooth parametric deformation from the single tone response.

the propagation direction as a function of reduced pump frequency δ for increasing pump powers ξ, and a fixed signal power corresponding to $\bar{n}_{\mathrm{ph}} \simeq 9$ in Fig. 5.12. We find that this behavior is accurately captured with the model described above, and overlay the results. For increasing pump powers, we observe smooth parametric deformation from the single tone response in the complex plane. For higher powers, we observe increasing gain with decreasing linewidth similar to Refs. [61, 222], up to a maximum measured forward efficiency of +16 dB. The slight curvature in the complex plane is a result of the finite pump-signal detuning Δ.

5.6 Conclusion

The demonstration of degenerate four-wave mixing realizes an important milestone for the development of quantum devices at millimeter-wave frequencies and temperatures above 1 K. For NbN films thicker than 25 nm, we measured millimeter-wave resonators with internal quality factors exceeding 2×10^4 at single photon powers, and by reducing wire width to 500 nm achieved self-Kerr nonlinearities up to 1.21 kHz for linewidths ranging from 1-200 MHz. With some modification the devices in this chapter could easily be redesigned as parametric amplifiers, which at microwave frequencies have been shown to achieve near quantum-limited noise figures and quadrature squeezing [64, 164, 222, 222, 263]. While insufficient for implementing a millimeter-wave artificial atom, the Kerr nonlinearity we measure arising from high KI thin films can further be used for superinductors [23, 169], photon frequency conversion [181], parametric mode cooling [107, 267], phase slip junctions [17, 161], and mode squeezing [222] realized at millimeter-wave frequencies. This opens the door to a new generation of high-frequency quantum experiments at temperatures above 1 K.

5.6.1 Nonlinearity Relative to Linewidth

Since our goal is to realize a strong single-photon-resolved nonlinearity, a useful metric to examine is the nonlinearity strength K relative to the total device linewidth $\gamma_T = \kappa + \gamma$. In Figure 5.13 we plot the nonlinearity with respect to linewidth for all the millimeter-wave NbN resonators described above, grouped by wire width, which we showed is inversely proportional to the nonlinearity strength. For these devices we find that we achieve a maximum nonlinearity of nearly 0.1 percent of the device linewidth, suggesting that a stronger nonlinearity is desirable...

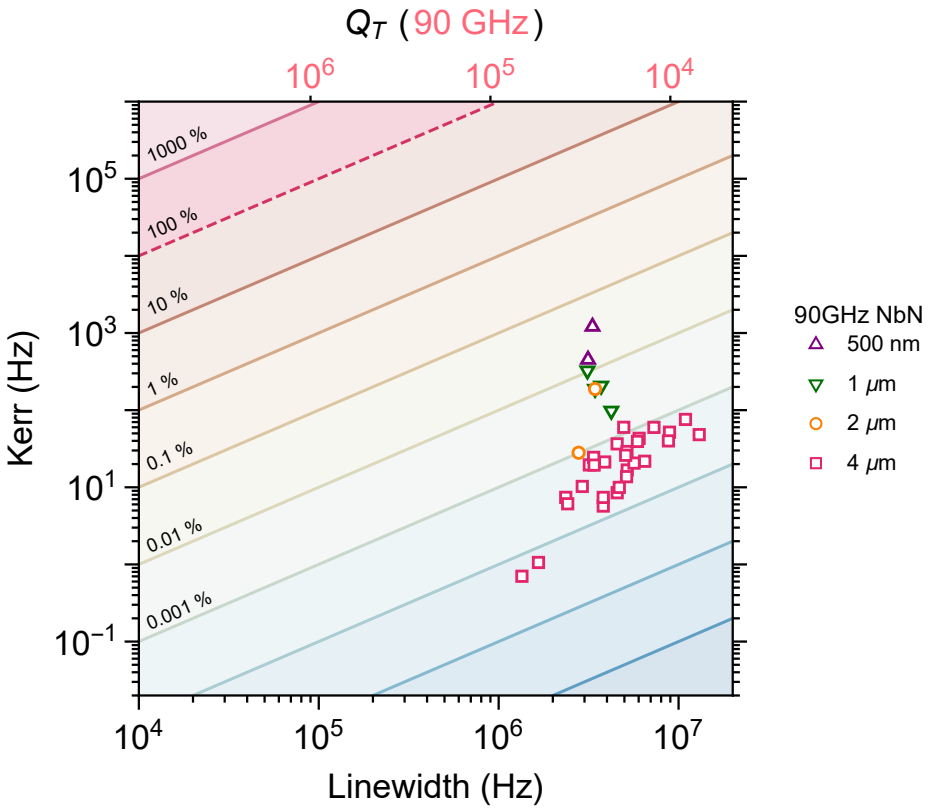

Figure 5.13: Nonlinearity (K) plotted as a function of total linewidth (γ_T) for all of the millimeter-wave CPS resonators made with NbN, grouped by wire width.

CHAPTER 6

IMPEDANCE-ENHANCED KERR NONLINEARITY

For an anharmonic superconducting system useful for quantum experiments, we would like a single-photon resolved nonlinearity. For microwave devices, this is frequently achieved with aluminum Josephson junction [111]. However at millimeter-wave frequencies, using an aluminum Josephson junction isn't practical, since millimeter-wave photons (which have energy above $2\Delta_{\mathrm{Al}}$) will directly generate quasiparticles, suppressing superconductivity and increasing sources of dissipation [47, 144]. Additionally, since we would ultimately like to operate near 1 K in an environment cooled with Helium-4, we would like a superconductor with a higher transition temperature (T_c) to avoid thermal decoherence from excess quasiparticles [145].

Kinetic inductance on the other hand offers an alternate source of the desired Kerr nonlinearity, arising from a current-dependent inductance in thin superconducting wires. In the previous chapter, we demonstrated that the kinetic inductance in thin Niobium Nitride films is a useful source of nonlinearity at millimeter-wave frequencies, and could show parametric oscillation and amplification. However a much more useful goal would be to realize a system where the self-Kerr is comparable to its linewidth[1]. Thin film kinetic inductance has also been proposed as a method to provide sufficient nonlinearity for realizing a millimeter-wave qubit [65]. Recently this concept was demonstrated at microwave frequencies and dilution temperatures (<50 mK) using granular aluminum [199, 252], which has a strong inherent nonlinearity arising from weak coupling between aluminum grains [137]. Similar to aluminum[2] however granular aluminum has a relatively low critical temperature, making it unsuitable for operating at elevated temperatures.

1. Here I will disagree with Ref.[99] and say that the total linewidth is the one that the nonlinearity should be compared to, not just the intrinsic linewidth (loss).

2. Actually granular aluminum has a slightly increased critical temperature (closer to 2 K) arising from disorder, which is somewhat counter intuitive.

Thin films of niobium nitride (NbN), and titanium nitride (TiN) have been shown to have very low microwave loss [82, 169, 204] while sustaining high levels of kinetic inductance. Kinetic inductance has also successfully been used as a three wave mixing nonlinearity: useful for decoupled parametric amplification [140] and frequency tunable "Ouroboros" resonators[3] [255].

However the nonlinearity strengths found in thin disordered films are still relatively weak when compared to the device linewidths, and typically lower than those seen in granular aluminum. Luckily it has recently been shown that even a weak source of Kerr nonlinearity can still be used to isolate the first two states of the weakly nonlinear system [92, 128], forming the desired artificial atom. As a result, we will attempt to maximize the inherent self-nonlinearity of a resonator with the goal of making a highly nonlinear resonator with an anharmonicity *nearly* that of the linewidth, which could then be used to create an adressable quantum system using a photon blockade such as proposed in Ref. [128].

In practice while our devices did achieve $K/\kappa = 0.17$, we found that

6.1 Increasing Nonlinearity from Kinetic Inductance

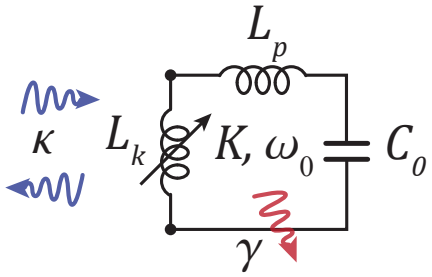

Figure 6.1: Nonlinear resonator consisting of a nonlinear inductor shunted by a single capacitance C_0. The parasitic inductance of the capacitor in this case reduces the self-Kerr strength so cannot be neglected.

Earlier on, we established that kinetic inductance can be expressed as a nonlinear induc-

3. I actually think they unfortunately got bullied out of using "Ouroboros" in their publication...

tance [70, 269]:

$$L(I) = L_k \left(1 + \frac{I^2}{I_*^2}\right) \tag{6.1}$$

Consider a simple circuit formed by shunting this nonlinear (kinetic) inductor L_k with a single capacitor C_0 (this can be the sum of several capacitances, so long as each has its voltage nodes shunting the nolinear inductor). Of course any real capacitor will have some associated parasitic inductance $L_p \ll L_k$ that contributes to the total inductance $L = L_k + L_p$. The power dependent resonant frequency of the circuit can be approximated for small currents (since we care about the zero and single photon separation)

$$\omega(n) = \frac{1}{\sqrt{C_0(L_k + L_p + L_k I^2/I_*^2)}} \approx \omega_0 \left(1 - \frac{1}{2}\frac{L_k}{L_k + L_p}\frac{I^2}{I_*^2}\right) \tag{6.2}$$

Labelling the kinetic inductance fraction $\alpha = L_k/L$ we can extract the frequency-shift per photon in terms of the zero point fluctuation current $I_{\text{ZPF}}^2 = \hbar\omega_0/L$. This means that the self-Kerr nonlinearity of the circuit will be:

$$K = -\omega_0 \frac{\alpha}{2}\frac{I^2}{I_*^2}\frac{1}{n\hbar} = -\alpha\frac{\hbar\omega_0^2}{LI_*^2} \tag{6.3}$$

Now using the convenient fact that inductance can be written as $L = Z_0/\omega_0$ in terms of the circuit impedance $Z_0 = \sqrt{L/C_0}$, the Kerr nonlinearity can be simplified to

$$K = -\frac{\alpha\hbar\omega_0^3}{Z_0 I_*^2} \tag{6.4}$$

From Chapter 1 we found the critical current I_* can be expressed in terms of a material-dependent critical current density along with the wire cross section area A

$$I_* = J_* A \approx A\sqrt{\frac{\pi}{\hbar}\sigma_n N(0)\Delta(0)^3} \tag{6.5}$$

Putting all of this together:

$$K = -\frac{\alpha}{Z_0 A^2}\frac{\hbar\omega_0^3}{J_*^2} \tag{6.6}$$

So given a nonlinear resonator at frequency ω_0, we can maximize the nonlinearity by maximizing the kinetic inductance fraction α, decreasing the wire cross section area A. However from Equation 6.6 we see that the circuit nonlinearity can also be increased by reducing the circuit impedance Z_0.

Counterintuitively, this means for a highly nonlinear circuit we want a very small inductance and very large capacitance! Since we need to keep the kinetic inductance fraction α high however, we need to ensure the capacitor has minimal parasitic inductance L_p.

6.2 Parasitic Inductance in a Finger Capacitor

To simplify the fabrication process we would like to make a microwave structure with high capacitance using the same on-chip thin films that will form the high-kinetic-inductance nanowire. We could make this by placing two metal pads next to each other (as shown on the left of 6.2). However to achieve a practical capacitance[4] above a few hundred femtoFarads, this would require gigantic structures on the order of millimeters, which isn't very tractable for compact microwave circuits[5]. Since the capacitance of the two pads scales effectively scales with the length of their adjacent surfaces [78], a practical solution is to meander this surface: folding the overlapping region back and forth on itself quickly increases the overlap length while taking up the same area. This process is illustrated in Figure 6.2 and results a structure referred to as a finger capacitor (see Ref. [19] for a very nice model for estimating its capacitance).

Finger capacitors are quite popular in microwave design, particularly in superconducting designs [19, 77, 113], since they can be patterned in a single fabrication step, and don't

4. For millimeter-waves, ~ 10 fF is practical already so we don't usually need to worry about this!

5. Or maybe not? In some cases[99] they didn't mind and made gigantic capacitors anyways.

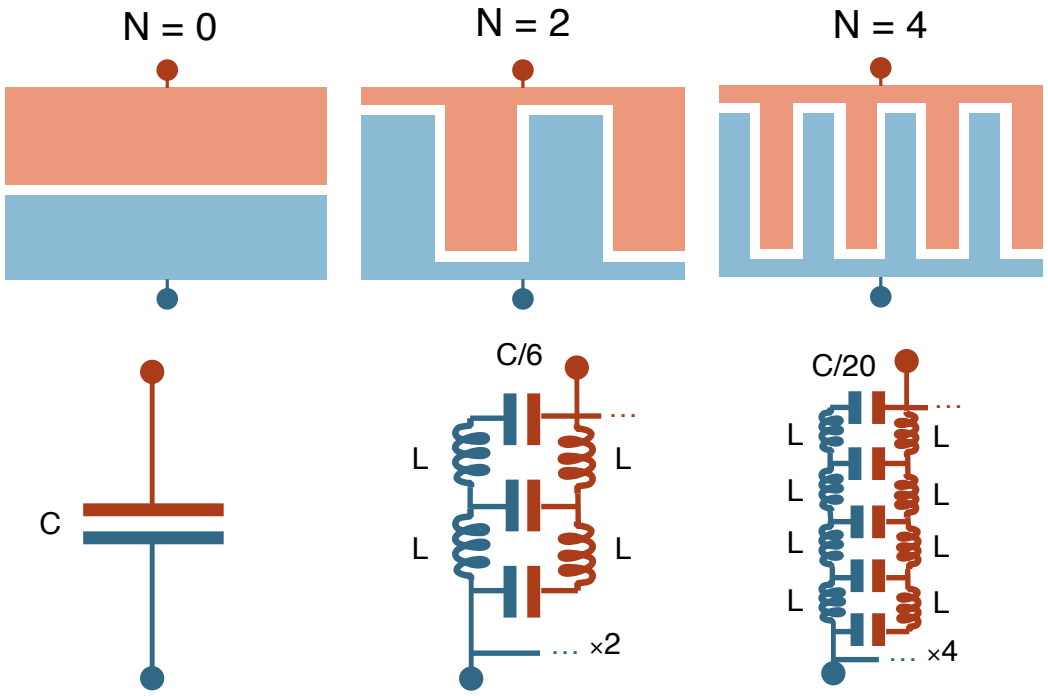

Figure 6.2: Illustration of a finger capacitor with increasing number of fingers N, along with the equivalent parasitic circuit network (the circuit is symmetric, so is shown for a single finger pair).

involve additional lossy dielectrics. However consider a capacitor that needs to be mode-free at millimeter-wave frequencies (such as for the circuit outlined in Ref. [181]), or a capacitor made from a high-kinetic inductance material [99]: in both of these cases the parasitic contributions to the finger capacitor will be significant and shouldn't be ignored. Following the same procedure outlined above for the simple example capacitor, we can break down the finger capacitor into subdivided capacitance and inductance networks, as outlined in Figure 6.2. As we increase the number of fingers and the amount of detail (or subdivisions) the effective parasitic circuit quickly becomes a messy grid of reactive components, as shown in Figure 6.2 (where we actually only show the circuit for a single finger for brevity).

Let's take a more reasonable approach to solving this problem, and simplify it with the assumption that the two ends of the finger capacitor (A and B) are electrically large, so we can neglect the effects of lateral parasitics. We are working with metal films that have

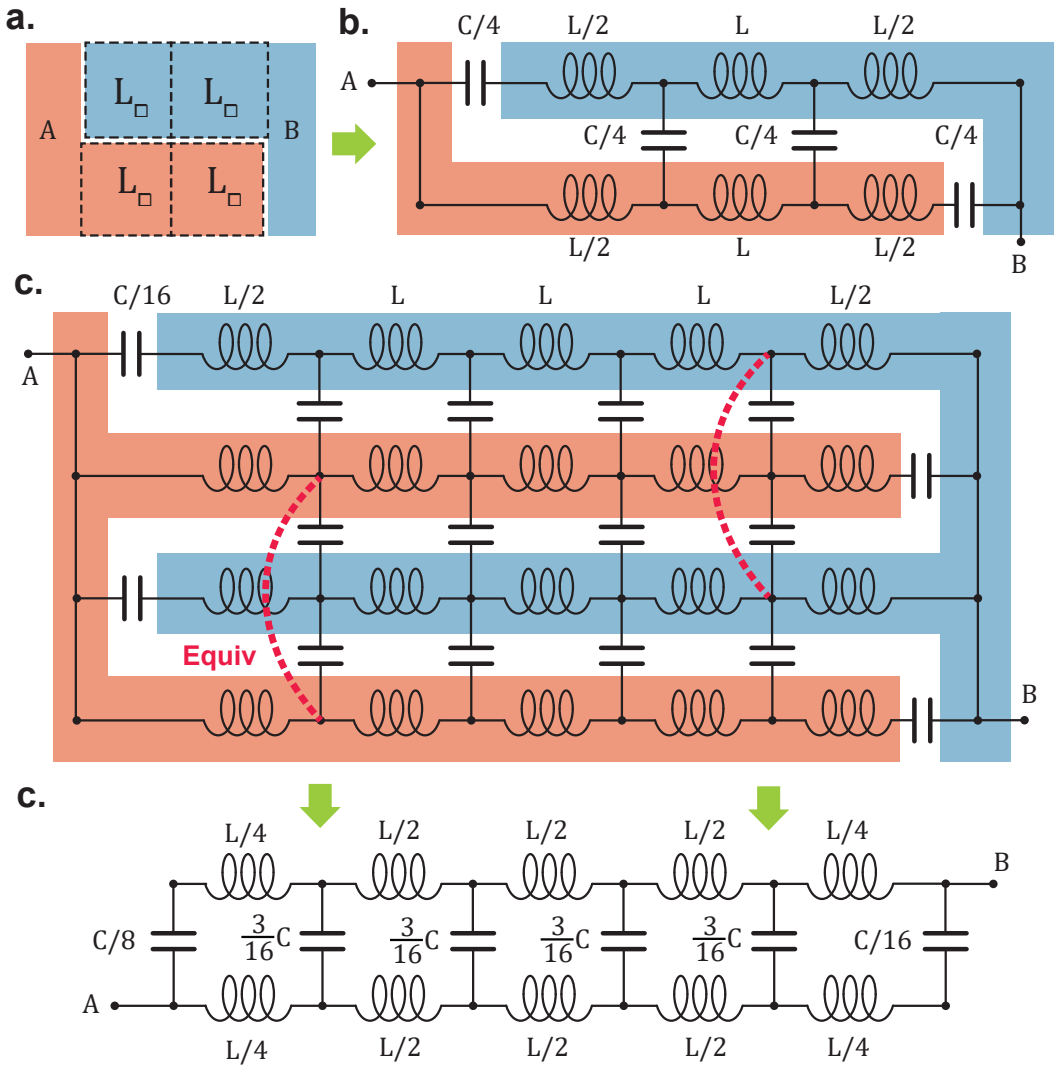

Figure 6.3: Subdivision of a finger capacitor. (a) We begin with a simple two-finger capacitor and split up the fingers into squares, each with sheet inductance $L_\square$ and a fraction of capacitance on their edges. This circuit is drawn out explicitly in (b), showing the parasitic network. The capacitances all still add up to the original capacitance C. (c) The parasitic network is drawn out for a four-finger capacitor. Assuming both ends of the capacitor A and B are well connected, we can use a symmetry argument to locate equivalent points on the finger capacitor and fold it up into the simplified circuit shown in (c). This process can be repeated for any number of fingers, which will only change the coefficients, but the equivalent circuit will look stay the same. This circuit only changes depending on how many squares we divide each finger into.

sheet inductance (eg. kinetic inductance), which is commonly defined as an inductance per

square[6]. So we will break the fingers down into squares, as shown in Figure 6.3a. Each unit cell will contribute a sheet inductance $L_\square$ (or L for short), and each interface between the unit cells will contribute some fraction of the total capacitance C/N, where the sum of all the sub-capacitors should sum up to the total capacitance C. The resulting circuit for a finger capacitor consisting of two fingers is shown in Figure 6.3b. From a tiling argument, we have put the capacitors in the middle of the unit cell, and split the unit cell inductance into two parts ($L/2$) on either side of a node.

A two finger capacitor is rather small though- higher capacitance density can be easily achieved by increasing the number of fingers [19]. In Figure 6.3c, we double the number of fingers, and simultaneously increase the number of subdivisions for better detail[7]. This produces the messy grid of inductors and capacitors which we touched on in Figure 6.2. However if we ignore lateral fluctuations (having assumed well-connected leads on either side of the finger capacitor, the voltage on either ends will at least start laterally homogenous), we can assume the voltage along each pair of fingers will be identical, allowing us to fold the fingers into a single pair. The resulting circuit is shown in Figure 6.3d, and is much more similar to the situation in Figure 6.3a-b, except with different weights on the sub-capacitors. The only fundamentally different thing about the new circuit comes from the number of squares we included. Going a step further we can use the same lateral symmetry argument the circuit of a n-fingered capacitor with m subdivisions into an equivalent two-fingers circuit with modified weights:

$$C_{\text{outer}} = C\frac{n/2}{m(n-1)+n} \qquad C_{\text{inner}} = C\frac{(n-1)}{m(n-1)+n} \tag{6.7}$$

$$L_{\text{outer}} = \frac{L}{n} \qquad L_{\text{inner}} = \frac{L}{2n} \tag{6.8}$$

6. We defined this explicitly for kinetic inductance in Chapter 1

7. This is actually a relic from where we were designing a capacitor that would have a fixed aspect ratio (I don't remember why)

With the finger-number scaling established, we can solve for the effective impedance of the two-finger equivalent circuit. Consider the circuit with four subdivisions, as shown in Figure 6.4. Unfortunately since this circuit has loops, we can't use simple algebra for determining the impedance. Instead, we consider a current I_0 flowing from $A \to B$, and use Kirchoff's law to establish a network of dependent equations for the currents on each branch I_1, I_3 and I_4 with respect to the impedances of the circuit (Z_C and Z_L are impedances of the center capacitor and inductor pieces, while $Z_R = i\omega L_{\text{outer}} + 1/i\omega C_{\text{outer}}$). We won't write out the equations, but there are enough of them that we can find a solution for V_A. Setting $V_B = 0$ for simplicity, the effective impedance of the finger capacitor can then be calculated from Ohm's law $Z_{\text{eff}} = V_A/I_0$. Wow you're actually reading this? You deserve a beer. With that, we can skip the math (it gets gnarly) we can summarize the results for $m = 2$–4:

$$Z_{\text{Cap}}^{m=2} = \frac{Z_c(3Z_L + Z_R) + 2Z_L(Z_L + 3Z_R)}{2(Z_C + Z_L + Z_R)} \tag{6.9}$$

$$Z_{\text{Cap}}^{m=3} = \frac{3Z_C^4 Z_L + Z_C^4(Z_C + Z_L)}{3Z_C^3(Z_C + Z_L) + Z_C^3(2Z_C + Z_L)} \tag{6.10}$$

$$Z_{\text{Cap}}^{m=4} = \frac{Z_C^2(7Z_L + Z_R) + 8Z_C Z_L(5Z_L + 3Z_R) + 8Z_L^2(3Z_L + 5Z_R)}{2(Z_C^2 + 4Z_L(Z_L + Z_R) + 2Z_C(3Z_L + Z_R))} \tag{6.11}$$

And so on and so forth (the higher order expressions started not fitting on the page). From

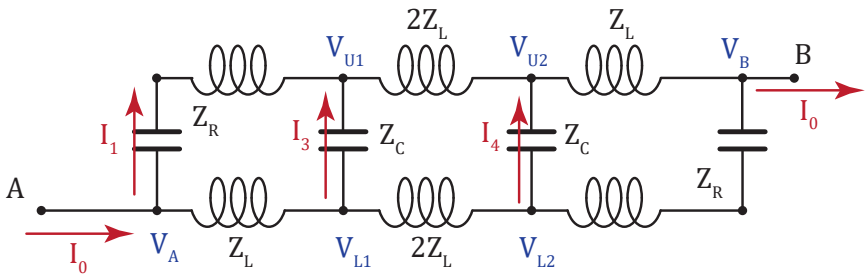

Figure 6.4: Simplified parasitic circuit for fingers divided into two squares. Since the circuit has loops we must use Kirchoff's theorem and solve for the effective impedance given an applied current I_0 and voltage $V_A - V_B$.

the expressions above, we can substitute the respective weights found earlier. As it turns out, for all of these, we can extract a single inductive correction term proportional to $L_\square$ (remember our simple model from the previous section?) and write the following:

$$Z_{\text{Cap}}(n,m) = \frac{1}{i\omega C_0} + i\omega \frac{m}{n} L_\square \left(1 + \mathcal{E}_{m,n}[C_0, L_0]\right) \tag{6.12}$$

Here, we have bunched the remaining correction terms into a function $\mathcal{E}_{m,n}$ which has a

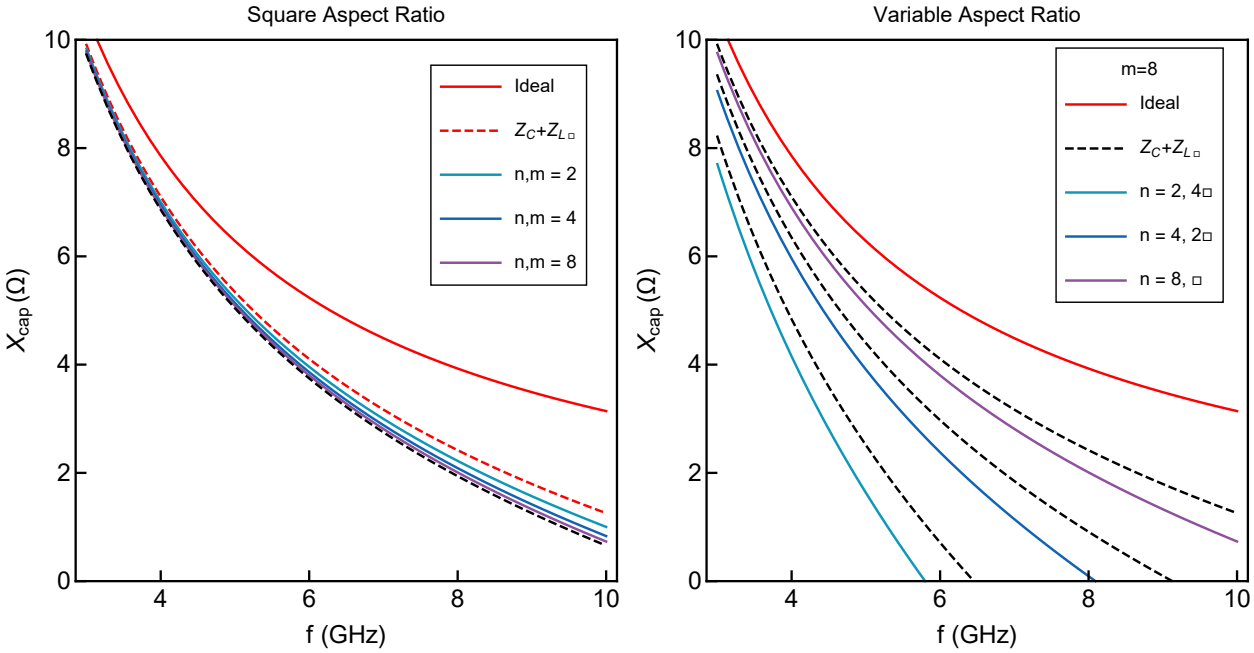

Figure 6.5: Effective reactance of a large (2 pF) finger capacitor, assuming a sheet inductance of 30 pH/□. For square aspect ratios ($n = m$ shown on left), the parasitics are largely described by adding a series inductor. Higher order subdivisions add small corrections to this, which are largely captured by $m = 8$ subdivisions. Since the added inductance depends on the aspect ratio of the capacitor multiplied by the sheet inductance, increasing finger number for a fixed finger length reduces the parasitic inductance as shown on the right.

small magnitude. From this expression we see that the capacitor impedance will be approximately modified by a series inductance, which can be calculated by multiplying the sheet inductance $L_\square$ by the aspect ratio of the capacitor m/n. In Figure 6.5, we plot the impedance of a finger capacitor based on the expressions we calculated (and higher subdivi-

sions) for varying subdivisions and finger numbers. We find that for a square aspect ratio, the parasitic corrections are largely summarized by a inductive term $L_\square$ in series with the capacitor impedance. Increasing the number of subdivisions reveals the small additional effects of the correction function $\mathcal{E}_{m,n}$, however we can get pretty good approximations with the $m = 4$ or $m = 8$ expressions. Using the fairly accurate impedance expression $Z_{m=8}$, we can explore what happens to the capacitor as we adjust the aspect ratio by keeping a fixed sheet inductance and varying the number of fingers. As we expect, longer and thinner narrower capacitors (such as the $n = 2$ finger case) behave more like wires and have more and more significant parasitic effects.

For the finger capacitor to successfully save space, we would like the aspect ratio to be as close to square as possible. From Figure 6.5 we see that unless we keep the finger capacitor aspect ratio low (which would mean it takes up a lot of room, just in the wide dimension), our capacitor will inherently have inductive effects associated with it. So the moral of the story[8] is be wary of finger capacitors made from metals with high sheet inductance! This decomposition method can help figure out how geometry will actually behave, but nowadays it's much nicer to let a computer simulation do the work.

6.3 High Capacitance Density with Fractals

Evidently finger capacitors, which are the traditional method used for increasing capacitance in a small footprint, have an inescapable amount of parasitic sheet inductance. The downfall of this design is in the inherent inductance of the fingers, throughout which the charge must flow. A simple solution which effectively reduces each finger length while also increasing their total number is achieved by instead sequentially branching each finger. This sequential branching forms a fractal pattern!

For our capacitor design, we use a modified square Sierpinski curve [50, 207]. The two-

8. The real parasitics were the friends we made along the way?

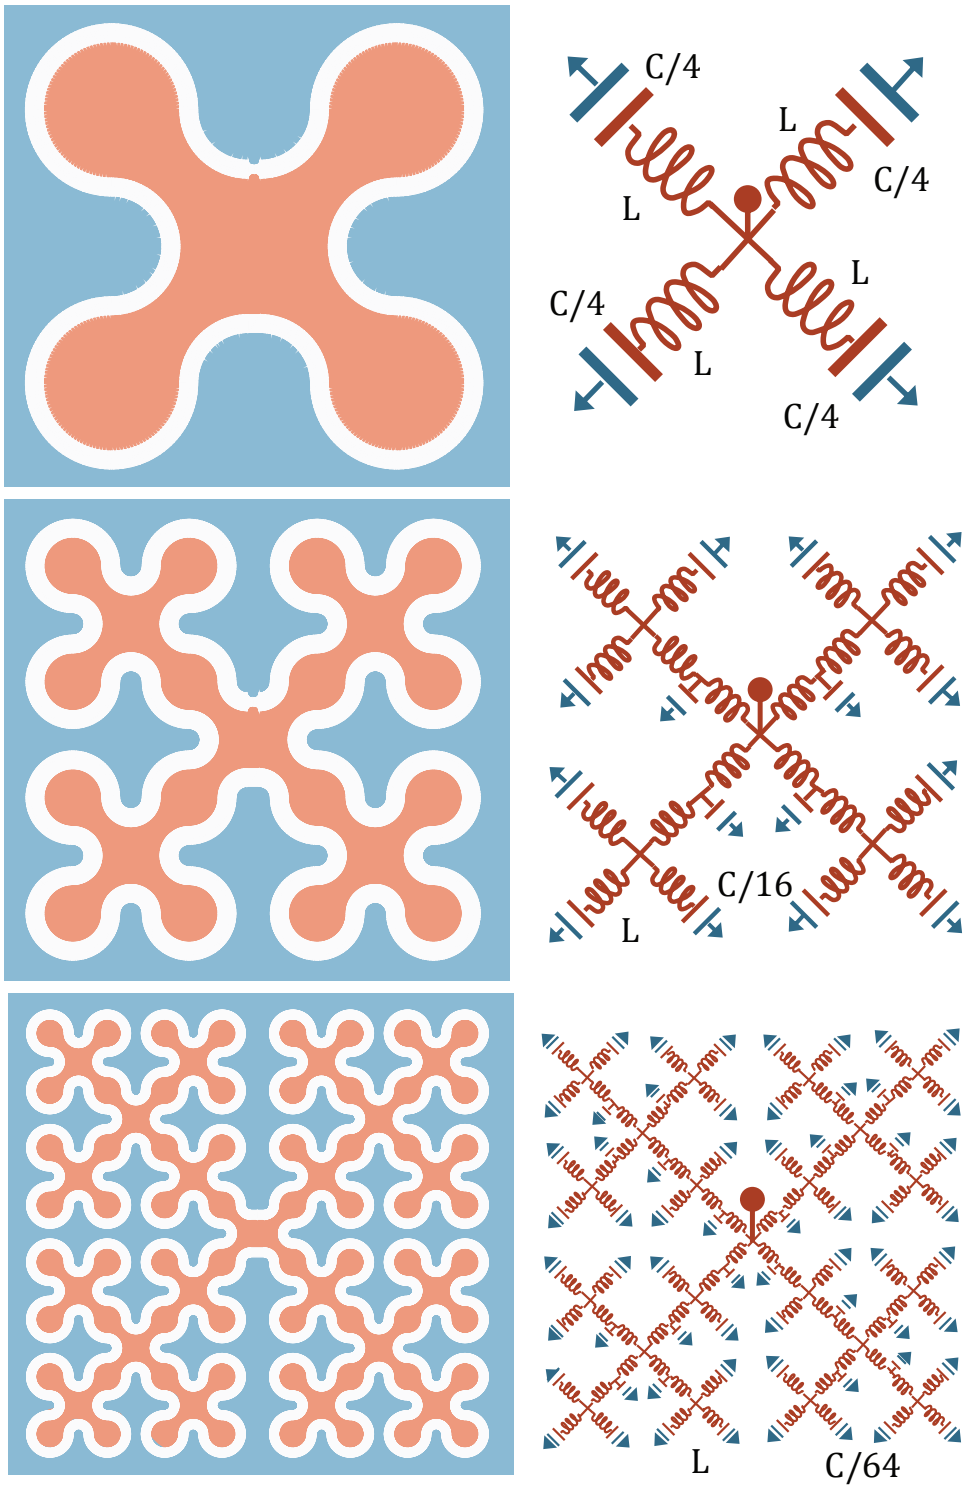

Figure 6.6: Parasitic networks for capacitors whose edges are defined by a Sierpinski fractal curve, shown for the first three fractal iterations.

dimensional fractal curve defines the gap between the capacitor electrode and the ground plane: since the Sierpinski curve is space-filling, each iteration of the fractal will significantly increase the capacitance density. The resulting geometry is shown in Figure 6.6. The first order fractal[9] is also the fractal generator: for each subsequent fractal iteration, each node is replaced by a copy of the generator. The resulting shapes are illustrated for the first three iterations in Figure 6.6.

Each node will contribute some capacitance, and the connections between nodes will contribute approximately one square of sheet inductance. Thus the equivalent parasitic circuit model can be described by a tree of successive capacitive and inductive elements, as shown in Figure 6.6. Notably each iteraction of fractal will have 4^n nodes, meaning that each node contributes a capacitance of $C/4^n$, and each node connection contributes one square of sheet inductance $L_\square$. The effective impedance of each iteration can be expressed in terms of the inductive impedance $Z_L = i\omega L_\square$ and capacitive impedance $Z_C = 4^n/i\omega C_0$ of each node:

$$Z_{n=1} = \frac{1}{4}(Z_L + Z_C) \qquad Z_{C1} = 4Z_{C,0} \tag{6.13}$$

$$Z_{n=2} = \frac{1}{4}\left(Z_L + \left(\frac{1}{Z_C} + \frac{1}{Z_L + \frac{1}{3}(Z_L + Z_C)}\right)^{-1}\right) \qquad Z_{C2} = 4^2 Z_{C,0} \tag{6.14}$$

And so on. The effective impedances for the first three fractal iterations shown in Figure 6.6 can be summarized as follows:

$$Z_{n=1} = \frac{Z_L + Z_{C1}}{4} \tag{6.15}$$

$$Z_{n=2} = \frac{Z_{C2}^2 + 8Z_{C2}Z_L + 4Z_L}{16(Z_{C2} + Z_L)} \tag{6.16}$$

$$Z_{n=3} = \frac{Z_{C3}^4 + 64Z_{C3}^3 Z_L + 288Z_{C3}^2 Z_L^2 + 256Z_{C3}Z_L^3 + 64Z_L^4}{64(Z_{C3} + Z_L)(Z_{C3}^2 + 8Z_{C3}Z_L + 4Z_L^2)} \tag{6.17}$$

9. It looks like the old Cingulair logo!

The impedance can be written out for higher fractal iterations as well, however even $n = 3$ fractal iterations already results in significant capacitance density for our purposes.

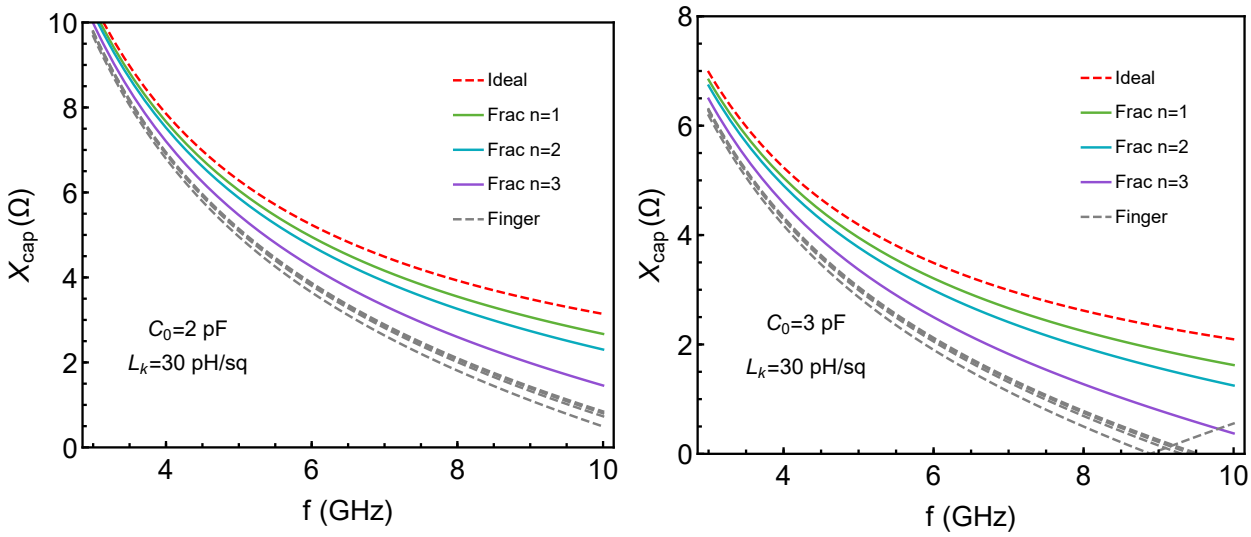

Figure 6.7: Effective impedance of a fractal capacitor with parasitic inductance taken into account, for the first three fractal iterations. Ideal capacitor model (dashed red) and equivalent footprint square finger capacitor (dashed gray) are also shown for comparison.

We compare the effective impedances for the first three iterationons of this fractal in Figure 6.7 for two large capacitances (2 pF and 3 pF) fabricated with a film with sheet inductance $L_\square = 30$ pH/□, which is typical for a high kinetic indutance material such as titanium nitride [204]. For the first three iterations, we see that each fractal capacitor still has lower parasitic inductance than finger capacitors with the same square profile[10]. This suggests that the Sierpinski fractal curve is an efficient method of introducing geometric complexity into the capacitor geometry while minimizing contributions of parasitic inductance (achieved by maximally branching the connections between capacitor nodes).

Indeed by the same logic a high-density capacitor with low parasitics could be created

10. This argument gets a lot more complex if we consider different aspect ratios of finger capacitors. For example you could reduce the parasitic inductance of a finger capacitor if you split it up into four finger capacitors each with lower aspect ratio on each side of a square center island. Equivalently you could choose a shape with more arms as the generator for the fractal (instead of a square snowflake it would become a hexagonal snowflake), while still applying the same fractal rules- while less space-efficient, this could help reduce the parasitic contributions.

with several other types of space-filling curves (such as the Peano curve [180], Hilbert Curve [88] or the flowsnake curve [142]). However, care must be taken to ensure that the fractal generator has second-order curvature (ie it should double back on itself), since otherwise the resulting capacitor island will not branch and the parasitic problem will become even worse.

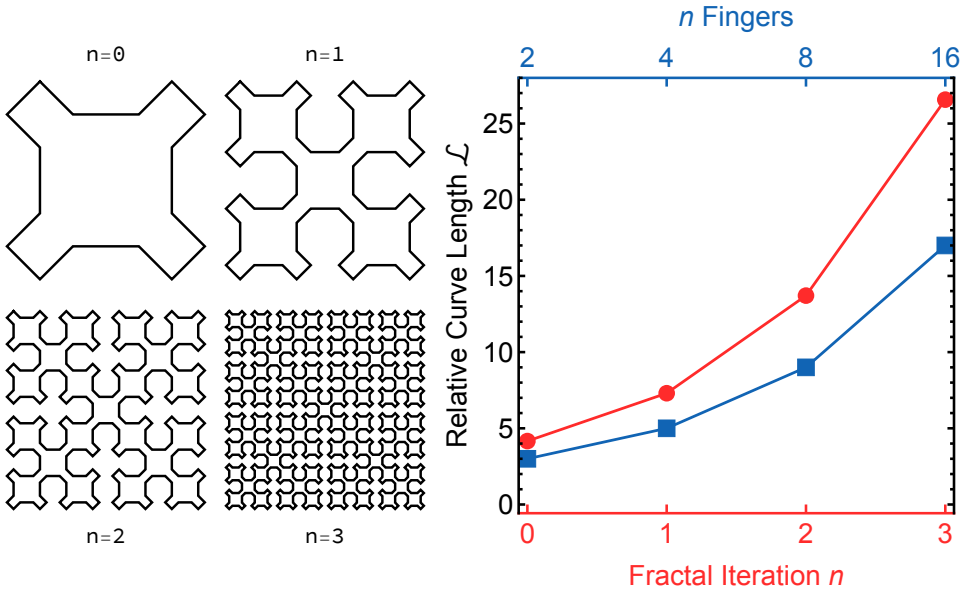

Figure 6.8: Generator and first three iterations of the square Sierpinski curve, on which our fractal is based, along with effective curve length relative to the overall width $\mathcal{L} = L_{\text{curve}}/W$ for fractals along with an equivalent square finger capacitor with n fingers. The curve length grows faster per iteration compared to the curve length of meander in a finger capacitor.

In the limit of infinite iterations, a space-filling fractal curve will completely fill its bounding area [50, 142], or in other words, the length of the curve will become infinite. This is a convenient property for us, since the capacitance between two islands separated by a curved gap will scale with the gap length: thus a fractal is an excellent way to maximize capacitance density. Consider the ideal square Sierpinski curve shown in Figure 6.8 (our design is based on this shape with the corners rounded). If we explicitly calculate the normalized fractal curve length $\mathcal{L}$ (defined as the curve length relative to the overall footprint width), we expect that the capacitance (proportional to $\mathcal{L}$) will scale exponentially with the number of fractal

iterations, as shown in Figure 6.8 The curve length (and thus capacitance) of a square aspect ratio finger capacitor scales with the number of fingers $\mathcal{L} = 2 + (n - 1)$ resulting in a power-law scaling which is slightly slower given a similar narrowest width. Combined with the reduced parasitic contribution to the capacitor, this makes fractals an ideal solution for increasing capacitance density in both high-frequency applications (such as for millimeter-wave circuits [65] or for hybrid microwave and millimeter-wave circuits [181]) as well as in cases where parasitic contributions reduce performance [99, 128].

6.4 Microwave Titanium Nitride Fractal Resonators

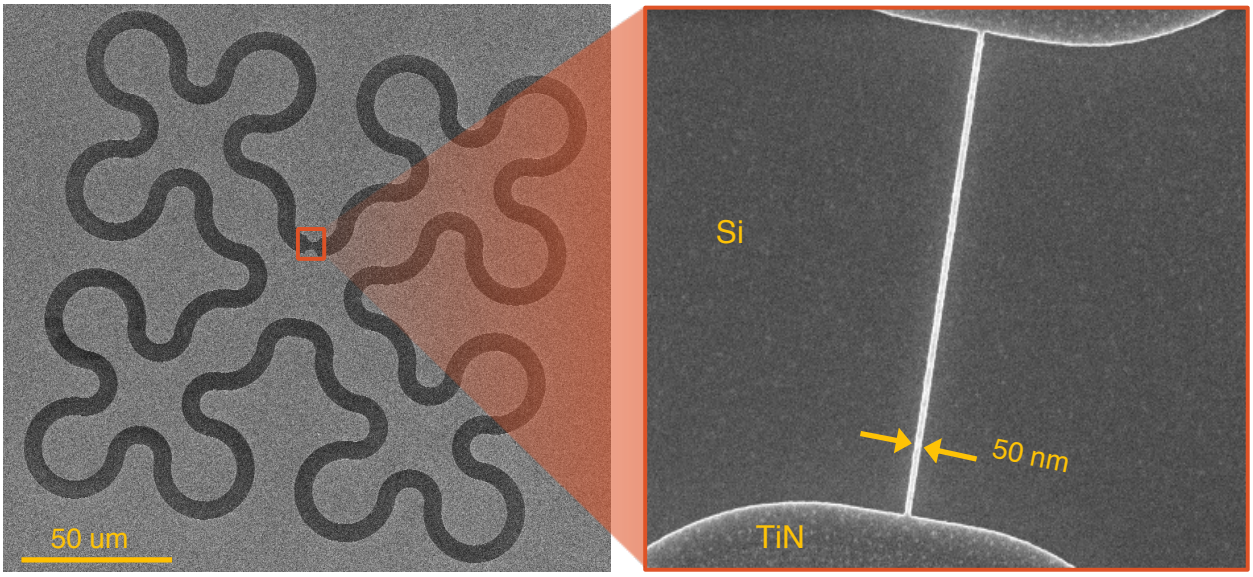

Figure 6.9: Scanning electron micrograph of a titanium nitride fractal capacitor (second order) shunted by a 50 nm wide nanowire. The circuit is patterned in one step on a crystalline silicon substrate.

We may have alluded to this already, but to make it absolutely clear, the devices we are discussing in this chapter are **microwave-frequency (4-8 GHz)** nonlinear resonators. Capitalizing on this well-established measurement platform, we can test out the impedance enhancement of low-impedance high-kinetic-inductance resonators (and potentially compo-

nents of the photon blockade) at microwave frequencies, with the idea that this can act as a model system for potentially similar applications at millimeter-wave frequencies [65].

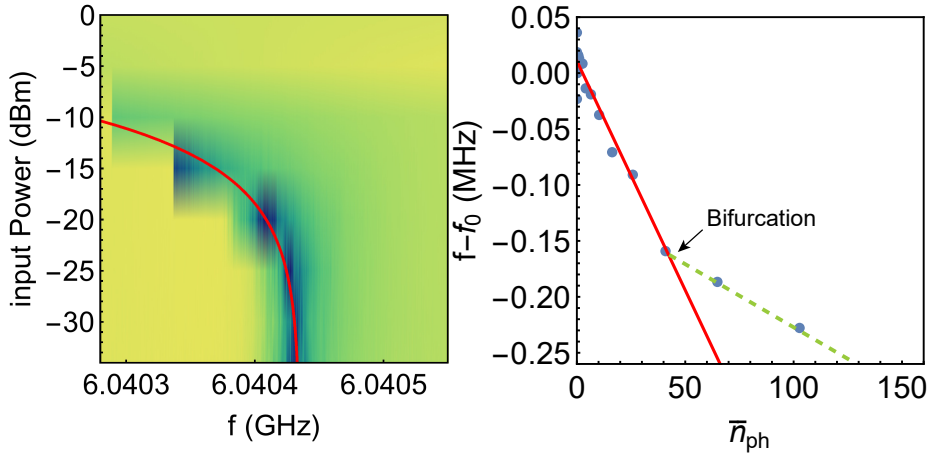

Figure 6.10: Transmission of a nonlinear resonator plotted for increasing powers. The frequency of the transmission dip shifts down as the circulating power in the resonator increases. By fitting the frequency of the transmission minimum can be used to extract the frequency shift, which is linear in resonator photon number, and can used to extract the Kerr nonlinearity. Above the bifurcation point, the transmission minima are determined by the bifurcation point instead of the resonant frequency (green dashed line).

We realize these nonlinear microwave resonators by patterning the fractal capacitor geometry described above, shunted by a narrow wire, as shown for a second-order Sierpnski fractal geomtry shown in Figure 6.9. The geometry is fabricated by first depositing titanium nitride using 300 cycles of atomic layer deposition (see Appendix B.2) on a crystalline silicon[11] susbstrate that has been cleaned and etched immediately prior to deposition. The capacitor geometry along with the impedance-optimized[12] coplanar waveguide (CPW) transmission line are patterned in a single exposure using electron-beam lithography, and the thin titanium nitride film is patterned by a chlorine reactive ion etch (see Appendix B.4). A rinse in deionized water helps dilute adsorbed chlorine. A finished resonator along with the

11. The wafer surface is the Si-111 orientation, since this is found to grow slightly nicer titanium nitride films [204].

12. See Chapter 2 for how kinetic inductance impacts coplanar waveguide impedance.

nonlinear nanowire is shown in Figure 6.9.

Using microwave spectroscopy, we characterize the resonators by measuring transmission through the central feedline. On resonance, we see a dip in transmission, which can be fitted to extract the resonator frequency and linewidth. As explored in Chapter 4, the frequency of a Kerr nonlinear resonator will shift as the applied power increases. We confirm this effect in Figure 6.10a. By measuring the fitted resonance frequency as a function of average resonator photon number (determined by the applied drive power) we can extract the nonlinearity per photon or the self-Kerr strength, as shown in Figure 6.10b. These measurements are particularly sensitive to the exact microwave power incident on the resonators, so a careful calibration experiment to verify the exact microwave attenuation is essential.

6.5 Nonlinearity Relative to Linewidth

As we discussed earlier, the quantity of interest is the nonlinearity strength relative to the total device linewidth K/κ. Having verified the nonlinearity of these high-kinetic inductance resonators using power-dependent microwave spectroscopy, we can investigate this in context of their linewidth. We summarize resonator nonlinearity relative to total device linewidth in Figure 6.11, grouped by the nanowire width w. Along with the low impedance fractal capacitor resonator designs, we have also included regular high impedance Jellyfish resonator designs (similar to those used in Ref. [204]). Consistent with Equation 6.6, we observe that thinner wires demonstrate higher nonlinearity strengths.

The effect of resonator impedance Z_0 is made clear in Figure 6.11b, where we plot resonator nonlinearity as a function of Z_0. We find that the nonlinearity indeed scales with impedance consistent with Equation 6.6, confirming that low-impedance resonators provide higher nonlinearity. For the highest fractal iterations with the narrowest 50 nm wide wires, we find a nonlinearity as high as 6.7 percent of the total device linewidth, which is substantial! The nonlinearity ratios in our devices are just under those measured in Ref. [99],

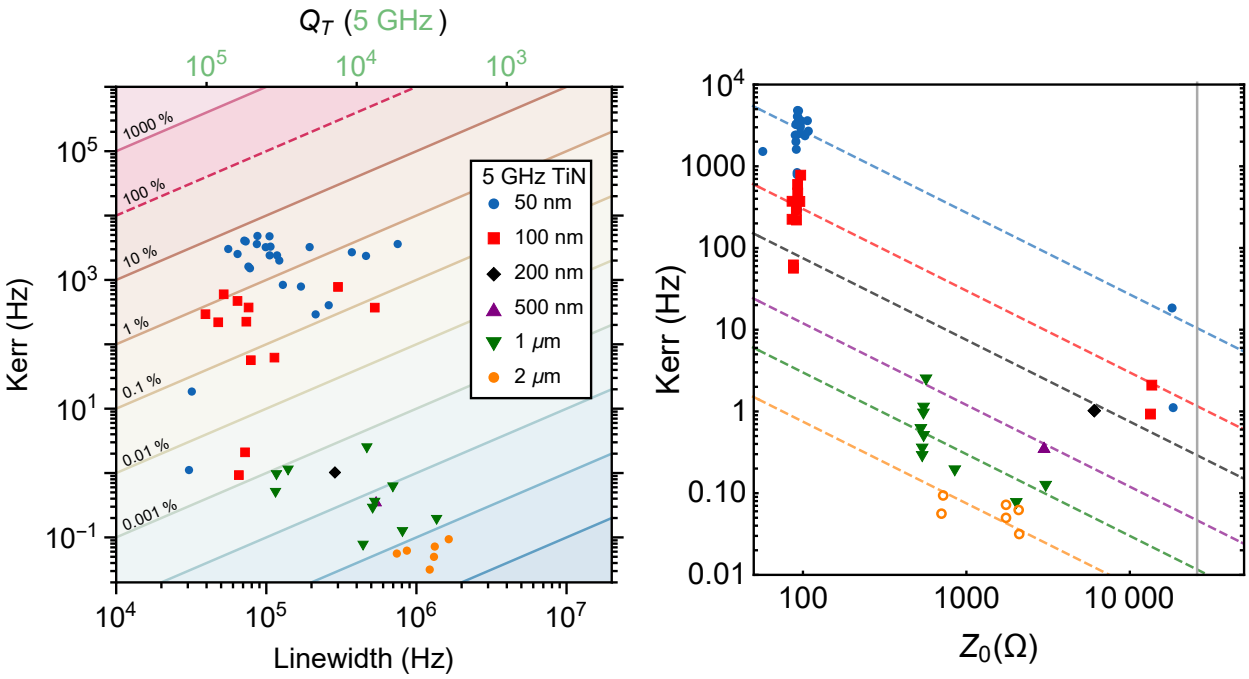

Figure 6.11: Left: Measured nonlinearity strength K for the titanium nitride devices in this chapter plotted as a function of total device linewidth γ_{Tot}, which is the combined internal and coupling linewidth. Right: Measured nonlinearity strength K for the titanium nitride devices in this chapter plotted as a function of device impedance, grouped by wire geometry. Dashed lines are the expected scaling from Equation 6.6, which takes into account wire geometry.

however accomplish this in a significantly smaller footprint[13].

We note that Figure 6.11 also portrays the same worrying trend for devices with high kinetic inductance explored in Chapter 5, where increased nonlinearity also results in increased linewidth arising from decoherence. It is likely that this is a manifestation of the same nonlinear resistance [264]. However this could also be the macroscopic result of incoherent phase slips in the nonlinear nanowire, which is suggested as a dissipation mechanism limiting the devices in Ref. [99]. While it is likely possible to push the inherent nonlinearity of kinetic inductance further, doing so in a coherent fashion will require careful attention to material purity.

13. What can fractals do for you?

6.6 Nonlinear Decay of a Coherent State

These high-kinetic inductance resonators do not have sufficiently high nonliearity to resolve individual energy transitions as done in [199]. However based on the nonlinearity strengths measured above, which puts us in an intermediate regime, we can investigate the classical or quantum mechanical nature of these devices.

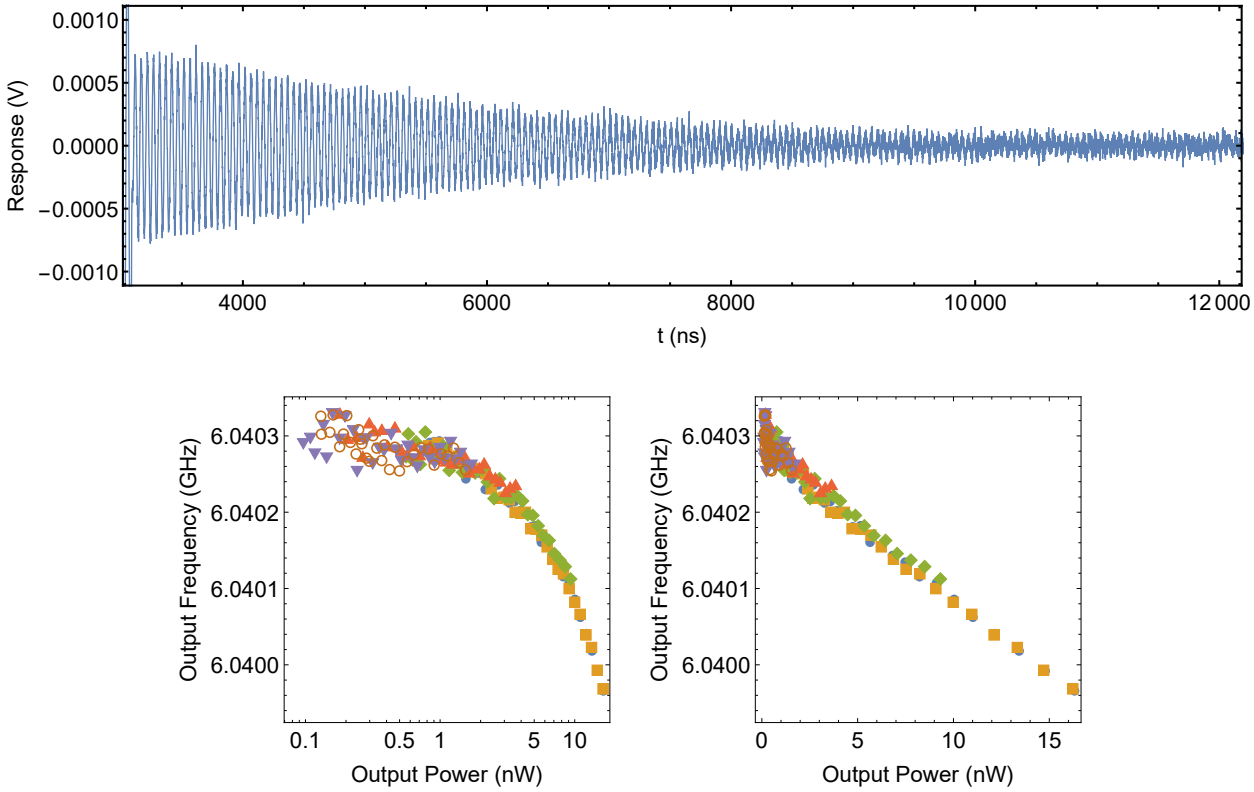

Figure 6.12: Top: Direct measurement of the ringdown chirp emitted by a decaying Kerr nonlinear resonator prepared in an initial coherent state α. The microwave ringdown was measured directly using an oscilloscope. Short incremental sections of the ringdown response can be fitted to a sinusoid function, which lets us track output frequency with respect to the instantaneous power of the emitted signal. For a weakly nonlinear system measured here, the resulting frequency is proportional to the emitted power, as expected for a Kerr nonlinearity.

With the blockade scheme from Ref. [128] in mind, consider our strongly nonlinear resonator prepared in a large coherent state $|\alpha_0\rangle$. The resonator photon number can be

expressed as $n = \langle\alpha(t)|\alpha(t)\rangle = |\alpha(t)|^2$. From a purely linear classical argument, we expect the resonator state to exponentially decay on a time scale set by the resonator linewidth $\alpha(t) = \alpha_0 e^{-t/\gamma}$. Based on the linearized argument, we expect the frequency of the resonator to depend on the resonator photon number $\omega = \omega_0 - |K|n$. Thus as the coherent state decays exponentially, so too should the resonator frequency. In other words, the resonator will emit a "chirp" as it decays, where the frequency of the signal changes over time.

In Figure 6.12 we prepare a resonator with fairly weak nonlinearity in a coherent state using a long microwave pulse with large amplitude, then directly record the signal emmitted by the resonator using a microwave oscilloscope. As expected, the emitted signal is indeed a chirp, with a time-dependent frequency. In agreement with the approximate classical argument used above, we find that the instantaneous frequency of the emitted pulse is consistent with the pulse amplitude.

Assuming the nonlinearity is relatively weak, we can use the mean field approximation [239] to solve for the system dynamics.

$$\dot{a} = [a, \mathcal{H}] - \frac{\gamma}{2}a \tag{6.18}$$

For a Kerr-nonlinear Hamiltonian without a drive term, this simplifies to

$$\dot{a} = i\omega_0 a - iKa^\dagger aa - \frac{\gamma}{2}a \tag{6.19}$$

For sufficiently weak nonlinearity we can linearize $a(t) \approx \alpha(t)$ and thus approximate the nonlinear term as $a^\dagger aa \approx |\alpha|^2\alpha$, which turns this into a simple differential equation, with the solution:

$$\alpha(t) = \alpha_0 e^{-\frac{\gamma}{2}t} e^{i\left(\omega_0 - K\frac{\alpha_0^2}{\gamma}(e^{-\gamma t}-1)\right)t} \tag{6.20}$$

Here α_0 is the initial state.

6.6.1 Heterodyne Measurements

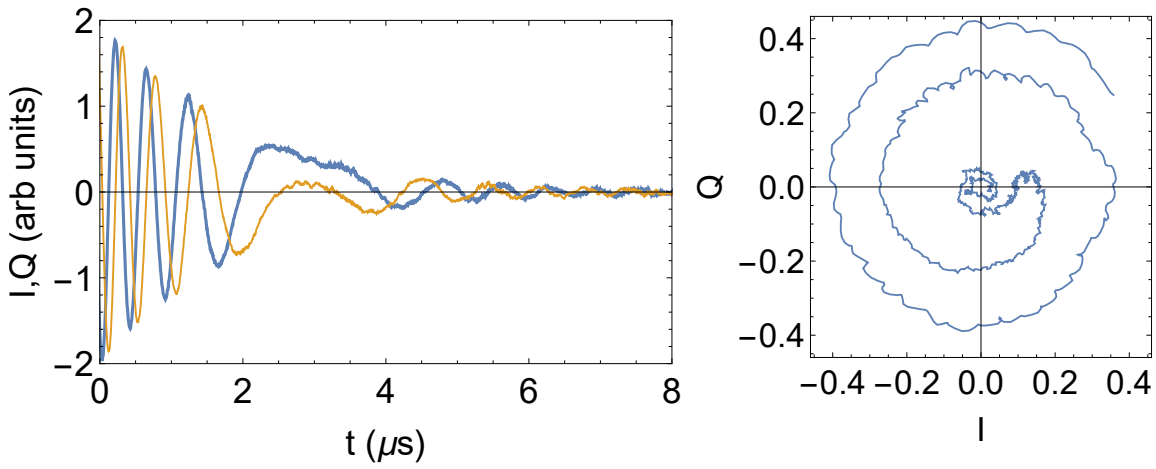

Figure 6.13: Heterodyne measurement of the ringdown chirp emitted by a decaying Kerr nonlinear resonator prepared in a coherent state α.

We now examine this behavior using more strongly nonlinear resonators. While directly measuring the emitted chirp was possible for the weakly nonlinear system, this proves experimentally challenging for high chirp frequencies, or fast chirp decays which rapidly change frequency. To remedy this, we perform a heterodyne measurement and measure the chirped response in the reference frame of the initial excitation pulse [14]. To maximize the initial coherent state (and thus initial detuning) the excitation pulse is chosen to be below the bare resonance frequency[15] The transformed chirp is shown in Figure 6.13. Initially, the resonator is excited, and its frequency starts below the reference frame frequency, so the initial signal oscillates. As the resonator decays, the chirp frequency returns toward its the bare resonant freuqency ω_0. As the frequency passes the refernce frequency, we observe the signal oscillations slow, then reverse their ferquency.

14. Technically there are two heterodyne steps: one to down-mix to an idler frequency and a second heterodyne experiment to determine the quadratures of the chirp

15. What is the optimal frequency for exciting the maximal coherent state? If we choose the bare resonant frequency the resonator will shift down, preventing power from building up. If we choose too low of a frequency we won't be able to pick up the resonator from its ground state! Really what we want is a chirp.

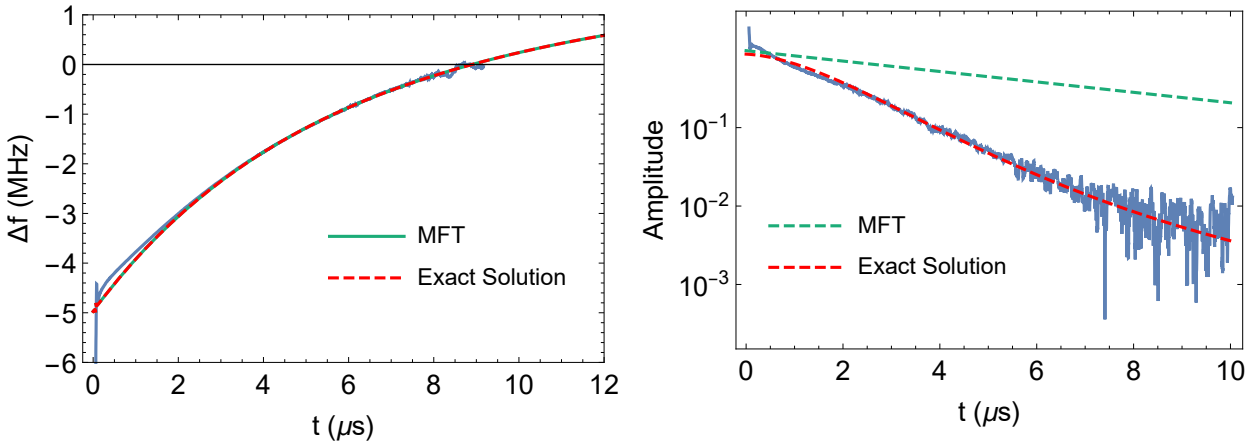

Figure 6.14: Ringdown Measurement of a highly nonlinear resonator prepared in a coherent state α. The amplitude decays much faster than expected from classical theory!

From these measurements we can extract the instantaneous frequency[16] and magnitude of the emitted chirp, and thus model the instantaneous resonator state and its frequency. The results are shown in Figure 6.14, with the classical solution (Equation 6.20) derived from mean field theory. Interestingly enough we find that while the classical theory describes the frequency of the chirp relatively well, the magnitude of the chirp decays significantly faster! Evidently we must use quantum mechanics to describe the nonlinear decay dynamics. The full solution (assuming a zero temperature bath) can be derived from Ref. [149]:

$$\alpha(t) = \alpha_0 e^{-\frac{\gamma}{2}t} e^{i\omega_0 t} \exp\left[-|\alpha|^2 K \left(\frac{K+i\gamma}{K^2+\gamma^2}\right)\left(1-e^{-(\gamma+iK)t}\right)\right] \tag{6.21}$$

As we see in Figure 6.14, the full solution does a much better job of describing the oscillator decay!

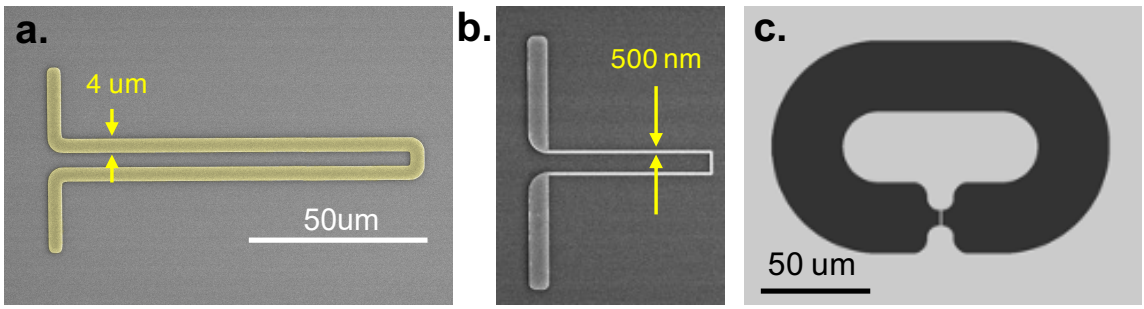

Figure 6.15: **a)** Regular millimeter-wave CPS resonator made from NbN on sapphire, with a standard 4 μm wire width. **b)** A similar CPS resonator with a 500 nm wide wire. The rest of the geometry is largely similar, so the impedance is not significantly changed from the standard design. **c)** A low-impedance millimeter-wave resonator, inspired by the titanium nitride fractal resonators. The required capacitance at 90 GHz is significantly smaller so only a zeroth-order fractal (in other words not a fractal at all) is required. Similar to the microwave design, a relatively thin wire shorts the capacitor to ground. However since this design is fabricated with optical lithography, the nanowires are between 1–2 μm wide.

6.7 Impedance Enhanced Kerr at 100 GHz

Having successfully demonstrated at microwave frequencies that lowering the resonator impedance increases the single-photon nonlinearity strength of the resonator, we attempt to replicate this at millimeter wave frequencies. Whereas our previous millimeter-wave resonator designs have used relatively long wires (even for the thinner case) as shown in 6.15a-b, we apply the low impedance methods by modifying the microwave fractal resonators for use at millimeter-wave frequencies, as shown in 6.15c.

Fabricating resonators with this geometry using similar niobium nitride films, we measure the nonlinearity of these devices, and compare them to the results from Chapter 4, as shown in Figure 6.16. We find that this results in a minor increase in nonlinearity. However the increase is less than a factor of 10, likely since the original millimeter-wave resonators already had relatively low impedance. Nevertheless by continuing to push in this direction, we could forseeably improve the millimeter-wave nonlinearity by a factor of 100 or 1000. Unfortunately due to the relatively large linewidths of the millimeter-wave devices caused

16. Doing this actually takes some thought, as the signal envelope also has a time derivative.

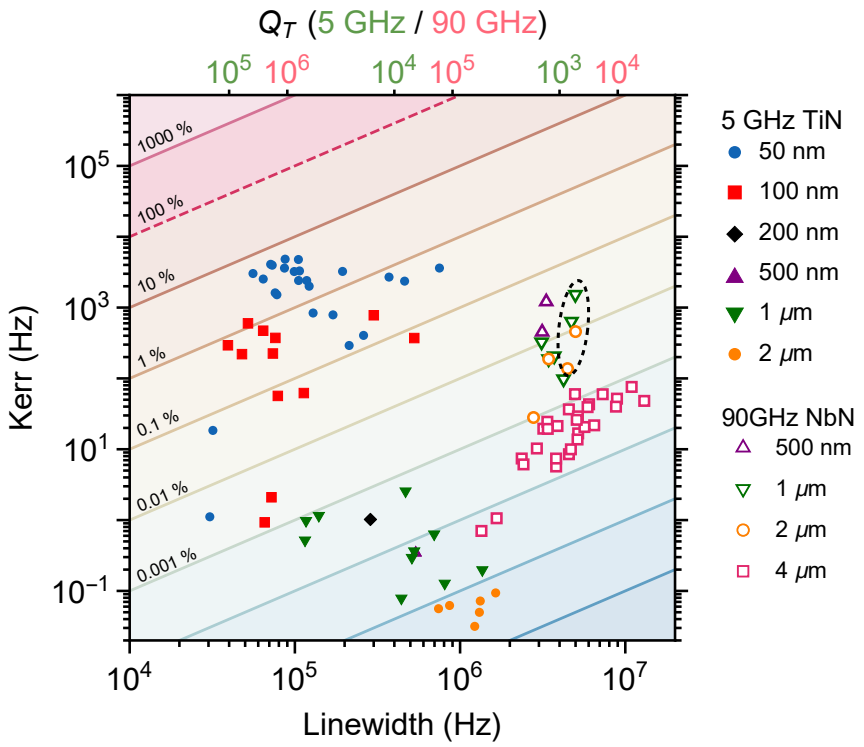

Figure 6.16: Nonlinearity (K) plotted as a function of total linewidth (γ_T) for the microwave titanium nitride resonators, the millimeter-wave NbN coplanar stripline resonators from last chapter, as well as millimeter-wave resonators with reduced impedance (highlighted by the dashed circle), all grouped by wire width. Notably the reduced impedance does help increase the self-nonlinearity of the millimeter-wave resonators. However this increase is less than one order of magnitude, likely since the original millimeter-wave resonators already had relatively low impedance.

by decoherence, this would probably still leave us in the same boat as the microwave devices in terms nonlinearity with respect to linewidth.

CHAPTER 7

NIOBIUM TRILAYER JUNCTIONS

Evidently kinetic inductance only goes so far. In the previous chapters we saw that kinetic inductance is an excellent tool for achieving weak nonlinearity: this property makes it a perfect choice for realizing nonlinear parametric amplifiers with very high compression powers [140]. However even with nonlinearity increasing techniques, using kinetic inductance to realize a single-photon-resolved nonlinearity appears to be a challenging task. So instead we turn to the workhorse of superconducting quantum devices: the Josephson junctions.

These superconducting circuit elements have enabled a wide variety of superconducting devices, with applications ranging from quantum-limited amplification and metrology [31, 95, 238] to digital logic [127, 226, 237] and they are an attractive platform for scalable quantum computing architectures due to their design flexibility and wide range of coupling strengths. Increasingly complex and robust quantum circuits have been demonstrated with aluminum junctions [111], however niobium is a tantalizing alternative superconductor due to its larger energy gap (and thus higher critical temperature and pair-breaking photon frequency) [66]. Taking advantage of this wider operating regime, niobium trilayer Josephson junctions became standard for single-flux-quantum circuits operating at liquid helium temperatures [127, 226, 237]. Employing these well-established fabrication processes, some early implementations of superconducting qubits were developed with niobium junctions [60, 100, 129, 143, 177, 247, 261, 262]. However, these initial niobium qubits only retained quantum state coherence for less than 400 ns, diminished by coupling to sources of dephasing and dissipation in the junction and the qubit environment.

Minimizing these loss sources is crucial in all sensitive quantum systems, but particularly for qubits, which must remain coherent over the duration of many gate operations. Significant effort has since been dedicated to investigating and reducing sources of decoherence [148], demanding either adjustments of circuit geometry to limit or dilute coupling to

spurious channels, or reducing the use of lossy amorphous dielectric materials. The need for insulated wiring contacts in these niobium trilayer junctions required growing passivating amorphous dielectric material in direct contact with the junction barrier, which likely degraded early qubit coherence [236], and limited their use in quantum devices. Higher temperature junctions with low loss promise a transformative source of strong nonlinearity for high-frequency quantum devices [12], and have since seen renewed interest from efforts to integrate digital and quantum logic [126, 131, 147], and the exploration of tunnel barrier materials beyond the limitations of aluminum [166, 260, 268]. Notably, by removing amorphous insulating scaffolding and increasing the circuit volume to reduce junction participation, qubits with epitaxially grown NbN junctions with crystalline AlN barriers have increased coherence times to 16 µs [109]. We will apply similar improvements to traditional Nb/Al/AlO_x processes, which are attractive due to the simpler deposition methods required.

In this chapter, we use an improved fabrication method to revisit niobium trilayer junctions as the core component of transmon qubits and explore their coherence properties. We describe a method to form a temporary self-aligned sidewall-passivating spacer structure based on Ref. [81], which limits the amorphous spacer material to the smallest necessary region, and can later be chemically removed to further reduce dielectric loss. We find that high-temperature spacer growth methods greatly reduce the critical current density of the junction barrier, allowing us to utilize exclusively optical lithography to fabricate high-nonlinearity junctions for microwave qubits. We find that our all-niobium qubits have lifetimes as high as 62 µs with an average qubit quality factor of 2.57×10^5: much closer to state-of-the-art qubits than past Nb/Al/AlO_x devices [60, 100, 129, 143, 177, 247]. We further observe that the higher superconducting gap energy results in reduced sensitivity to quasiparticles, particularly above 160 mK, where conventional aluminum-junction qubit performance deteriorates [47, 178, 189]. These results demonstrate the reemergent relevance of niobium junctions for pushing the boundaries of superconducting devices.

7.1 Trilayer Fabrication

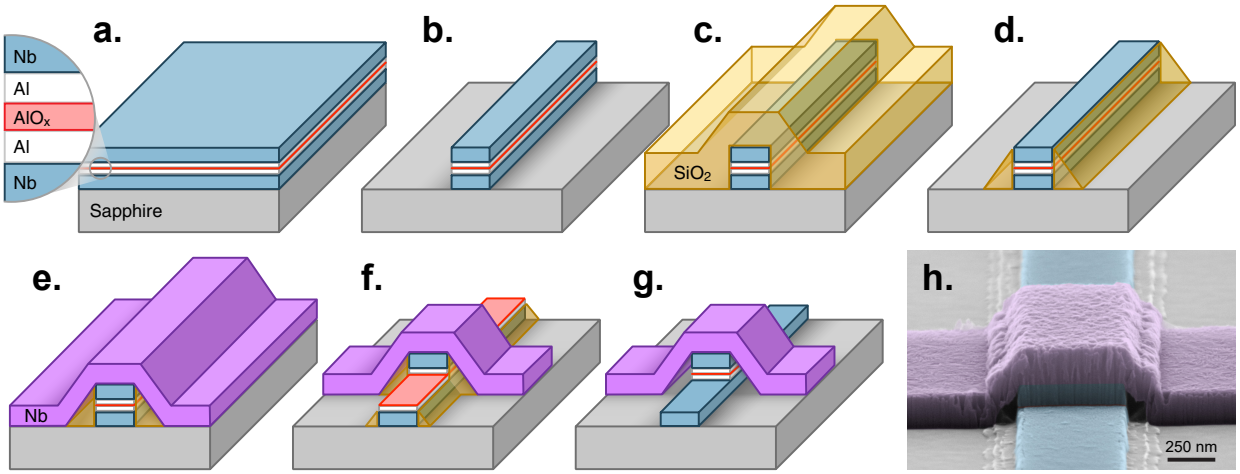

Figure 7.1: Junction fabrication process. (a) Trilayer is deposited and oxidized in-situ. (b) First layer is etched with a chlorine RIE. (c) SiO_2 is grown isotropically. (d) Sidewall spacer is formed by anisotropic etching with fluorine chemistry. (e) Surface oxides are cleaned in vacuum and wiring layer (purple) is deposited. (f) Second junction finger (and other circuit elements) are defined by a fluorine plasma etch selective against Al. (g) Final devices undergo a wet etch to further remove SiO_2, exposed Al and some NbO_x. (h) Color-enhanced electron micrograph of a finished trilayer junction with an area of $\sim (500\text{ nm})^2$.

Despite niobium's attractive electrical properties, in thin layers its oxides are imperfect insulators with high dielectric loss [236], resulting in very poor natural tunnel junction barriers. Aluminum, on the other hand, forms a thin self-terminating oxide with low leakage and loss, but has a low critical temperature. The trilayer method leverages the strengths of both of these materials by using a thin layer of oxidized aluminum as the tunnel barrier and encapsulating it with niobium: through the proximity effect the Josephson junction inherits desired electrical properties and a clean tunnel barrier. This trilayer structure is typically grown on a wafer-scale as the first step in fabrication, enabling excellent uniformity [37, 228] and high purity growth methods. The junction fabrication process is illustrated in Figure 7.1. But you probably came here for the details right?

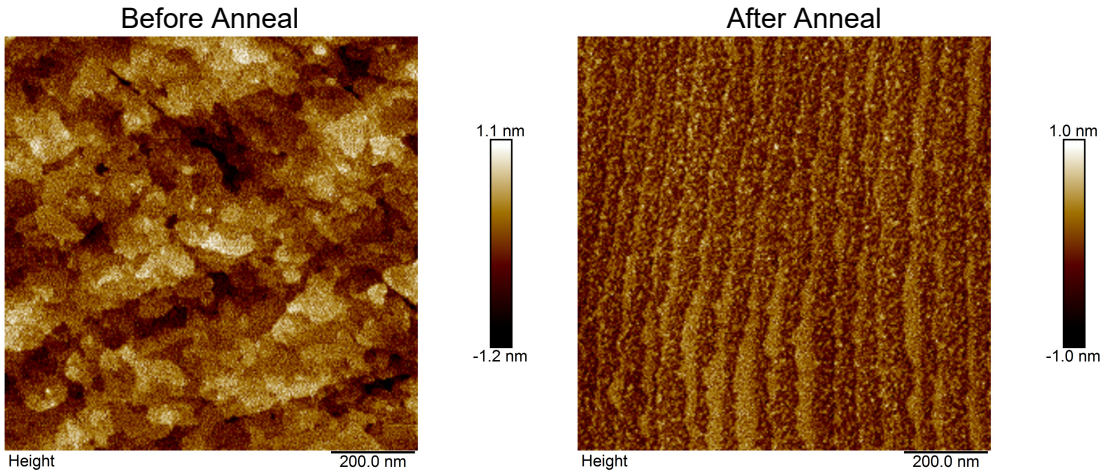

Figure 7.2: Atomic-force microscopy measurements of the C-plane sapphire surface. As received (left), the surface is still fairly rough after the epitaxial surface polish, however after a 1.5 hour anneal at 1250 °C in air (right) atomic terraces are visible, meaning the surface is extremely flat. Annealing in nitrogen instead results in oxygen depletion on the sapphire surface, creating a coral-like surface. Data courtesy of F. Zhao.

7.1.1 Wafer Preparation and Trilayer deposition

Similar to methods using sputtering, our trilayer is formed in a shadow-evaporation-compatible electron-beam system. High-purity single-crystal C-Plane sapphire wafers that have been annealed (See Figure 7.2) and chemically etched to remove surface damage are ultrasonically cleaned, then etched in a 40 °C piranha solution. Immediately following, the wafers are loaded into a Plassys MEB550S electron-beam evaporation system, where they are baked by heating the stage to >200 °C under vacuum for an hour to help remove water and volatiles. When a sufficiently low pressure is reached ($< 5 \times 10^{-8}$ mbar), titanium is electron-beam evaporated to bring the load lock pressure down even further. We then form the trilayer by depositing 80 nm of Nb (at >0.5 nm/s) then 8 nm of Al. The deposition rate is kept high to maximize film quality (see section 7.3). To reduce defects and promote aluminum oxide formation [30], the aluminum is first ion milled then oxidized with an O_2-Ar mixture. To prevent oxygen diffusion into the Nb layer and the formation of lossy NbO_x [227, 236], the oxidized Al surface is protected by a thin (3 nm) capping layer of Al. The aluminum is lightly

etched with a 400 V Ar^+ beam for 10 s, then oxidized with a mixture of 15% O_2:Ar at a static pressure between 2–50 mBar for 1.5–40 min. After pumping to below ($< 10^{-7}$ mbar), titanium is again used to bring the vacuum pressure down to the low 10^{-8} mbar range. We note that the pressure for the remainder of the trilayer deposition is higher than for the first Nb layer. The second 3 nm layer of Al is evaporated vertically while rotating the substrate to minimize void formation in the following layer. A (150 nm) thick counter electrode is then deposited on top, forming the trilayer in-situ, without breaking vacuum. The substrate is allowed to cool in vacuum for several minutes, and we attempt to form a thin protective coating of pure Nb_2O_5 by briefly oxidizing the top surface at 3 mbar for 30 s. This layer is deposited while rotating the substrate at an angle for complete coverage while keeping it thin enough to avoid affecting junction properties.

7.1.2 Trilayer Etch 1

After mounting on a silicon handle wafer with photoresist, the trilayer is patterned with I-line (375 nm laser) photolithography and the entire trilayer structure is plasma-etched in one step with Cl_2, BCl_3 and Ar to define the bottom electrode (Etch 1 in Table 7.1). The plasma conditions are optimized to be in the ballistic ion regime, which gives high etch rates with minimal re-deposition. Immediately after exposure to air, the wafer is quenched in DI water: this helps prevent excess lateral aluminum etching by quickly diluting any surface HCl (formed by adsorbed Cl reacting with water vapor in the air). The remaining photoresist is thoroughly dissolved in a mixture of 80 °C n-methyl-2-pyrrolidone with a small addition of surfactants, which also removes the substrate from the handle wafer.

7.1.3 Spacer Growth

As it is necessary to make contact to the counter electrode without touching the base electrode, we then form an insulating sidewall-passivating spacer structure [81]. Amorphous

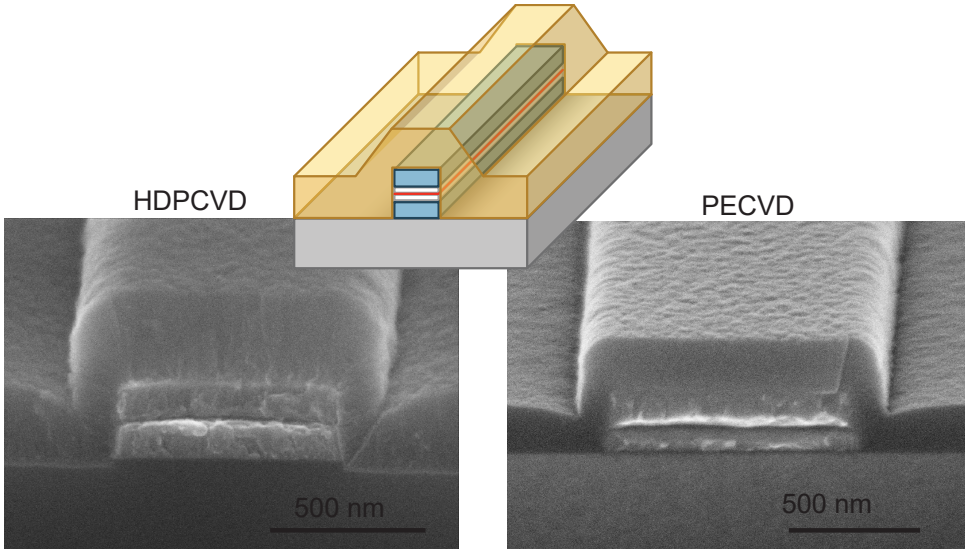

Figure 7.3: Spacer Growth, imaged in cleaved samples from HDPCVD growth methods and PECVD growth methods. Note the superior conformality of PECVD, while HDPCVD has two phases and a breadloaf cross section.

SiO_2 is grown isotropically (Figure 7.3) by either plasma-enhanced chemical vapor deposition (PECVD) which heats the wafer to 300 °C for 16 min or high density plasma-enhanced chemical vapor deposition (HDPCVD) (90 °C). For PECVD, SiH_4 and N_2O are reacted in a 100 W plasma with the chamber at 300 °C. The complete process (including chamber cleaning pumping and purging steps) takes approximately 15 minutes. For HDPCVD, the wafer is mounted on a silicon handle wafer using Crystalbond 509 adhesive softened at 135 °C, then the spacer is deposited with a SiH_4 O_2 and Ar plasma, with the substrate heated to 90 °C.

7.1.4 Spacer Etch 2

The SiO_2 is now etched anisotropically with a highly directional CF_4, CHF_3 and Ar plasma (Etch 2 in Table 7.1), which forms the spacer structure when the bulk material has been etched away (Figure 7.1d). This etch is optimized to be directional but in the diffusive regime to promote chemical selectivity while enabling the formation of the spacer structure.

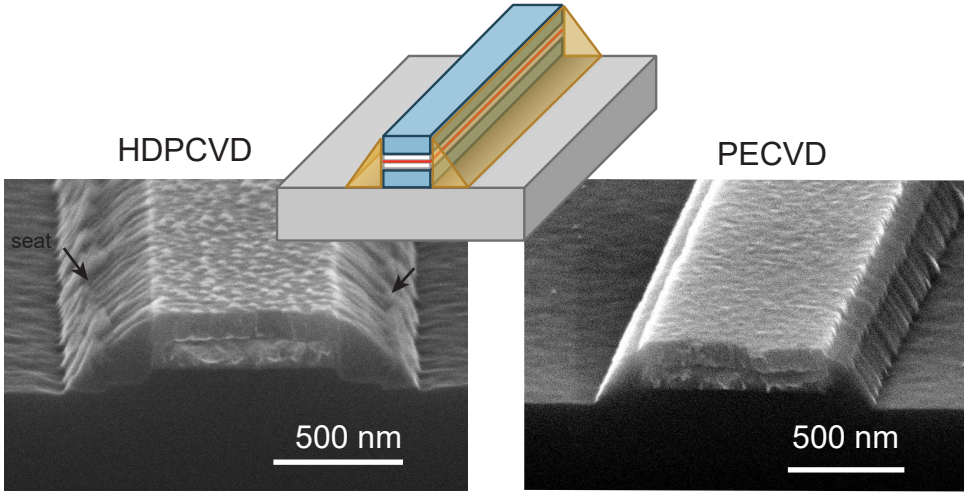

Figure 7.4: Sidewall-passivating spacer after etching, imaged in cleaved samples from HDPCVD growth methods and PECVD growth methods. The smooth conformal PECVD growth translates to a smooth spacer profile, while the HDPCVD results in a chair-like structure, with potentially lossy discontinuities in the seat area.

The etch is terminated using visual endpoint detection[1]

At this point minimizing oxide formation is crucial since the top surface of the trilayer is exposed and will need to form a good contact to the wiring layer, so immediately following the completion of the etch, wafers are separated from the handle wafer by heating[2] to 135 °C, ultrasonically cleaned of remaining adhesive in 40 °C acetone and isopropanol, then immediately placed under vacuum in the deposition chamber, where they are gently heated to 50 °C for 30 min to remove remaining volatiles.

7.1.5 Wiring Layer Deposition

The contaminated and oxidized top surface of the counter electrode is etched with a 400 V Ar^+ beam for 5 min, which is sufficient to remove any residual resistance from the contact (see Section 7.3). After pumping to below ($< 10^{-7}$ mbar), titanium is used to bring the vacuum pressure down to the low 10^{-8} mbar range. The wiring layer is now formed by

1. In other words I determined when to stop by looking at the color!

2. This heating step can affect the junction barrier!

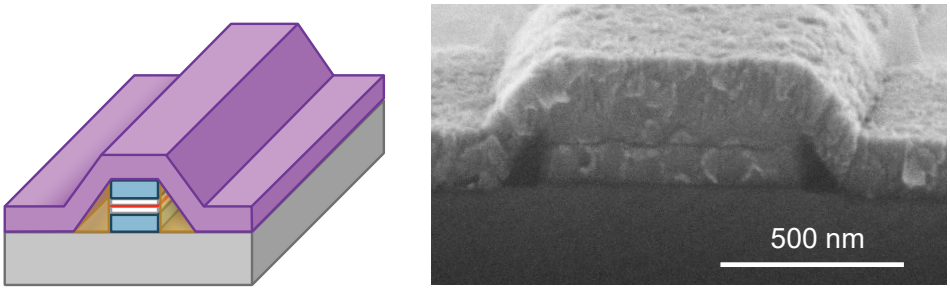

Figure 7.5: When cleaving a wafer I got lucky and a fault line ran right through a junction, allowing us to see what the trilayer stack looks like after the wiring layer is deposited. The spacer is still intact here, and you can see no interface between the wiring and counterelectrode indicating good contact.

evaporating 160 nm of Nb at > 0.5 nm/s. The substrate is allowed to cool in vacuum for several minutes, and the wiring layer is briefly oxidized with 15% O_2:Ar at 3 mbar for 30 s to promote a thin protective coating of pure Nb_2O_5.

7.1.6 Wiring Etch 3

The wiring layer is patterned and a selective SF_6, CHF_3, O_2 and Ar plasma etch removes the wiring layer and the counter electrode, defining the perpendicular top junction electrode (Figure 7.1f). This etch is carefully optimized to minimize the formation of lossy fluorocarbon polymers [45] (see Appendix 7.4) while preserving chemical selectivity: and although the plasma etches the Al layers far slower than Nb, the etch is still timed to finish a few seconds after the counter electrode is fully removed to limit excessive polymer deposition.

This step proves to be highly problematic as it easily forms inert residues, and needs to be highly chemically selective in order to avoid etching through the aluminum, so the plasma is operated in a low-density ballistic regime with the addition of O_2 which helps passivate exposed aluminum and increase selectivity. The etch time is calculated for each wafer based on visual confirmation when the bare wiring layer is etched through. In some cases we attempt to remove crosslinked polymers from the photoresist surface with a mild 180 W room temperature oxygen plasma that minimally oxidizes the exposed Nb (though

find this is not very effective).

Etch	T(°C)	Pressure	ICP/Bias Power	Cl_2	BCl_3	Ar	CF_4	CHF_3	SF_6	O_2
1	20 ± 0.1	5 mT	400 W / 50 W	30	30	10	-	-	-	-
2	20 ± 0.1	30 mT	500 W / 60 W	-	-	10	30	20	-	-
3	20 ± 0.1	5 mT	400 W / 60 W	-	-	7	-	20	40	4

Etch	etch time	etch rate
1	50-60 s	~ 4.5 nm/s
2	120-140 s s	~ 2 nm/s
3	65-90 s	~ 4.5 nm/s

Table 7.1: Plasma etch parameters used in the ICP-RIE etches described in the process. Etches are performed in an Apex SLR ICP etcher. Gas flows are listed in sccm.

7.1.7 Spacer removal

After dicing into 7 mm chips, the qubits are cleaned in 80 °C n-methyl-2-pyrrolidone with surfactants (we find this can also help remove stubborn organic residues from previous steps). Finally, to further remove the lossy amorphous materials present in the junction, a solution of NH_4F and acetic acid [3] are used to dissolve the remaining SiO_2 then carefully dried from isopropanol to preserve the now partially suspended wiring layer. This process additionally removes any exposed Al and a small amount of surface NbO_x (Figure 7.1g). As this step can dissolve aluminum in the junction as well, etch times are kept below 15 s. This final treatment could likely be improved with a HF vapor etch, which has shown good results forming similar contact structures [59]. However we found this vapor process difficult to optimize, since liquid frequently builds up on surfaces and begins to etch niobium, particularly around corners.

The finished chips are packaged and cooled down within a couple hours from this final etch to minimize any NbO_x regrowth from air exposure.

3. Transene AlPAD Etch 639

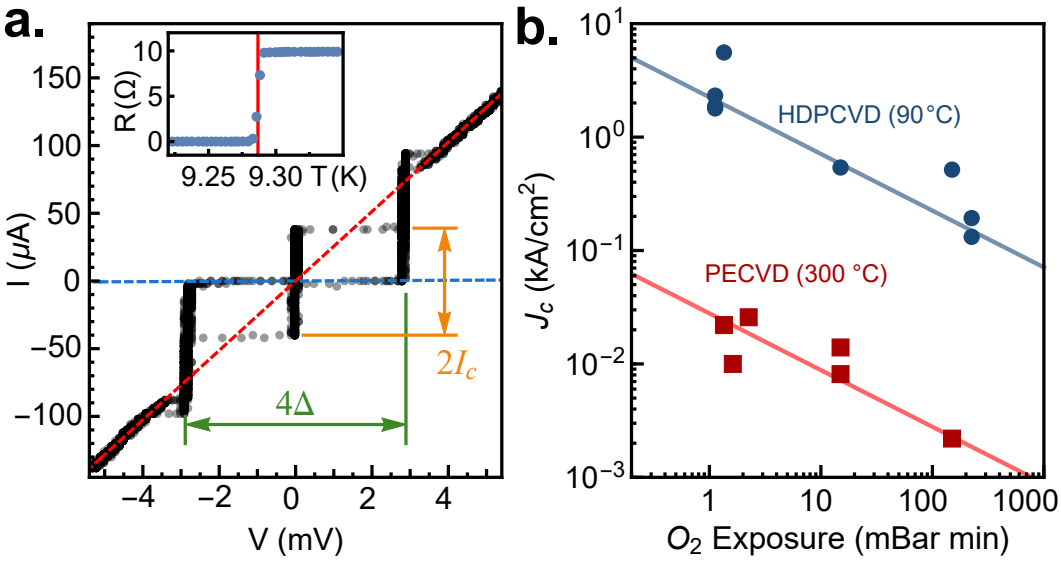

Figure 7.6: Junction properties. (a) Current-voltage relations for an un-shunted junction at 860 mK with $I_c = 38\,\mu$A and an energy gap $2\Delta = 2.89$ meV. Bulk resistivity measurements (inset) give a critical temperature of $T_c = 9.28$ K. Above 4 mV, a linear fit (red dashed line) gives $R_n = 39\,\Omega$, and a fit to the sub-gap region (blue dashed line), estimates sub-gap resistance $R_s > 8\,\text{k}\Omega$. (b) Critical current density J_c (found by fitting room-temperature junction resistance as a function of junction area) as a function of oxygen exposure E measured for various wafers made with two deposition processes. The expected empirical $E^{-1/2}$ relationships are plotted as guides to the eye.

We verify the expected Josephson junction behavior [214] in our devices by measuring their hysteretic current-voltage curves in Figure 7.6a. When cooled to 860 mK, the un-shunted junction shows a zero-resistance superconducting branch up to the critical current I_c, and an energy gap $2\Delta = 2.89$ meV. By comparing this value with critical temperature measured with resistivity, we find a relationship $2\Delta/k_BT_c = 3.61$: slightly lower than reported values for pure Nb [173, 231]. Measuring the asymptotic normal state resistance R_n above the energy gap we find a I_cR_n product of 1.5 mV, similar to values reported previously for Nb trilayer junctions [37, 81, 162, 225, 228]. Although measurements of the subgap region were limited by the experiment hardware, no excessive subgap leakage currents are observed.

Using the I_cR_n product found above, we can use room-temperature junction resistances

to predict low-temperature properties [10, 76]. Fitting the measured resistance for junctions of varying areas with two free parameters, specific resistivity and junction critical dimension bias (see Appendix 7.5), we obtain the effective junction areas and the specific critical current density J_c for each wafer. This method allows us to easily investigate effects of the fabrication process on junction electrical parameters. For Nb trilayers, the critical current density is sensitive to temperature [162] as well as oxygen exposure E, the product of oxygen partial pressure and oxidation time: this relationship has been empirically found to match $J_c \propto E^{-0.5}$ [81, 112, 162, 216, 248]. In Figure 7.6b we plot J_c as a function of E for wafers with trilayers grown using various oxidation parameters and fabricated with two spacer deposition methods. For the HDPCVD junctions, we find critical current densities in the $\mathrm{kA\,cm^{-2}}$ range, comparable with other methods [11, 81, 162, 225], and observe reasonable agreement with the oxygen exposure dependence described above. The effect of process temperature is readily apparent when we examine junctions with high-temperature-grown PECVD spacers: compared to HDPCVD junctions, we observe nearly a factor of 50 reduction in J_c. We find this temperature-annealing effect activates above 200 °C (see Appendix 7.6), in agreement with [162], and is likely the result of reduced barrier transparency [202] from diffusion.

7.3 Junction Superconductor Properties

Josephson junction properties are largely determined by the characteristics of the two superconductors and the insulating oxide barrier that separates them, so the initial formation of the trilayer materials is crucial for the device quality. As niobium sets the limit of superconducting properties and losses in our junctions and qubits, it is crucial to begin with a high-quality and thus high-purity material. Maintaining material purity presents a challenge for any thin film deposition technique, made difficult in particular by the incorporation of contaminants into the film during growth. This contamination can be addressed with two main approaches: first by reducing the flux of contaminants (achieved by reducing the vac-

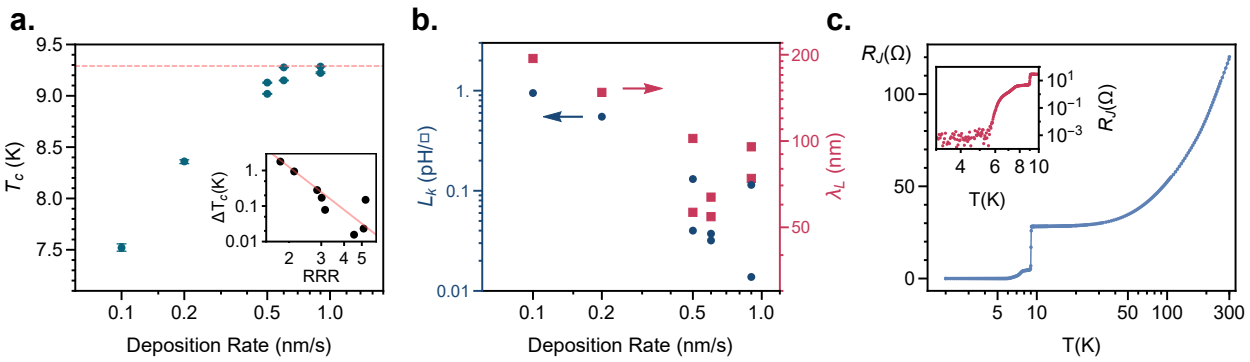

Figure 7.7: Superconductor material quality. (a) Niobium superconducting critical temperature T_C extracted from resistivity measurements as a function of metal deposition rate. At rates above 0.6 nm/s, T_C approaches bulk value (dashed line). The inset shows deviations from bulk $\Delta T_C = T_C^{\text{bulk}} - T_C$ are correlated with the residual resistivity ratio, implying high deposition rates result in high-quality films. (b) Sheet kinetic inductance L_k and observed London penetration depth λ_L plotted as a function of deposition rate suggesting that films deposited at higher rates are closer to the clean superconductor limit. (c) Specific junction resistance $R_J = R/N$ obtained by measuring the resistance R of a chain of $N = 12$ junctions as a function of temperature. A sharp drop in resistance is observed above 9 K as the niobium electrodes begin to superconduct. As the temperature decreases, the junction critical currents increase above the excitation current (10 µA), and below 5 K the measured resistance drops to zero as the excitation is confined to the superconducting branch, indicating proximitization of the aluminum and superconducting contact between the counterelectrode and wiring layers.

uum pressure during the deposition process), but also by reducing the duration of exposure, which can be controlled by the deposition rate.

For electron beam evaporation (the deposition technique used here) vacuum pressures are reduced as low as possible during deposition, however are limited to the 10^{-8} mbar range by the hardware. With the contaminant flux fixed by the deposition system vacuum pressure, we explore the effect of deposition rate on Nb purity. By measuring the resistivity of a film with a known geometry at varying temperatures, we obtain a wealth of information about the film properties. In Figure 7.7a we plot the superconducting transition temperature T_C (proportional to the superconducting gap Δ_0) as a function of metal deposition rate. We observe that higher rates yield increased transition temperatures, which approach those found in bulk high-purity Nb [251], indicating that the films are increasingly pure. Indeed,

we can also correlate the residual resistivity ratio $RRR = \rho(300\,\mathrm{K})/\rho(T_C)$, an indicator of superconductor quality, with deviations of measured critical temperature the bulk value T_C^{bulk}, supporting the notion that higher deposition rates yield higher-quality films. Due to the extreme local temperatures required, practical considerations and stability concerns put a limit on feasible deposition rates. Nonetheless, despite variations induced by vacuum conditions, we find that rates above 0.6 nm/s are required to deposit a film with high purity.

We can go further to examine the degree of disorder in the superconductor by probing the kinetic sheet inductance $L_K = \hbar R_\square/\pi\Delta_0$ where $R_\square = \rho_0/t$ is extracted from the film thickness t, and the resistivity just above the superconducting transition. The sheet inductance also yields the London magnetic penetration depth $\lambda_L^2 = tL_K/\mu_0$. In Figure 7.7b we find that both L_K and λ_L are also reduced with films deposited at higher rates. Lower kinetic inductance and shorter London lengths indicate a lower degree of disorder in the superconductor, suggesting that increased deposition rates bring the material further away from the disordered dirty superconductor limit ($\lambda_L \gg \xi$) [134].

We verify the superconducting contact quality between the wiring layer and the counter-electrode, as well as the junction tunnel barrier transparency by measuring the voltage accross a chain of 12 junctions in series, through which we send a fixed excitation current of 10 µA. In Figure 7.7c we plot the per-junction specific resistance R_J as a function of temperature, showing the immediately apparent superconducting transition above 9 K. Immediately below the transition, the superconducting gap is still relatively low, and the junction critical currents fall below the excitation current, so a small resistance is observed. However as we decrease the temperature, we find that the resistance shrinks by several orders of magnitude (below the noise floor of the instrument). This indicates that the sum of any remaining resistance channels in a single junction is likely well below the mΩ range, suggesting a superconducting contact between the Nb wiring layer and the Nb counter-electrode.

In a superconductor the residual resistivity ratio is also correlated with grain size in the

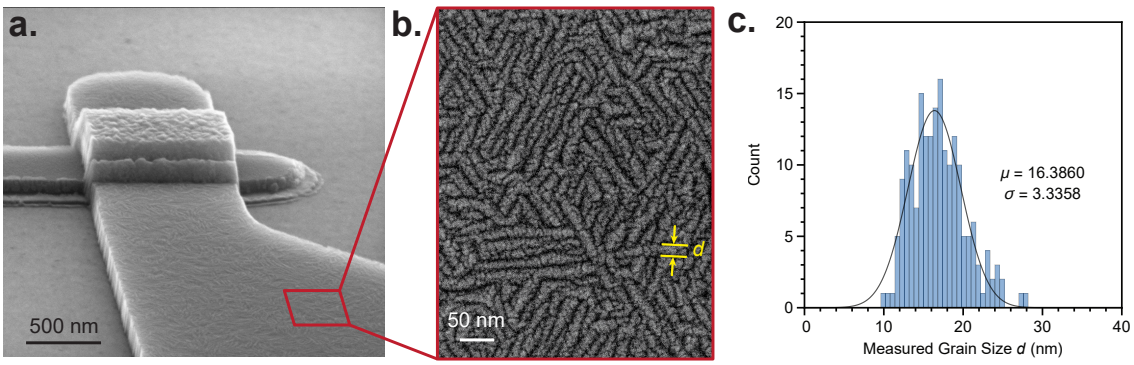

Figure 7.8: Superconductor grain size. (a) In a tilted scanning electron microscope image of a junction, microscopic grains are observed on the metal surface. In regions of the wiring layer that lie directly on the sapphire substrate, the columnar grain growth is uninterrupted, and the grain pattern is transferred to the top surface of the metal. (b) A top down high-resolution scanning electron micrograph reveals the hexagonal arrangement of the grains. The grain size can be estimated by measuring the narrow dimension of a grain, marked d. (c) A histogram of repeated measurements of grain width are fitted to a normal distribution which suggests an average grain width of 16.386 nm.

film [27, 28, 185]. For a junction wiring layer deposited at 0.9 nm/s we observe a RRR of 4.45, indicating good quality relative to the films we produce (see Figure 7.7a). From the scanning electron microscope image shown in Figure 7.8a we observe a short-range ordered microscopic grain structure in the regions where the wiring layer is deposited directly on the exposed sapphire substrate. A high-resolution top-down SEM image shown in Figure 7.8b reveals a network of thin grains with a visible hexagonal arrangement. Interestingly since the crystal structure within niobium grains is expected to be cubic [27] this suggests the long-range hexagonal order is a reflection of the C-plane sapphire substrate surface. The individual grains (distinctly larger than the 1-4 nm grains of the Pt and Pd film used to reduce charging in the image) are significantly elongated in one dimension. To get a sense of the grain size, we measure the short dimension of a grain (as shown in Figure 7.8b) for a number of grains visible in the image, and summarize the results in Figure 7.8c. By fitting to a normal distribution, we find an average grain width of 16.386 nm, with some skew towards longer widths. Notably we do not see the expected T_c reduction from this grain size [27, 28] since we find the measured T_c for this film is relatively close to the bulk value [66, 251]. This suggests

that the shortest dimension of the grains does not limit superconductor performance. We can instead extract an average grain area of $d \times l \simeq 163.8\ \text{nm}^2$ for our film, which corresponds to a effective grain size of $d_{\text{eff}} = \sqrt{ld} \simeq 40.47$ nm, from which we expect properties similar to bulk [27, 28]. Further investigation using X-ray diffraction or transmission electron microscopy [28, 185, 236] could reveal even more details about the microscopic properties of the niobium.

7.4 Lossy Plasma Etch Residues and Treatment

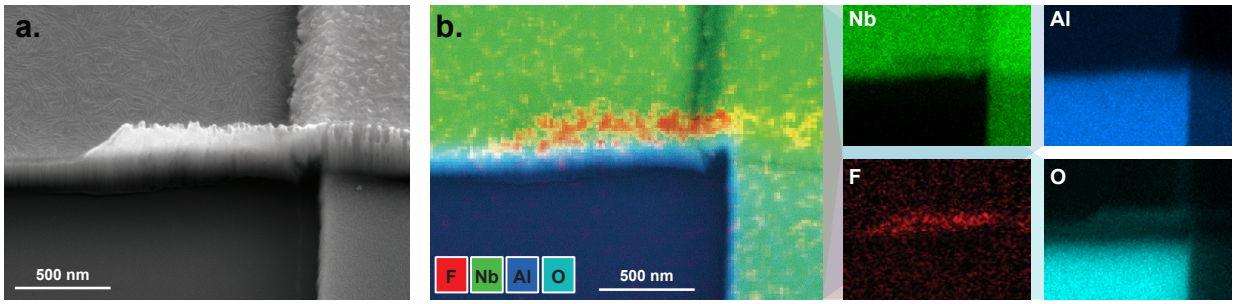

Figure 7.9: Etch residue chemical analysis. (a) Scanning electron micrograph of a plasma etch residue located on the wiring layer near a junction. (b) Composite Energy Dispersive Spectroscopy (EDS) image overlaid on the image in (a) showing normalized element density regions for F, Nb, Al, and O, with individual element density maps shown in their respective color on the right. Along with clear Nb and sapphire (Al_2O_3) regions, a high concentration of fluorine relative to the background is found in the residue region, suggesting the residue is composed of fluorinated polymers.

By virtue of size, the electric field concentration in a junction is orders of magnitude higher than in the qubit capacitor (or any planar structure such as the resonator capacitor), meaning the participation ratio [241] of the junction side surfaces will also be much higher. As such, our junction loss is likely still limited by the presence of lossy dielectrics formed on the sides of the junction, which for our design are primarily either spacer material, metal oxides, or residues left by the reactive ion etching process. As we cannot use more aggressive spacer [59] or oxide removal methods [236] without further risking the integrity of the aluminum junction barrier, we instead study the etch residues and discuss mitigation strategies.

Alongside the desired chemical and mechanical processes that remove niobium, reactive ion etching hosts a variety of simultaneous mechanisms that can grow material: etched material can either be re-deposited by sputtering, low-energy reaction products can re-adsorb onto exposed surfaces, and components in the plasma can react with exposed material [45]. The products of all of these mechanisms tend to be much more difficult to remove, so end up staying behind after the photoresist is dissolved, particularly on vertical walls not directly exposed to plasma bombardment during the etch. While the deposited material passivates the walls of the etched region during the etch and can produce high-aspect ratio features, for our junctions its critical to reduce any excess dielectrics, so we explore ways to understand and mitigate these residues in order to reduce loss.

In Figure 7.9a we show an example of a dielectric residue located on the side of a junction, which has not been removed throughout the entire fabrication process. This material must be formed during the third dry etch (Figure 7.1f) since it covers and extends off the sides of the Nb wiring and counterelectrode layers. The residues appear to be present on all vertical surfaces exposed by the etch, visible as striations on the junction sides. To determine the deposition mechanism for this residue, we probe the chemical composition of the residue using energy dispersive spectroscopy (EDS). A composite map of normalized element composition is overlaid on the same image of the residue in Figure 7.9b, with individual normalized element concentration maps shown to the right. As expected, we observe high Nb concentrations in the metal regions, and high aluminum and oxygen concentrations in the sapphire region, but more importantly we observe a significant concentration of fluorine in the residue (carbon is also observed in this region as well, but cannot be quantified due to high background carbon levels). This heavily suggests the residue is some kind of fluorocarbon polymer.

7.4.1 Optimizing Etch Conditions to Reduce Residue Formation

Regardless, the best way to remove the residues is to not form them in the first place, which is achieved by optimizing the etch plasma conditions. First, we remove obvious residue sources by ensuring the plasma chamber is thoroughly cleaned with oxygen, and no fluorinated vacuum oils are present in the system. We find that using gas constituents with low hydrogen and carbon content (eg. SF_6 or CF_4) significantly reduces the residue growth: in particular we find CHF_3 and C_4F_8 readily polymerize. However, we note that using too much SF_6 can lead to the incorporation of sulfur [223] into any exposed SiO_2, which forms an even more inert residue and should be avoided at all costs, as it presents an even greater source of loss than the fluorocarbon residues.

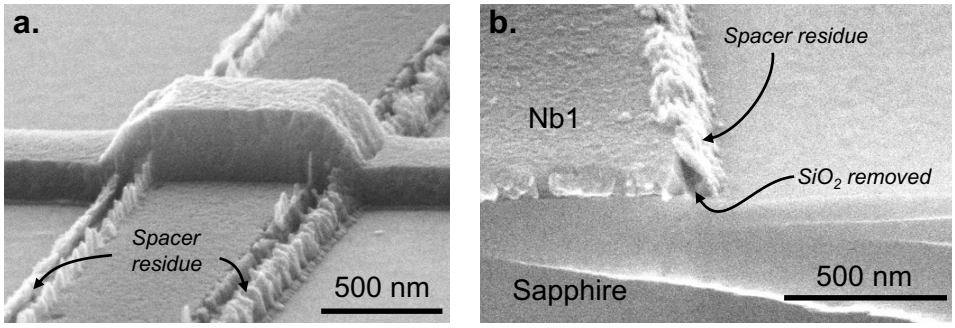

Figure 7.10: Lossy silicon oxide spacer residue, which is insoluble in HF of NH_4F.

The addition of O_2 in the plasma can also help increase the carbon-fluorine ratio of the plasma [157], but also increases resist etch rate [75] and may passivate exposed metal [197], which affects the etch profile. Using a low density plasma with a long mean free path for the radicals is key to increasing the etch rate and reducing re-deposition, as it increases the effective reactant and product temperature. Residue formation is also particularly sensitive to substrate temperature. With the substrate too cold, the reaction product temperature becomes low enough to allow recondensing, leading to increased fluorocarbon deposition. If the substrate is too hot, reactivity of the photoresist polymers is increased, promoting crosslinking, polymerization, and fluorination: thus good thermal contact between the substrate and

the carrier wafer is essential, as the high temperature plasma can otherwise significantly heat the substrate. Finally, we observe the residue formation accelerates when the insulating substrate is exposed (likely a result of screening charges focusing the plasma towards remaining metal), so we ensure the etch is stopped within 15 s of completion.

We find that the between samples, etch conditions can fluctuate significantly however. This is partially a result of our photoresist mounting method, which yields inconsistent thermalization to the carrier wafer (which is backside cooled with helium). As a result, inevitably and unpredictably the fluorocarbon residues can still form even under more ideal plasma conditions. To this end, it is vital to find another way to remove these fluorocarbon resides: in combination with an optimized etch, this should give a safeguard to fluctuating process parameters and ultimately provide the best surface conditions with consistency.

7.4.2 Residue Defluorination with Sodium and Potassium

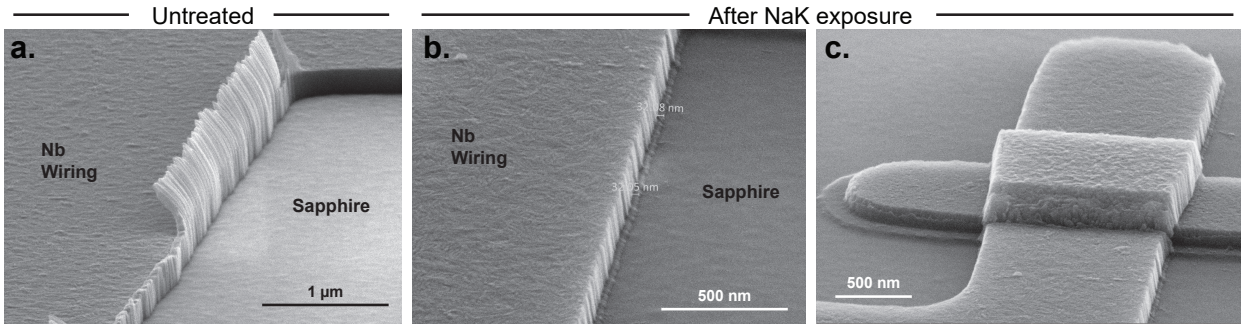

Figure 7.11: Etch residue NaK reactivity. (a) Scanning electron micrograph of a plasma etch residue on the edges of the wiring layer. A closer inspection of the bottom left reveals that the residue extends to cover the sides of the metal, even where the top crust has been mechanically removed. (b-c) The wiring layer and a junction from the same wafer imaged after a 15 min exposure to sodium-potassium amalgam (NaK) showing nearly complete removal of the etch residue.

Fluorocarbons are chemically inert and robust against most standard solvents, acids, or oxygen plasma, and the residues remain largely unaffected by these conditions. However, fluorocarbon polymers are susceptible to defluorination by strong alkali reductants such

as sodium-potassium amalgam (NaK) [46, 221]. In Figure 7.11a we show a device with particularly extensive residues covering and extending off the sides of the wiring layer. In an oxygen-free dry nitrogen glovebox, we immerse the sample surface in a sodium-potassium amalgam (NaK) for 15 min, rinse with tetrahydrofuran, move the sample into air, finish rinsing with acetone and isopropanol, then image the residues. In Figure 7.11b-c we observe that the residue material is largely removed: the overhanging features have been removed, as well as the material on metal sides, with the original extent of the residue (about 30 nm) apparent by the indentation left on the sapphire by the residue during the etch. This corroborates the hypothesis that these residues are composed of fluorocarbons, since the material could be removed upon treatment with NaK, wherein the amalgam cleaves the problematic C-F bonds and allows the remaining residues to become soluble in organic solvents.

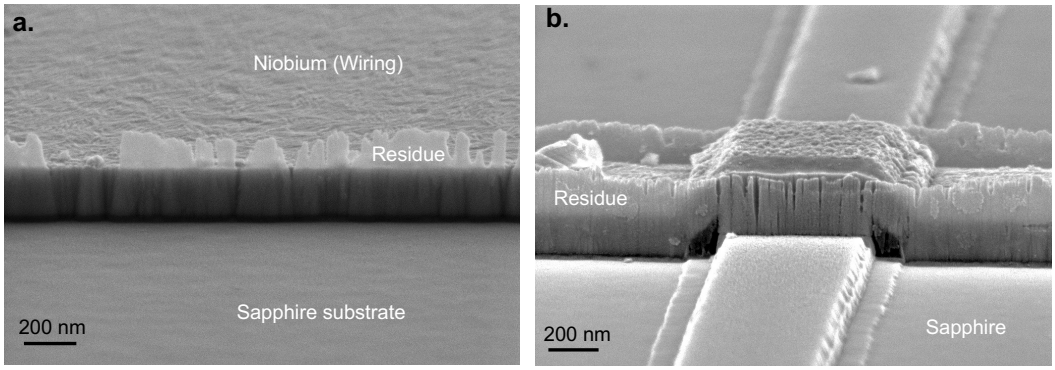

Figure 7.12: Residue treatment with potassium napthalenide K[Nap]. (a) Scanning electron micrograph of a wiring layer residue after immersion in a room temperature K[Nap] solution for 15 minutes. The residue remains, however is thinner and slightly damaged. (b) A junction from the same sample shows residue damage visible as vertical striations especially near the junction. While the residues are partially attacked by the K[Nap] solution, this treatment is not sufficient for full lossy residue removal.

While this NaK treatment appears promising on the microscopic scale, in practice the amalgam is difficult to keep clean, and leaves behind significant quantities of dust and salt deposits on the chip surface. A more practical method to post-clean any residues left behind by the etch might be to instead use a liquid solution with a high reducing potential such as

sodium napthalenide [46], commonly used as a surface treatment for PTFE. This method could preserve the microscopic cleanliness of the

7.4.3 Dissolving Fluorinated Organometallic Compounds

Ultimately we land on a commercially-available cleaning solution (Dupont EKC 265) specifically optimized to remove organometallic compounds. This alkanolamine solution consists of a polar solvent (H_2O), hydroxylamine, 2-(2-aminoethylamine)-ethanol (AEEA), and catechol [122, 123]. The etch mechanism for this mixture begins with the reduction of any organometallic compounds by the hydroxylamine, allowing the compounds to become more soluble in the water and alkanolamine solution [122, 123]. This also reduces any exposed metal oxides, which conveniently also removes contaminated surface metal oxides, which are significant sources of loss [185, 236]. It is unclear how this reactivity translates to silicon oxide in residual spacer material (or spacer residues).

The catechol, primarily a chelating agent or ligand, allows the now-soluble metal ions to form complexes and avoid precipitating out of solution. The hydroxylamine can also serve as a ligand in solution. Catechol also assists in metal protection and metal oxide solubilization as it can have a reducing potential[4] at neutral pH levels. The alkanolamine of choice, AEEA, is an alcohol amine with a relatively high boiling point, high flash point, and nearly nonexistent metal or substrate etch rates under standard process conditions. The two-carbon linkage is key towards reducing attack on metals and alloyed substrates while enhancing reactivity with organic and organometallic residues. Importantly, this mixture is only moderately acidic, with amines with hydroxyl groups both enhancing solubility and not increasing acidity greatly. This has the added benefit of reducing surface hydrogen content, which helps limit niobium hydride precipitation at cryogenic temperatures [69], which is known to adversely affect niobium superconducting properties [98] and coherence properties

4. At pH 7, catechol has a redox potential between 1.0 V to -2.0 V relative to the standard hydrogen electrode.

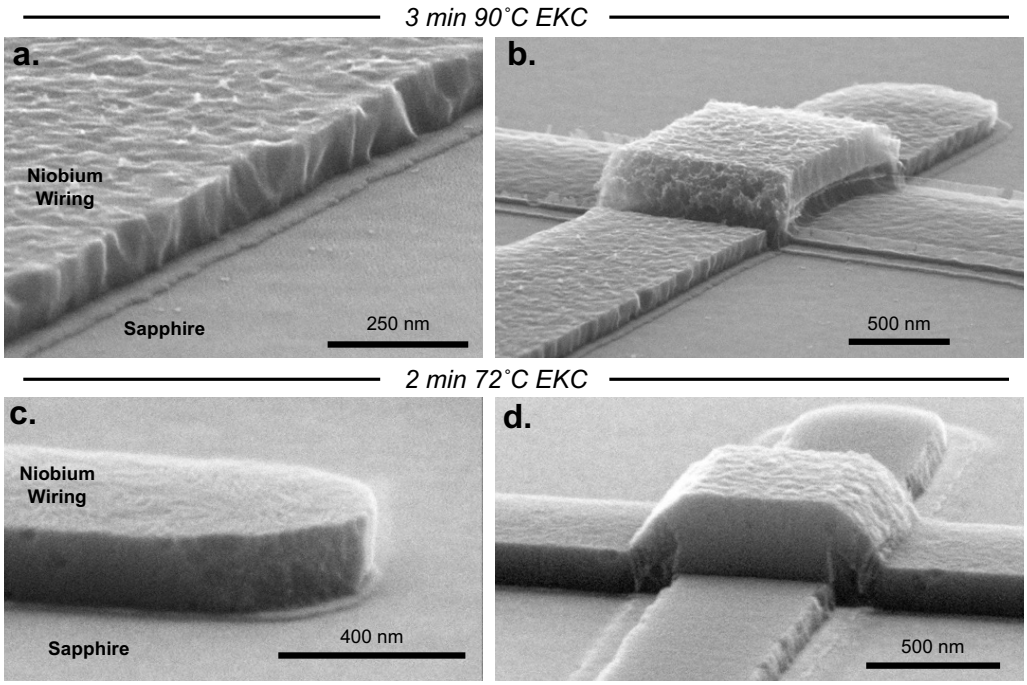

Figure 7.13: Residue treatment with EKC. **(a-b)** Finished junction treated with the EKC mixture nearly 20 degrees above the target etching temperature, showing significant metal attack (nearly 60 nm). Notably no sign of the fluorocarbon residues on the edges of the Nb wiring layer remain. Traces of material remains on the spacer-niobium interface, which warrants further study. **(c-d)** When treated between 70–75 °C, the metal etch rate is reduced to a reasonable level, while the organometallic residue is still efficiently removed. This leaves incredibly smooth and virtually residue-free surfaces on the junction.

[69].

As shown in Figure 7.13 the mixture is effective at efficiently removing fluorocarbon residue. At elevated temperatures (approaching the boiling point) the metal etch rate is increased, removing nearly 60 nm of niobium as shown in Figure 7.13a-b. In practice this consumes too much of the junction, leaving behind very little of the first niobium layer. At high temperatures the etch rates are less homogeneous, and the etch leaves behind ridges and pillars, producing a rougher surface. Interestingly enough while the niobium surface is significantly affected, the aluminum barrier appears to have a lower etch rate, visible in Figure 7.13b as some aluminum remaining in the junction extends out past the etched niobium. This may be useful for selectively etching away residues while preserving delicate

aluminum structures. We also note that a thin layer of material remains at the niobium-spacer interface: this may suggest that the spacer is either not completely removed by the ammonium fluoride wet etch, or perhaps forms an interfacial Nb-Si alloy [9] with different etch chemistry (and loss characteristics).

Reducing the etch temperature to 70–75 °C yields much more reasonable etch rates and surface profiles, as shown in Figure 7.13c-d. Only about 10 nm of niobium is consumed in the process, while the etch demonstrates selectivity towards the oranometallic residue materials (which can be up to 30 nm thick). These etch conditions completely remove the fluorocarbon residues, and leave a very smooth and lightly etched finish on the junction metal surfaces. Additionally, the mild surface etch ensures lossy surface interfaces and oxides [185, 241] are removed, which should improve surface loss characteristics.

7.5 Junction Area Dependence, Variation and Stability

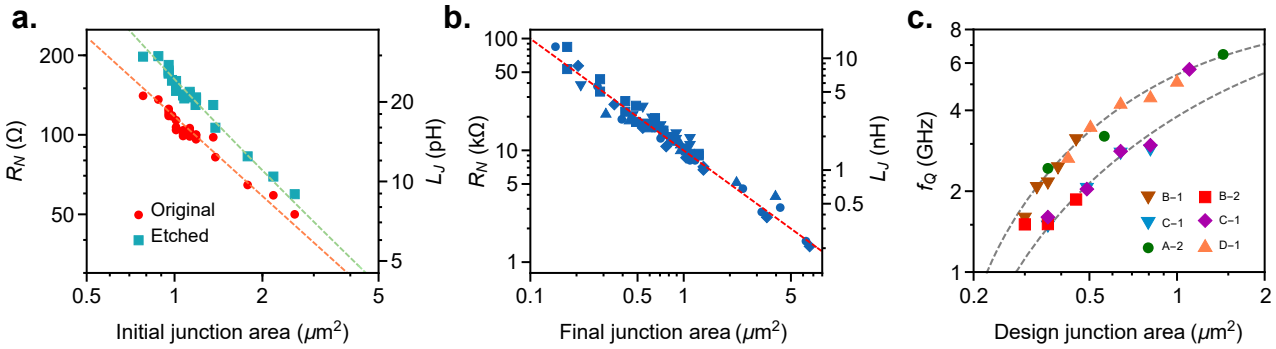

Figure 7.14: (a) Room temperature junction resistance and junction inductance plotted as a function of junction area (corrected for lithographic reduction). Original un-treated junction resistances are shown in red, and etched junctions in teal, with fits to an inverse relationship to area (dashed lines) yielding the original critical current density J_c and an etch dimension reduction of approximately 160 nm. (b) Junction resistances as a function of the final junction area with a inverse fit (dashed line) which gives the critical current density. For illustrative purposes we have shown PECVD junctions in (a) and HDPCVD junctions in (b). (c) To estimate reproducibility, spectroscopically measured qubit frequencies are plotted as a function of design junction area, labelled by wafer and cooldown. Expected values for the two different qubit capacitor designs (120 and 160 fF) are shown with dashed lines.

Having verified the relationship between the normal state resistance R_n, the critical current and the gap energy [10] (see Figure 7.6), we can use room temperature resistance measurements to efficiently predict cryogenic junction properties. In Figure 7.14a, we show room temperature junction resistance and junction inductance (calculated from resistance using the I_cR_N product), plotted as a function of junction area (corrected for lithographic reduction). The original un-treated (see Figure 7.1f) junction resistances are in good agreement with the expected inverse dependence on junction area, enabling us to fit the original critical current density. After etching the spacer (see Figure 7.1g) some of the aluminum is removed as well, and the resistance increases since the effective junction dimensions have shrunk. By fitting the etched junctions, we extract a dimension reduction of approximately 160 nm, which corresponds to about 80 nm of aluminum that gets removed by the etch. We note that this sets a practical limit on how small the junction can be before etch effects become more significant than lithographic definition of junction area.

Fitting junction resistances as a function of the final junction area (taking into account the dimension reductions) yields the true critical current density for the final junctions (Figure 7.14b). We repeat these measurements for wafers with different processing conditions to populate Figure 7.6b. A spread (typically between 5-10%) is noticeable in our junction resistance for a given junction area. While higher than typical niobium trilayer junction non-uniformity [37, 228], our junction variance can primarily be attributed to relatively large geometric deviations due to the limits of our lithographic resolution, which is compounded by fluctuations in the etch dynamics that determine the final junction area. This implies that our junction parameter spread could likely be reduced with higher resolution lithography methods and a more selective spacer removal technique. We test the functional limits of our junction reproducibility by measuring deviations of qubit frequencies across different chips from different wafers. In Figure 7.14c, we show spectroscopically measured qubit frequencies (determined by junction inductance) as a function of design qubit junction area for devices

with two different qubit capacitor designs. After determining the qubit capacitance and applying the estimated junction area reductions, we find the measured frequencies are self-consistent within 10 percent or so, even across separate wafers.

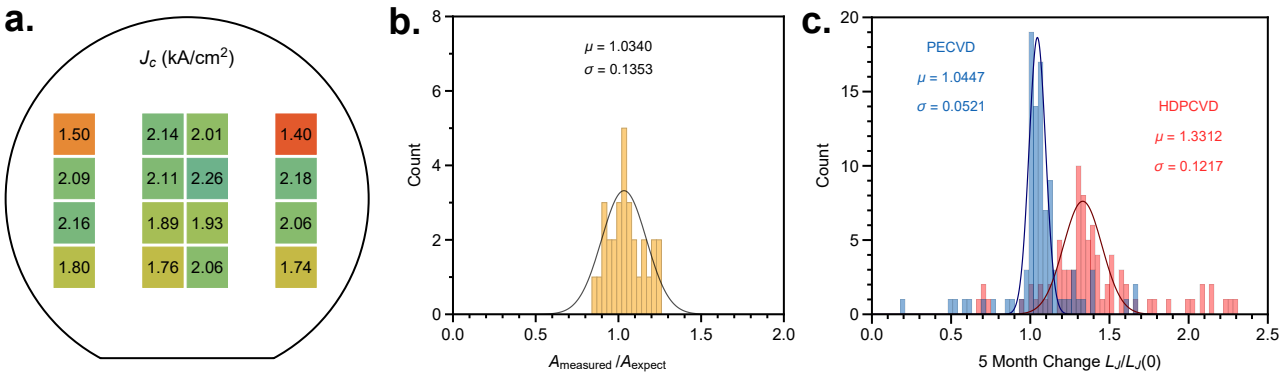

Figure 7.15: (a) Average junction critical current density on an individual chip measured across several chips across a wafer, with deviations from nominal values (2.088 kA/cm^2) highlighted with color. (b) Junction area measured with optical microscopy relative to the expected design area, highlighting the distribution of deviations resulting from lithography. (c) Long term stability of junctions measured by the relative change in Josephson inductance for 5 month old junctions relative to their original values. Notably the change in high temperature PECVD junctions is much lower than HDPCVD junctions.

We can investigate the variation of junction properties in more detail by repeating the measurements in Figure 7.14b for chips in different physical locations across a wafer. We plot the results by their original position in the wafer and summarize the results in Figure 7.15a. We find that the fitted critical current density fluctuates from chip to chip, consistent with the typical 5-10% junction variation observed in Figure 7.14b. Additionally, we observe a wafer-scale radial dependence in extracted critical current density, with noticeably lower values near the edge of the wafer. This is likely caused by a combination of dimension deviations from optical lithography and RIE etch rates, both of which have a wafer-scale radial dependence in our process. To estimate this lithographic dimension variation, we examine the statistics of measured junction area within a single chip relative to the expected area (with critical dimension bias taken into account), summarized in Figure 7.15b. Notably the measured areas are distributed with a standard deviation of 13.53% relative to the

expected area: when accounting for the 5% accuracy of the area measurement, the remaining spread accounts for a significant amount of the fluctuations observed in junction parameters. Thus we estimate that the dominant source of junction parameter variation is a result of dimension variation from optical lithography along with further dimension variation from fluctuation in etch dynamics. Process uniformity and lithographic dimension conformity are extensively studied topics [37, 162, 228], so we believe that applying these techniques or moving to higher-resolution lithography (stepper or electron-beam) [37, 109] could help decrease junction parameter variation.

Josephson junctions are known to change with age [162], so it is also important to investigate the long-term stability of junction parameters, particularly for our design which leaves the junction barrier exposed from the side. To this end, we re-measure junctions after 5 months of storage in air. In Figure 7.15c we show the relative change in calculated junction inductance

7.6 Junction Annealing Mechanism

The effect of process temperature is readily apparent when comparing the resulting critical current densities of junctions with PECVD spacers (deposited at 300 °C) and those with HDPCVD-grown spacers (90 °C). In Figure 7.6b, for the high temperature PECVD junctions, we observed an approximately 97.7% reduction in J_c. We investigate this effect in more detail by annealing finished low temperature (HDPCVD) junctions with initial $J_{c0} \sim 3\,\mathrm{kA\,cm^{-2}}$ in a dry Ar atmosphere, then re-measuring their critical current density.

In Figure 7.17a, we plot the annealed J_c as a percentage of the untreated J_{c0}, and confirm that the annealing effect activates above 200 °C, in agreement with [162]. In Figure 7.17b we show the critical current density of junctions annealed at 300 °C for various lengths of time. After about 20 min (the approximate time wafers spend at 300 °C during PECVD) we find that the current density reduction approaches the measured ratio between the PECVD and

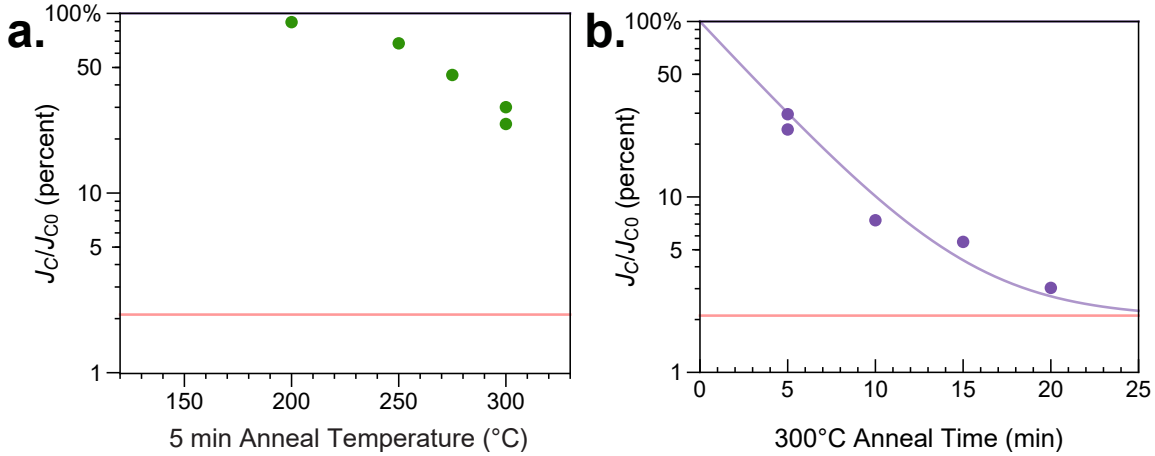

Figure 7.16: (a) HDPCVD Junction critical current density reduction after annealing for 5 min plotted as a function of anneal temperature showing activation at 250 °C. (b) Critical current density reduction as a function of anneal time at 300 °C, which approaches the factor of 50 reduction observed in the main text (red lines). The purple line represents an exponential fit saturating at the observed reduction factor.

HDPCVD junctions. This suggests the high-temperature process dynamically anneals the junction barrier, likely increasing mobility and in the oxide barrier which enables diffusion and reduces pinhole density [202]. Qualitatively, this process appears to be exponential in time, so we overlay a saturating exponential fit of the form $J_c/J_c^0 = (1-\alpha)e^{-t/\tau} + \alpha$, where α is the observed reduction factor, and obtain a critical time $\tau \approx 4\,\text{min}$. The observed annealing effect is consistent with the critical current densities measured in Ref. [81] which do not exceed 190 °C during the fabrication process.

7.6.1 High Jc with Low Temperature PECVD

With this in mind, our PECVD process could be modified to produce high critical current density junctions by either reducing the deposition temperature below 200 °C or to a lesser extent by limiting the time spent at elevated temperatures. This would allow for improved process stability by providing control over a wide range of critical current densities in a unified process, eliminating the need for switching between PECVD and HDPCVD deposition methods.

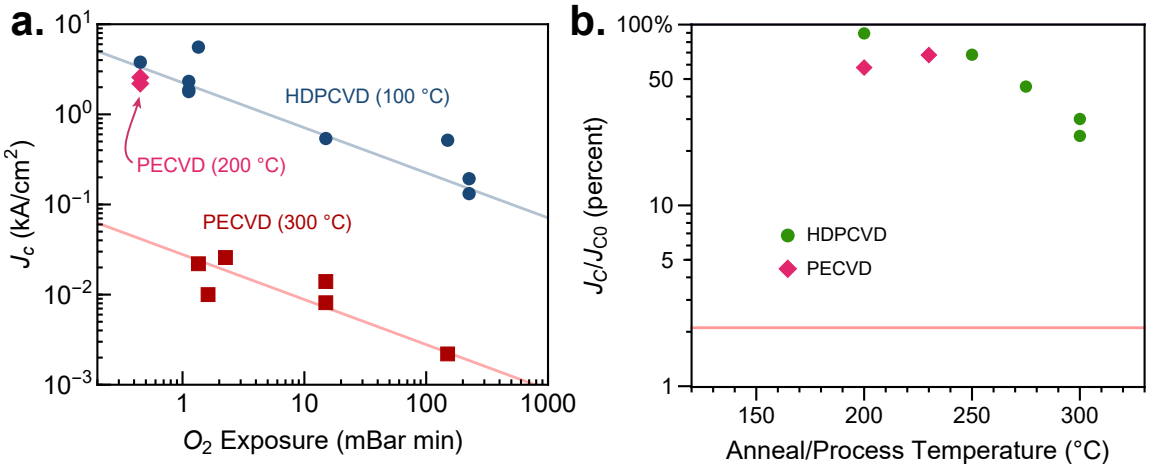

Figure 7.17: Low temperature PECVD junctions (primarily used in Chapter 8). (a) Junction critical density as a function of oxygen exposure (same as 7.6b) with the addition of low temperature PECVD junctions, which still have high critical current density. (b) Critical current density reduction as a function of anneal temperature, with the addition of the low-temperature PECVD junctions, which are only mildly annealed. The HDPCVD junctions were annealed for 5 minutes, while the junctions that went through PECVD spent approximately 15 minutes at the temperature.

7.7 Qubit Geometry and Experimental Setup

The qubit, readout resonator and other structures are formed in the same steps as the junction. We base our design on a qubit geometry [21] popular for its reduced radiation profile, a result of the cross-shaped coplanar qubit capacitor whose local electric dipole moments act to cancel each other out far away. In our case, the cross shape (typically used to implement qubit-qubit coupling or additional charge drives) isn't strictly necessary and a coplanar capacitor composed of any two-dimensional shape would work as well. We also make an effort to minimize coupling to lossy two-level systems in surface dielectrics by rounding sharp corners where possible in the geometry. This reduces electric field concentration at specific points in the capacitor, leaving a weaker and more homogenous electric field which should couple less strongly to individual two-level systems.

An example of our qubit geometry is shown in a composite microscope image on the top right of Figure 7.18, imaged after Etch 3 (Figure 7.1f). The niobium and un-etched

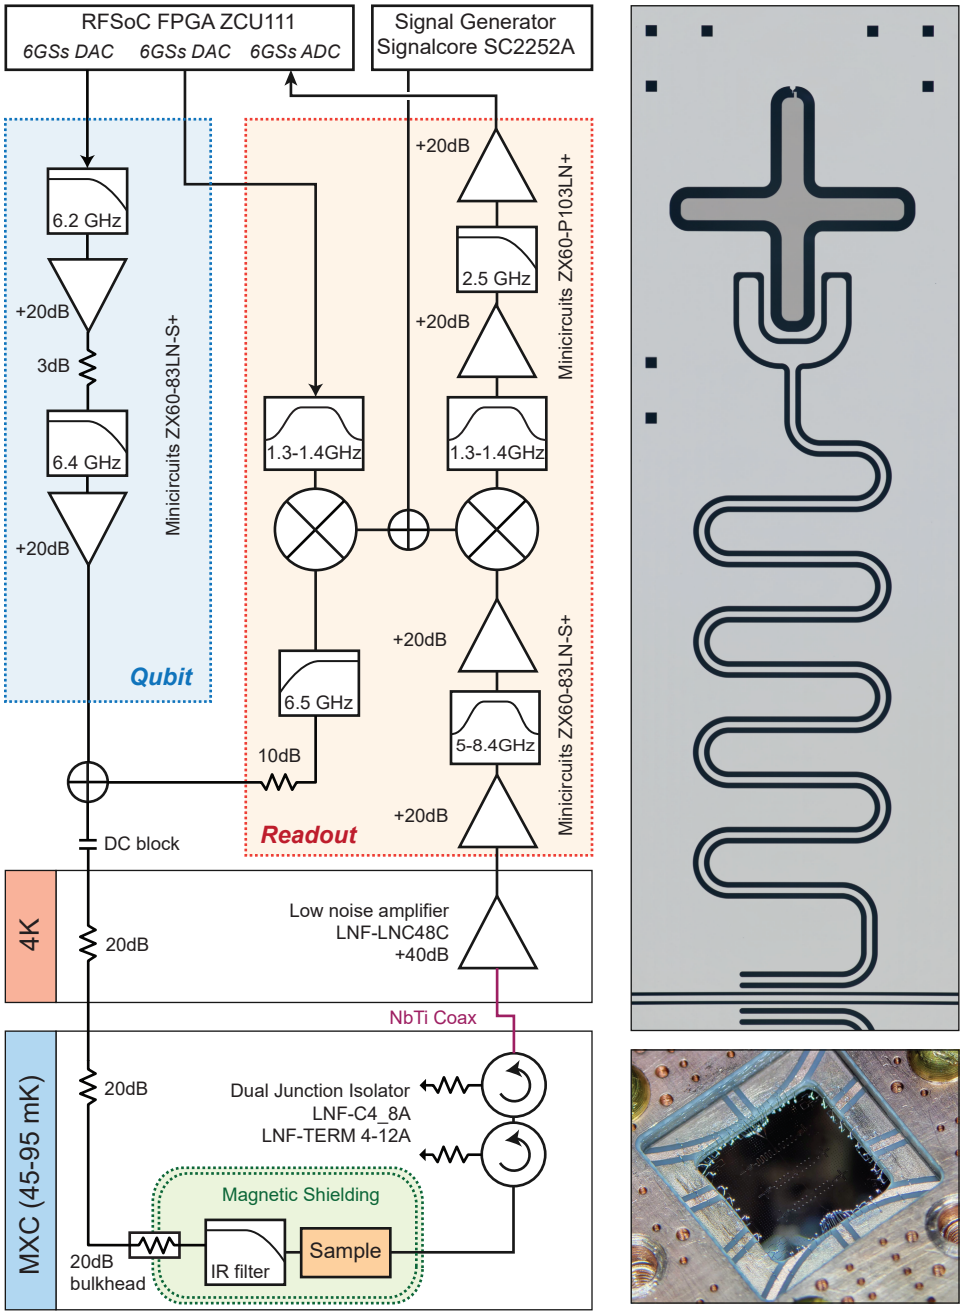

Figure 7.18: Schematic of the microwave measurement setup used for qubit characterization. Colored tabs show temperature stages inside the dilution refrigerator. A composite microscope image (top right) shows a single qubit and its readout resonator, coupled to a waveguide for measurement. A photograph (bottom right) shows the chip containing 6 qubits mounted in its copper circuit board.

aluminum have visibly different colors, allowing us to distinguish between the wiring layer and the base electrode. In our geometry, the qubit capacitor is formed with both layers, while the rest of the circuit and the majority of the chip (ground plane, readout resonator and coupling waveguides) is formed with just one layer. We find that the wiring layer readout resonators exhibit lower loss (See Appendix 7.8.2), so typically pick the wiring layer for the ground plane. However having measured devices with both configurations (majority wiring layer and majority base electrode), we don't find extreme differences in qubit properties, where the fields participate in both layers regardless of orientation. As an example, compare base-electrode ground plane wafer D with wiring ground plane wafer A in Figure 7.19a, whose qubit quality factors are similar.

The qubits are capacitvely coupled to a meandered quarter wave coplanar waveguide resonator, which is in turn inductively coupled to a transmission line for readout. For simplicity, we couple directly to the readout resonator without additional purcell filtering. Chips containing up to 6 qubits and resonators are mounted in a copper circuit board shown in the bottom right of Figure 7.18, which is in turn bolted to a copper post thermalizing the assembly to the base temperature of an Oxford Triton 200 dilution refrigerator with minimum mixing chamber temperatures between 45–95 mK. The mounted assembly is encased in two layers of Mu-metal magnetic shielding to reduce decoherence from stray magnetic fields, the qubits are isolated from microwave noise through an Eccosorb CR-110 high-frequency absorbing filter as 60 dB of cryogenic attenuation which keep the input noise close to the mixing chamber temperature. Transmitted microwave signals pass through two wideband circulators (isolating the qubits from microwave noise from the output side) into a low-loss superconducting NbTi coaxial cable, then are amplified by a low noise cryogenic amplifier followed by additional room temperature amplification.

Resonators and qubit transitions are characterized with single and two-tone spectroscopy using a Agilent E5071C network analyzer. For pulsed qubit measurements, we use a Quan-

tum Instrument Control Kit [213] based on the Xilinx RFSoC ZCU111 FPGA. Qubit pulses are directly synthesized by the FPGA, while measurement pulses are generated with a heterodyne conversion setup, as shown in Figure 7.18. With the spectral layout of each device determined, we select filter networks to minimize unwanted images and harmonics from the FPGA for both the qubit and readout pulses, with a broadband example configuration shown in Figure 7.18. The FPGA and carrier signal generator are clocked to a 10 MHz rubidium source for frequency stability.

7.8 Microwave Qubits

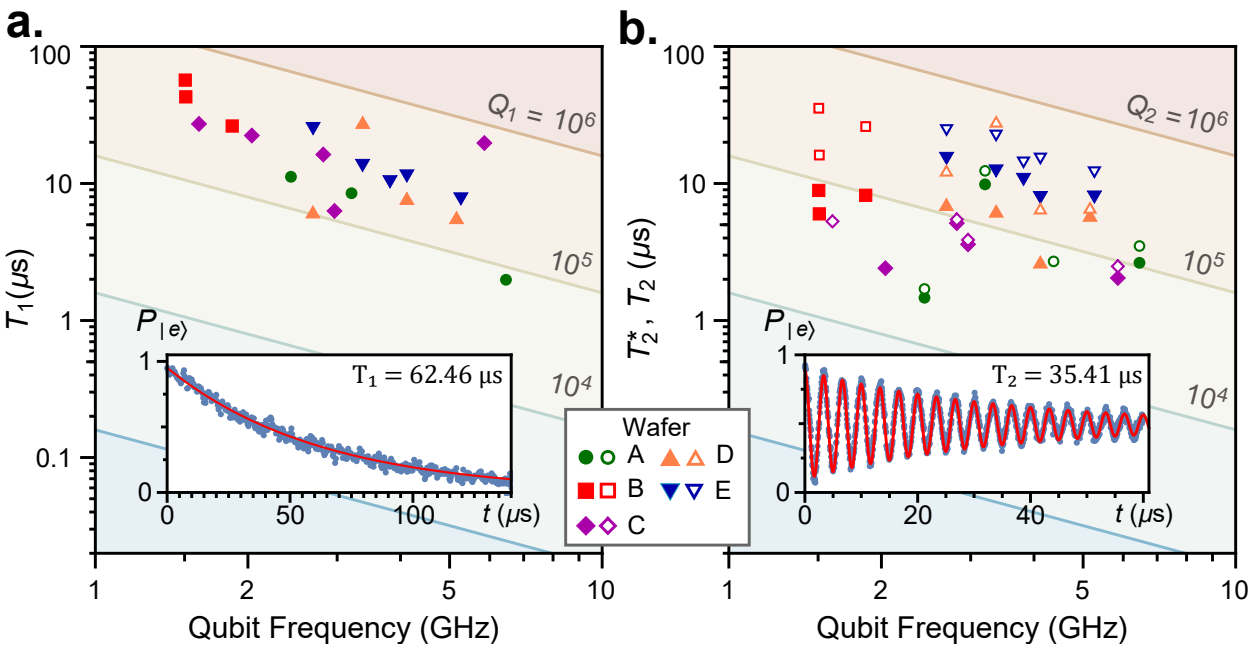

Figure 7.19: Qubit Properties. (a) Average qubit decay time T_1 extracted by fitting the exponential decay of excited state population in (b) plotted as a function of qubit frequency, grouped by wafer. Lines indicate qubit quality factor $Q_1 = \omega_q T_1$. We find an overall mean Q_1 of 2.57×10^5 with some wafer to wafer variation. (c) Ramsey dephasing time T_2^* (filled points) and Hahn-echo dephasing time T_2 (hollow points) extracted by fitting the exponential decay of oscillations in (d) as a function of qubit frequency. We find an overall average T_2^* and T_2 of 6.643 μs and 12.916 μs respectively. (e) Qubit quality factors as a function of their junction participation ratio plotted for devices in this chapter (reds) and in literature (blue, black, green). Lines and shaded confidence regions show $Q_1^{-1} = p_J/Q_J + p_0/Q_0$ as a guide to the eye.

With access to wide ranges of J_c, we can use PECVD-annealed junctions to realize qubit junctions with areas between 0.16–1.1 μm^2, which are large enough for optical lithography. We fabricate microwave transmon qubits [21, 113] with an otherwise standard geometry (see Appendix 7.7) capacitively coupled to a coplanar waveguide resonator for dispersive readout. The qubit capacitor, ground plane and readout resonator are defined on either the base electrode or wiring layer, so no additional fabrication steps are needed. Chips with several qubits and their readout resonators sharing a common microwave feedline are characterized at the base stage of a dilution refrigerator (45–95 mK). Using microwave spectroscopy [113] we verify our qubits have anharmonicities around 140 MHz, and couple to their readout resonators with a bare coupling strength $g = 30$–60 MHz.

7.8.1 Measuring Losses with Qubit Coherence

For superconducting qubits, the relaxation time and dephasing time are parameters of particular interest, as they dictate qubit limitations and act as sensitive probes for loss channels. We measure relaxation time by placing each qubit in its excited state and measuring it after time t: fitting the exponential decay gives the characteristic time T_1. We perform these measurements for each qubit and show averaged results as a function of qubit frequency in Figure 7.19a, finding $T_1 = 62.4\,\mu$s for our best device. To probe loss channels in detail we use the frequency-independent qubit quality factor $Q_1 = \omega_q T_1$, which we find for our devices is on average above 10^5: within an order of magnitude of recent aluminum qubits [111, 141, 182, 185, 198] and similar to readout resonator quality factors (see Appendix 7.8.2). We also perform a Ramsey experiment to measure the dephasing time T_2^*, and a Hahn-echo experiment to characterize the spin-echo dephasing time T_2. We find that T_2^* is within a factor of 2 of T_1, and particularly limited at lower qubit frequencies, where the environment two level system (TLS) temperature is higher and system $1/f$ noise is higher. The T_2 values, which decouple low frequency noise are noticeably higher, demonstrated in particular by the

qubits below 2 GHz from wafer B, which experience significant frequency drifts as they are in the charge sensitive regime [203].

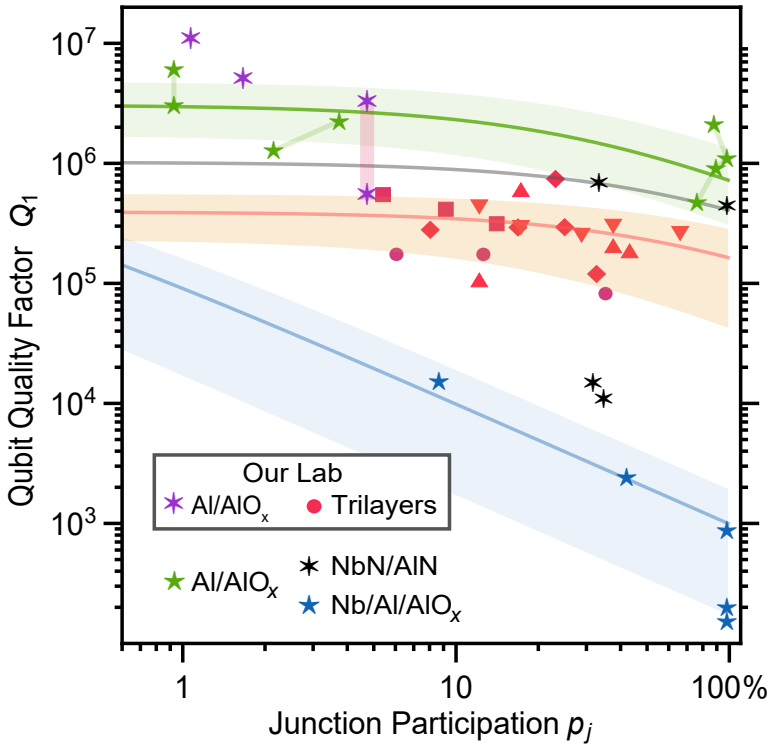

Figure 7.20: Qubit quality factors as a function of their junction participation ratio plotted for our trilayer qubits (reds) aluminum junction qubits from our lab (purple) and selected qubits in literature (blue, black, green). Lines and shaded confidence regions show $Q_1^{-1} = p_J/Q_J + p_0/Q_0$ as a guide to the eye.

Our qubits have relatively large junctions compared to typical qubits [141, 182, 185, 198] making them more sensitive to junction coherence properties. This allows us to study the junction loss contributions independent of other sources: in Figure 7.19c we use the junction participation ratio $p_j = c_j/C_\Sigma$ [141, 241] of our devices to examine the effects on Q_1. For our devices (red), we estimate an effective junction quality factor of 10^5: approximately 100 times greater than previous Nb/Al/AlO$_x$ qubits (blue) [60, 100, 129, 143, 177, 247, 261, 262], and much closer to epitaxial NbN junctions (black) [109, 166, 260] and modern aluminum-junction qubits (green) [141, 182, 185, 198]. Extrapolating to lower p_j values, we find our device loss is largely not limited by the junction, indicating that material refinements and

device engineering could further improve qubit performance.

7.8.2 Material Loss Probed by Resonator Quality Factor

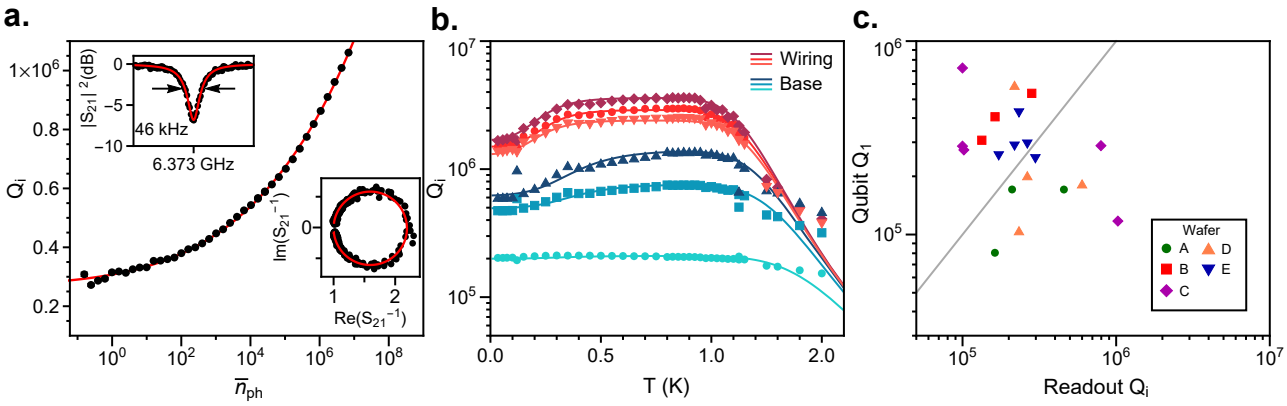

Figure 7.21: (a) Power dependence of the internal quality factor for a readout resonator ($Q_e = 2.6 \times 10^5$) with no qubit present. The red line is a fit to a model including loss from two-level systems (TLS). The insets show the lineshape and fits at an average photon occupation $\bar{n}_{\mathrm{ph}} \approx 0.96$. (b) Internal quality factor of resonators without qubits measured as a function of temperature. Solid lines are fits to a model including TLS loss and quasiparticle loss. The three red resonators are formed from the wiring layer, and the blue resonators from the base electrode. Measurements are taken at $\bar{n}_{\mathrm{ph}} \approx 10^4$ so some TLS loss is saturated. (c) Qubit quality factors Q_1 plotted as a function of their readout resonator quality factors Q_i (measured at $n_{\mathrm{ph}} < 1$). A grey line indicates a 1:1 relationship.

To compare qubit loss contributions from material sources with contributions from the junction itself, we measure quality factors for readout resonators subject to the same fabrication conditions, but with no qubits attached. A typical normalized transmission spectrum of a resonator taken at a low average photon number ($\bar{n}_{\mathrm{ph}} \approx 0.96$) is shown in the inset of Figure 7.21a. On resonance, we observe a dip in magnitude, which at low powers is described well by [106]:

$$S_{21} = 1 - \frac{Q}{Q_e^*} \frac{e^{i\phi}}{1 + 2iQ\frac{\omega - \omega_0}{\omega_0}} \tag{7.1}$$

where $Q^{-1} = Q_i^{-1} + \mathrm{Re}[Q_e^{-1}]$ and the coupling quality factor $Q_e = Q_e^* e^{-i\phi}$ has undergone a complex rotation ϕ due to minor impedance mismatches. We plot fitted internal quality factors in Figure 7.21a, finding that Q_i increases with power. This behavior is entirely

captured by a power dependent saturation mechanism [242], suggesting the dominant loss mechanism in the resonators arises from coupling to two-level systems.

We futher investigate limits on the resonator loss by using increased temperatures to further saturate the two-level systems. In Figure 7.21b we plot Q_i measured at $\bar{n}_{\mathrm{ph}} \approx 10^4$ as a function of temperature (grouped by fabrication layer), with solid lines corresponding to a model of the form

$$Q_i(T)^{-1} = Q_{\mathrm{other}}^{-1} + Q_{\mathrm{TLS}}(T)^{-1} + Q_\sigma(T)^{-1} \tag{7.2}$$

where Q_{TLS} is the saturating loss mechanism from two-level systems [242], Q_{other} is a temperature independent upper bound arising from other sources of loss, and the conduction loss Q_σ is given by [145]:

$$Q_\sigma(T) = \frac{1}{\alpha}\frac{\sigma_2(T, T_c)}{\sigma_1(T, T_c)} \tag{7.3}$$

where σ_1 and σ_2 are the real and imaginary parts respectively of the complex surface impedance, calculated by numerically integrating the Mattis-Bardeen equations for σ_1/σ_n and σ_2/σ_n [145]. T_c is constrained to the values measured in Appendix 7.3, and α is used as a fit parameter.

Comparing resonators formed during different steps in the fabrication process, we observe that resonators made from the wiring layer exhibit consistently higher quality factors, while resonators from the base layer are lossier and much more variable. Since the sides of the base layer have been exposed to more fabrication steps than the wiring layer, the surface niobium of this layer has a much longer chance to oxidize, and has the additional potential to host lossy dielectrics from un-removed spacer material. Thus, while we have improved losses in the wiring layer to about $Q_{\mathrm{TLS}} \sim 0.9 \times 10^6$ by reducing fluorocarbon formation, our devices are still loss-limited to approximately 2×10^5 by two-level systems in the surfaces of the base electrode.

To investigate the relationship between qubit and readout resonator decoherence, we also measure quality factors of the readout resonator for each qubit. At single-photon powers, the readout resonator is maximally susceptible to material-based loss from two-level systems in its surface, but due to the hybridization of its electric field with the qubit mode will also interact with the materials in the qubit. In Figure 7.21c we compare qubit quality factors Q_1 with the single-photon readout quality factor Q_i for each of the devices from Figure 7.19. On average, we observe that the two quality factors are close to a one to one relationship (as indicated by the grey line), with device variations within a factor of 3 or so. While a direct correlation between the two cannot be extracted from this data, this is to be expected for loss dominated by inhomogeneous material defect distributions between the resonator and qubit. Nevertheless, the similarity of the two quality factors leads us to conclude the qubit and resonator are likely limited by similar decoherence mechanisms.

7.8.3 Detailed Model of Junction Losses

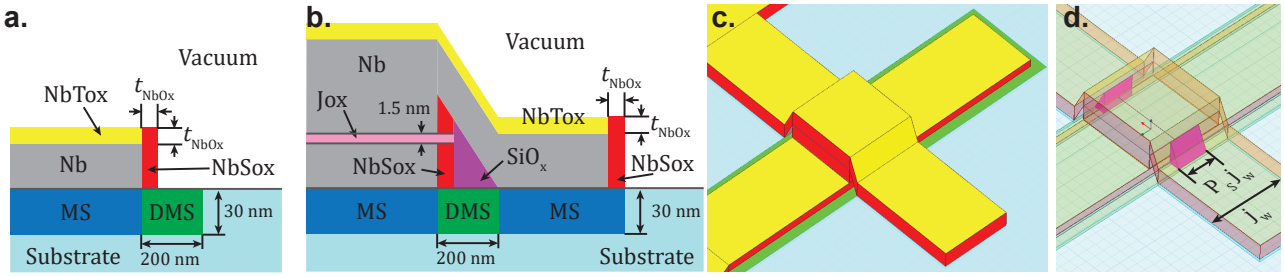

Figure 7.22: Junction loss regions (a) Cartoon showing regions defined for a resonator made with the first layer, with dimensions exaggerated. Niobium oxide (metal-air interface) is separated into top oxide (Tox) and side oxide (Sox) regions. For a wiring layer resonator, the dirty substrate region (DMS) is merged with the substrate layer. (b) Cartoon showing regions for a junction, which adds the junction barrier region (Jox) and the spacer region (SiOx). (c) Three dimensional rendering of the junction with realistic dimensions. Simulated regions are colored in the same way as in parts (a-b). (d) Transparent rendering of the junction visualizing the spacer remaining percentage P_S relative to the junction width j_w.

In the main text along with Appendix 7.8.2 we established that our junction quality factor ($Q_J \approx 10^5$) is similar to the single-photon quality factors of bare resonators, for which

we measured an average of 2.6×10^5 for base layer and 1.04×10^6 for the wiring layer. The fact that these loss rates are comparable suggest that some part of the qubit decoherence arises from same material losses probed by the resonators. To investigate the origins of these loss channels in more detail and elucidate important pathways for further improvement, we use finite element method simulations (Ansys HFSS) to examine the energy participation ratios [241] of different regions and interfaces in the junction.

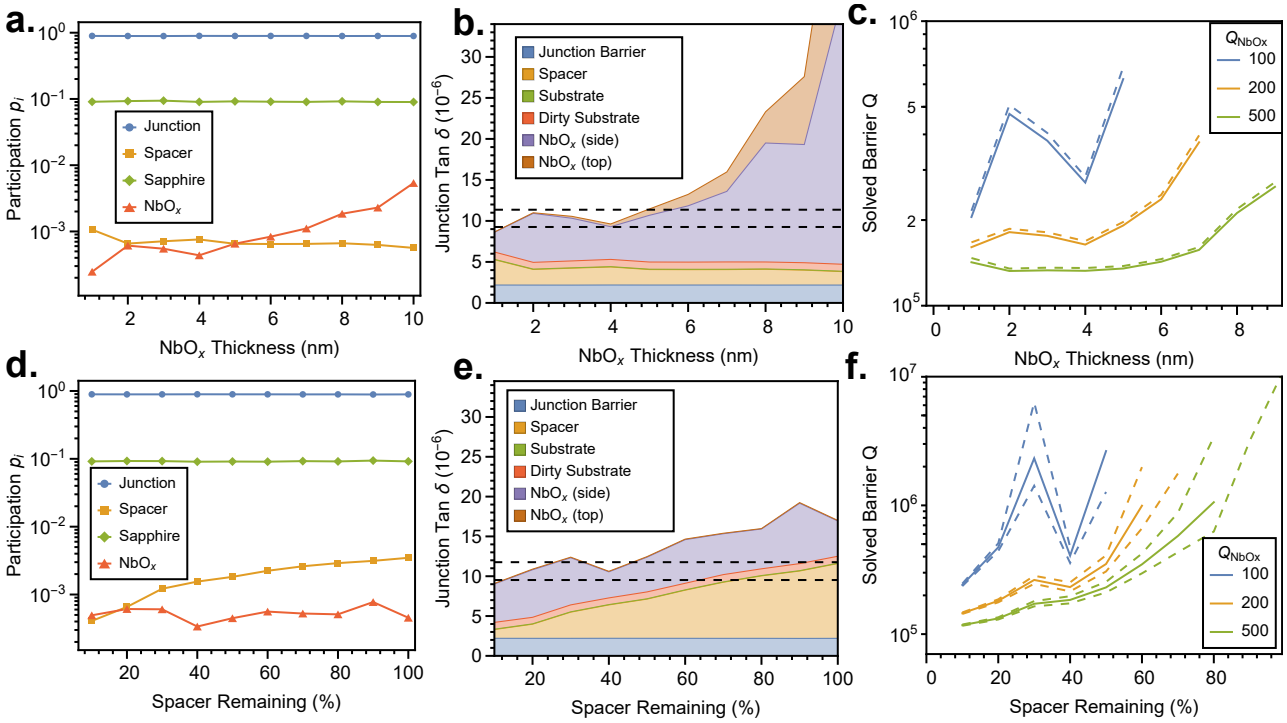

Figure 7.23: Junction losses by region. (a) Participation ratios of the primary lossy materials in the junction, plotted as a function of niobium oxide thickness t_{NbO_x}. As expected the niobium oxide participation ratio increases as the layer gets thicker. (b) Junction loss tangent expressed as visual sum of losses from various materials in the junction with assumed loss tangents, plotted as a function of niobium oxide thickness. For thicker oxide layers (eg. those used in anodization processes) niobium oxide loss dominates the junction loss. The junction loss calculated from Fig. 7.19c is shown in black dashed lines. (c) We can also solve for the barrier quality factor based on the junction quality factor and the calculated participation ratios for varying material quality factors. Solid and dashed lines correspond to a SiO_2 loss tangent of $\tan\delta = 2.7 \times 10^{-3}$ and 2.9×10^{-3} respectively. In (d-f) we repeat parts (a-c) but measure the effect of partially un-removed spacer material expressed as a fraction P_S of the junction width. We find that residual spacer material contributes a significant amount of loss. For both sets of simulations, the unswept variable is set to nominal values of $t_{\mathrm{NbOx}} = 2$ nm and $P_S = 0.2$.

Figure 7.22 illustrates the material regions studied. Similar to other surface participation studies [54, 241, 242] we consider the metal-substrate interface regions, which we further subdivide into the metal-substrate interface (MS) and the dirty metal-substrate region (DMS) which may contain some remaining spacer material. As the etched sapphire surface and bulk loss are both expected to be minimal [188] we combine the substrate-air (SA) interface with the substrate region for participation calculations. Based on the electric field density we choose the thickness of the surface regions of the substrate to be 30 nm: adjusting this thickness will simply rescale the effective participation and loss of the metal substrate regions. The bulk of the loss is expected to lie in the amorphous oxide dielectric regions of the junction. These can be separated into the aluminum oxide comprising the junction barrier (Jox) which we expect to be 1–2 nm [162], and the niobium oxide, which we further separate into a top oxide layer (NbTox) and side oxide (NbSox): the latter of these should be substantially lossier since it may contain fluorocarbons after exposure to the fluorine plasma (see Appendix 7.4). Finally we also consider the possibility of incomplete spacer removal and also include a portion of SiO_2 to model the spacer, as shown in Fig. 7.22d. With our imaging methods, we are unable to determine the amount or layout of the residual spacer material, so for simplicity we approximate the region as a uniform percentage of the original spacer volume $P_S \leq 1$.

Integrating the simulated electric fields in the junction geometry determines the participation ratios in each dielectric region [241]. The surface niobium oxide thickness t_{NbO_x} typically ranges between 1–5 nm [9, 236] but can vary to a greater depending on process conditions, so in Fig. 7.23a we study the participation ratios as a function of t_{NbO_x}. From this we conclude that most of the energy is stored in the junction barrier, followed by the sapphire regions, with the niobium oxide and spacer regions contributing less than one percent. As expected, the niobium oxide participation increases with t_{NbO_x}, however importantly the energy participation is dominated by the side oxide ($p_{\mathrm{Sox}} \gg p_{\mathrm{Tox}}$), especially

for thinner values of t_{NbO_x}. In the same manner, we can also simulate the participation ratios of a section of coplanar waveguide (the cross section of which will be the same as that shown in Fig. 7.23a). Comparing resonators fabricated from the first and wiring layer effectively amounts to the presence of the dirty substrate region (DMS) in our model. For the resonator geometry, we find this region has an average participation $p_{\mathrm{DMS}} = 0.366\%$, largely independent of oxide thickness (for which $p_{\mathrm{NbO}_x} \sim 0.005\%$, similar to Ref. [236]). Based on the single-photon resonator quality factors from both layers in Appendix 7.8.2, we solve for the material quality factors as a function of the niobium oxide quality factor, which is typically $Q_{\mathrm{NbO}_x} = 1/\tan\delta_{\mathrm{NbO}_x} \simeq 100$ [9, 236]. From this we conclude that $Q_{\mathrm{DMS}} \simeq 1.4 \times 10^3$ and $Q_{\mathrm{Sapphire}} \simeq 1.8 \times 10^6$, which is reasonably consistent for averaged bulk and surface measurements of sapphire loss [188] and closer to the loss values found in silicon oxide ($\tan\delta_{\mathrm{SiO}_2} \simeq 2.8 \times 10^{-3}$[176]) in the DMS region.

Combining the expected material losses with the calculated participation ratios, we can express the junction quality factor as a sum of loss contributions from each region to identify dominant sources of decoherence.

$$\tan\delta_J = \frac{1}{Q_J} = \sum_x \frac{p_x}{Q_x} = \sum_x p_x \tan\delta_x \tag{7.4}$$

We summarize the contributions for each material as a function of $t_{\mathrm{NbO}_x} = 2$ nm in Fig. 7.23b along with the average junction quality factors measured in the main text. Despite the high barrier participation, we find the dominant loss contribution is from the niobium oxide: specifically that on the sides of the metal (NbSOx) which is also more likely to be impacted by the plasma etch chemistry. For simplicity, we have determined the junction barrier quality factor from the average junction quality factor Q_J by assuming that $t_{\mathrm{NbO}_x} = 2$ nm [236] and conservatively estimating $P_S = 0.2$: this yields a junction barrier oxide quality factor of $Q_{\mathrm{Jox}} \simeq 4.7 \times 10^5$. We can repeat this calculation with varying conditions to estimate the effective barrier quality, as shown in Fig. 7.23c which suggests that the barrier Q may

exceed our estimate if the oxide thickness is thicker than 2 nm, or may be lower if the niobium oxide quality factor is in fact higher than expected. As the predicted junction loss cannot exceed the measured value, assuming standard oxide loss $Q_{\mathrm{NbO}_x} = 100$ [9, 236] implies $t_{\mathrm{NbO}_x} < 5$ nm, which helps validate the previous assumptions.

We can also perform a similar set of calculations for the remaining spacer amount P_S, summarized in Fig. 7.23d-f. As expected we observe the silicon participation ratio p_{SiO_2} increases with larger spacer volume. When more than half of the spacer remains, we estimate that the silicon oxide comprises the dominant source of loss in the junction. Similar to the niobium oxide thickness, the spacer percentage also affects the estimated junction barrier Q as shown in Fig. 7.23f. This also suggests an upper bound for the residual spacer percentage $P_S \lesssim 0.5$, indicating the final wet etch is at least somewhat successful in removing spacer material under the wiring layer.

Thus we have identified several key areas where junction loss could be further improved. As discussed in the main text, reducing the amount of lossy dielectrics (particularly the spacer material and niobium oxide) is key to increasing junction loss, as highlighted in Fig. 7.23b,e. While we have taken steps to reduce the volume of both niobium oxides and spacer material, further improvements on both these fronts could help improve junction quality. Further reduction of junction loss may require addressing losses in the dirty substrate region with improved cleaning methods. However from Fig. 7.23a,c we conclude the junction is most sensitive to the quality of the barrier dielectric. In this regard, atomically uniform barriers such as AlN deposited with molecular beam epitaxy in NbN junctions [109, 260] may provide even better performance.

7.8.4 Temperature dependence of decoherence

Next we investigate our qubits at increased operating temperatures, shown in Figure 7.24. We observe a mild decrease in T_1 with temperature above 160 mK, consistent with heat-

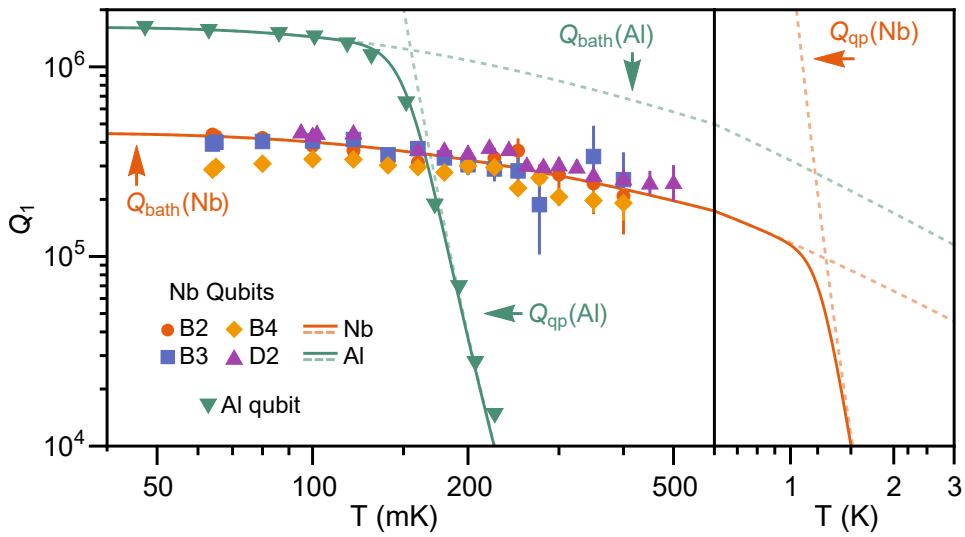

Figure 7.24: Qubit quality factors from wafers B, D as a function of temperature. A mild decrease is observed at higher temperatures consistent with the system bath temperature Q_{bath}, however lifetimes are virtually unaffected by quasiparticles Q_{qp} (red lines). We also plot quality factors of an Al junction qubit, whose performance is noticeably limited by quasiparticles above 160 mK (green lines), whereas the Nb junction wouldn't see an effect until 1.6 K.

ing from the environment bath [129], but importantly don't see the drastic temperature dependence expected for quasiparticle-induced loss [39, 40, 145], in line with expectations for niobium. The advantage of higher-temperature junctions is apparent when comparing our qubit performance to an aluminum counterpart: above 160 mK, the aluminum qubit is quickly overwhelmed with quasiparticle-induced decoherence, whereas our devices largely retain their properties.

7.9 Conclusion

We have described a Nb/Al/AlO_x/Al/Nb trilayer fabrication method demonstrating a 100-fold improvement in junction loss at the single-photon level. By removing lossy dielectric materials wherever possible, we use our low current density junction process to fabricate microwave transmon qubits using I-line photolithography demonstrating qubit quality factors within an order of magnitude of recent aluminum devices. Our qubits have relatively high

junction participation ratios, which could either be reduced to improve coherence through material optimization [236], or exploited further to significantly reduce qubit size [141, 268]. Together with this device footprint flexibility, our all-optical qubit process opens the door to large-scale direct integration of scalable quantum processors with digital superconducting logic [126, 131, 147]. Niobium's higher energy gap significantly reduces sensitivity to quasiparticles for our junctions compared to aluminum analogues, allowing operation at much higher frequencies, and resulting in far less decoherence above 160 mK where conventional qubit properties deteriorate. Combined with their low loss, these properties make our trilayer junctions a promising candidate for quantum architectures with lower cooling power requirements, hybrid qubit systems requiring elevated temperatures, and enable new possibilities for nonlinear elements at millimeter-wave frequencies [12, 119], paving the way for higher temperature, higher frequency quantum devices.

CHAPTER 8

K BAND QUBITS

In the previous chapter we explored how niobium trilayer junction qubits had reduced sensitivity to quasiparticles. Increasing qubit operating temperature requires a two-pronged approach addressing both quasiparticle loss [144], which is determined by the superconducting transition temperature of the materials that form the qubit, as well as decoherence from thermal photons. As we verified previously, Niobium's higher superconducting energy gap compared with aluminum's significantly increases the energy required to generate quasiparticles, making it a promising material for higher temperature qubits.

Addressing increased environment noise at higher temperature presents two challenges. First, thermal photons at the qubit frequency result in qubit heating, which requires active reset [84, 135] or extensive dissipation engineering [233, 245] to artificially cool the qubit to its ground state before experiments begin. Second, thermal photons at the readout resonator frequency, which cannot be easily cooled, increase qubit dephasing [43]. Increasing the system frequencies addresses both of these issues, reducing sensitivity to thermal fluctuations. Newly developed components for high-speed communications [89] enable us to directly scale traditional microwave transmon design to the K band (18-26 GHz), and explore coherence properties in a new regime.

8.1 K Band Qubit Design

8.1.1 High Frequency Trilayer Junction

The key element at the heart of our transmon qubit is a niobium trilayer Josephson junction with high critical current density (J_c), shown in Figure 8.1a. We use the self-aligned fabrication process desribed in Chapter 7 to form a Nb/Al/AlO_x/Al/Nb Josephson junction

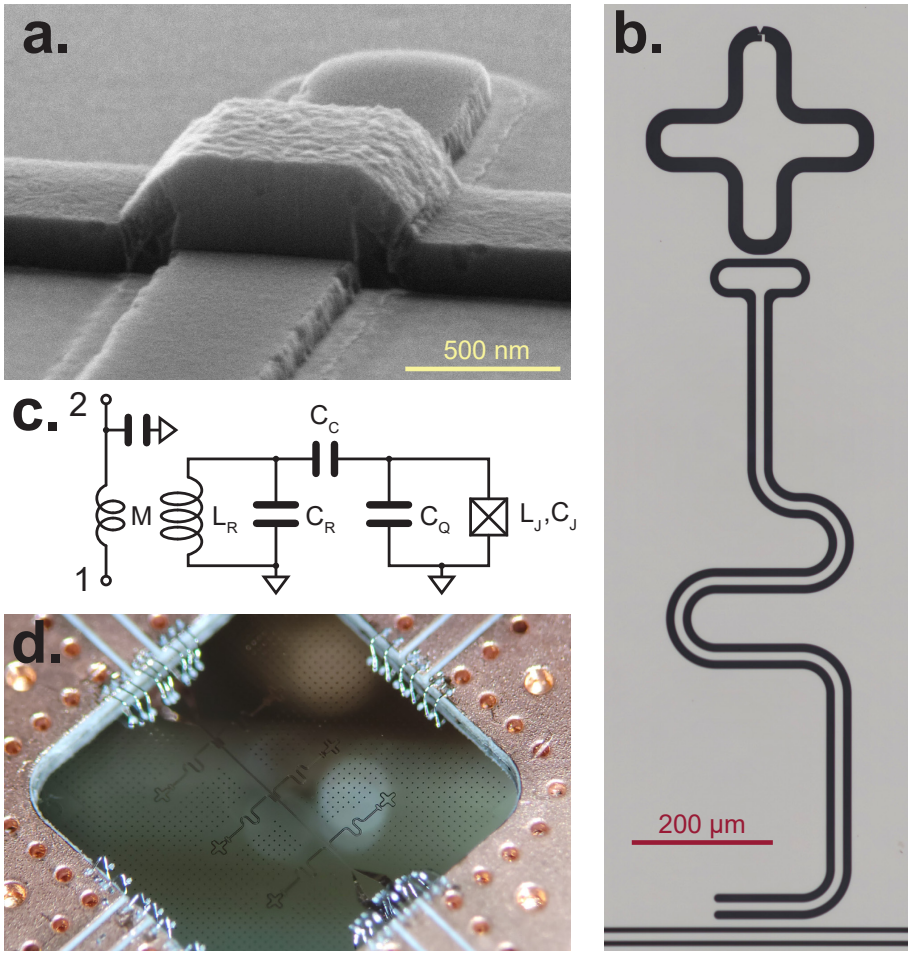

Figure 8.1: Qubit geometry. **a)** A scanning electron micrograph of a low-loss niobium trilayer junction at the core of the qubit. **b)** A micrograph of the qubit and readout resonator geometry, with the junction location marked at the top. **c)** Equivalent circuit of the qubit and readout resonator coupled inductively to a transmission line. **d)** Photograph of a chip containing six qubits mounted in a low-loss K band circuit board.

on a C-plane sapphire substrate[1], with two main changes from Ref. [13]. First, we use low-temperature PECVD for spacer growth [81]: keeping the wafer below 230 °C during the oxide growth creates a spacer oxide with lower loss than other methods, such as high-density PECVD[13] and preserves a high J_c density, needed for higher-frequency junctions.

1. HEMEX sapphire grown by heat-exchange method: this is actually very important

8.1.2 Post Fabrication Residue Removal

The second major improvement is a etch specifically targeting the fluorocarbon residues, investigated in detail in Chapter 7. Immediately prior to measurement we use a 72 °C solution of polar solvent (H2O), hydroxylamine, 2-(2-aminoethylamine)-ethanol (AEEA), and catechol [2] to remove lossy plasma etch residues along with a thin layer of the oxidized niobium surface, known to contain lossy oxides [185, 236]. The etching mechanism for this mixture is believed to begin with the reduction of any organometallic compounds or metal oxides by the hydroxylamine to first become more soluble in the water and alkanolamine solution [122, 123]. The catechol, primarily a chelating agent (or ligand), allows the dissolved metal ions to form complexes and not precipitate out of solution. The hydroxylamine can also serve as a ligand in solution. Catechol also assists in metal/alloy protection and metal oxide solubilization due to its redox potential[3]. The alkanolamine (AEEA), is an alcohol amine with a relatively high boiling point, high flash point, and a relatively low (or non-existent) etch rate under standard process conditions for metallic substrates or layers. The two-carbon linkage is key towards reducing attack on metals and alloyed substrates and enhances residue attack. Importantly, the etch mixture is only moderately acidic, with amines with hydroxyl groups both enhancing solubility without greatly increasing acidity greatly[4]. In summary, this treatment leaves a smooth metal surface with minimal defects and residues, visible in the smooth details of the junction in Figure 8.1a.

8.1.3 Qubit Circuit Design

For these PECVD junctions, we measure a J_c density between 2.19–2.57 kA/cm^2 (see Chapter 7). This is much higher than aluminum junctions (typically 0.3 kA/cm^2), allowing us

2. The solution is commerically available as Dupont EKC 265

3. The strong redox potential may actually reduce exposed surface niobium as well.

4. This has been shown to be beneficial for reducing niobium hydride precipitation, which is known to induce microwave losses.

to design qubits in the K frequency band using junction finger dimensions between 0.4-0.8 μm, realizable with optical lithography. Our qubit geometry, shown in Figure 8.1b, is qualitatively similar to conventional microwave transmon qubits [21, 113], with the primary difference that every feature is slightly smaller. The rounded cross-shaped qubit capacitor C_Q is capacitively coupled to a short quarter-wave coplanar waveguide resonator used for dispersive measurements. The other end of the resonator is inductively coupled (κ =2-9 MHz) to a common feedline. The system can be modelled by the circuit in Figure 8.1c, where the qubit frequency is determined by the Josephson inductance L_J and total shunt capacitance C_Σ. The qubit capacitor, ground plane and readout resonator are defined along with the junction, so no additional fabrication steps are needed.

8.2 K Band Packaging

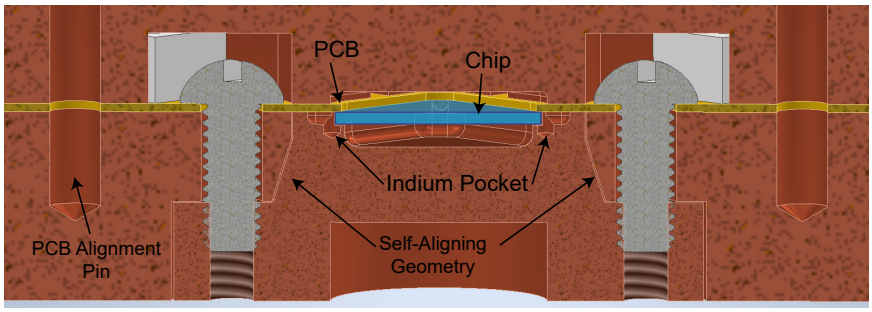

Figure 8.2: Cross section view of the K band packaging, showing the chip (blue) secured to the high-frequency circuit board (gold) by the copper enclosure.

Chips with several qubits and their readout resonators sharing a common microwave feedline are mounted in a K-band package (shown in Figure 8.1d), which is carefully engineered for low-loss operation up to 30 GHz. Inspired by microwave solutions [90], the package aims to shield the qubits from decoherence from external fields, thermalize the sapphire chip containing the qubits, while adding minimal unwanted resonant modes, and efficiently routing signals to and from the chip with minimal dissipation and reflection . We accom-

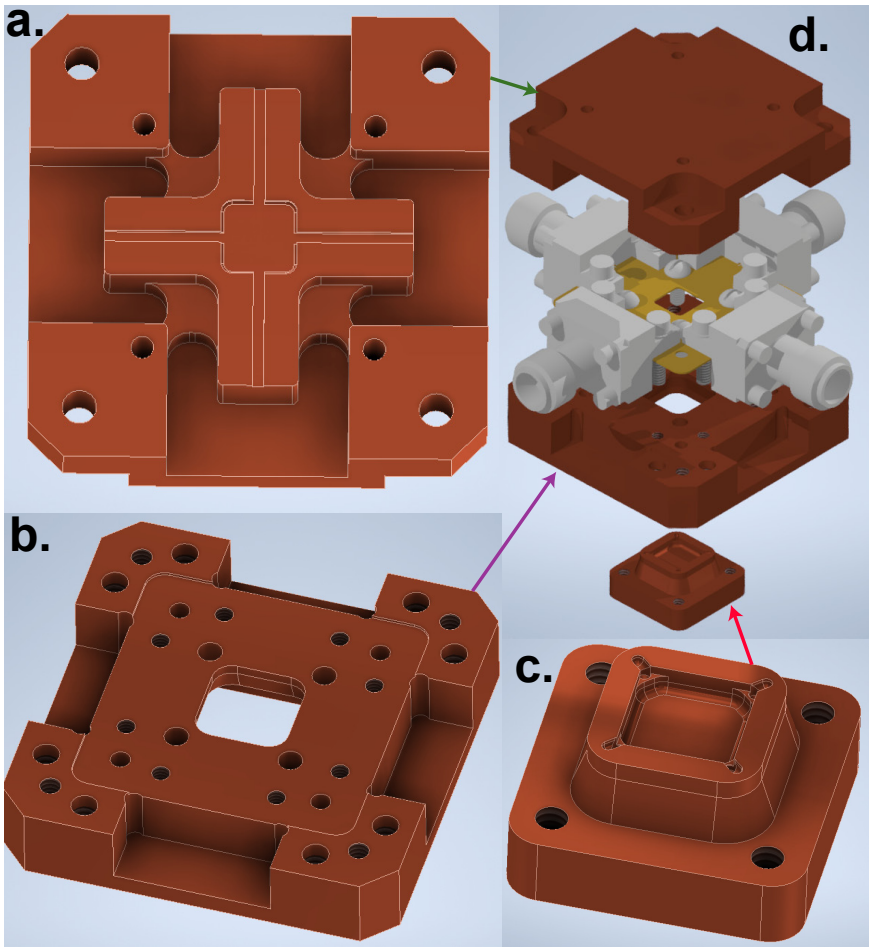

Figure 8.3: Mode-free K Band packaging. Three copper pieces align the chip containing qubits with the low-loss printed circuit board.

plish this with a copper-clad ceramic-filled PTFE dielectric [5] circuit board patterned with via-fenced [19] coplanar waveguides for signal routing. Using compressed indium wire, the sapphire substrate is aligned and clamped to the circuit board on each of its four corners by a high-purity copper enclosure, as shown in Figure 8.2. This enclosure, is composed of three precision machined pieces, illustrated in Figure 8.3, which also serve to thermalize and mechanically secure the sample, as well as prevent leakage by fully enclosing the chip on all sides. The enclosure is designed to naturally align the chip with the circuit board. This suspends the sapphire so that the surface qubits are well separated from the lossy copper

5. Rogers 3003

surfaces [90], with the resulting dimensions optimized such that the nearest package mode lies above 27 GHz. The signal is coupled on and off the sapphire substrate with wirebonds. Since each of these has a high inherent inductance of about 1 nH/mm [249], which is increasingly problematic at higher frequencies, we use sets of 3 wirebonds for each connection: their location and profile is optimized in conjunction with on-chip signal launch geometry for better than 0.2 dB insertion loss from 0-30 GHz. The optimized geometry is shown in Figure 8.4. The assembly is thermalized to the base stage of a dilution refrigerator (65-95 mK), where transmission measurements through the central feedline can be used to characterize resonators and qubits [113, 201].

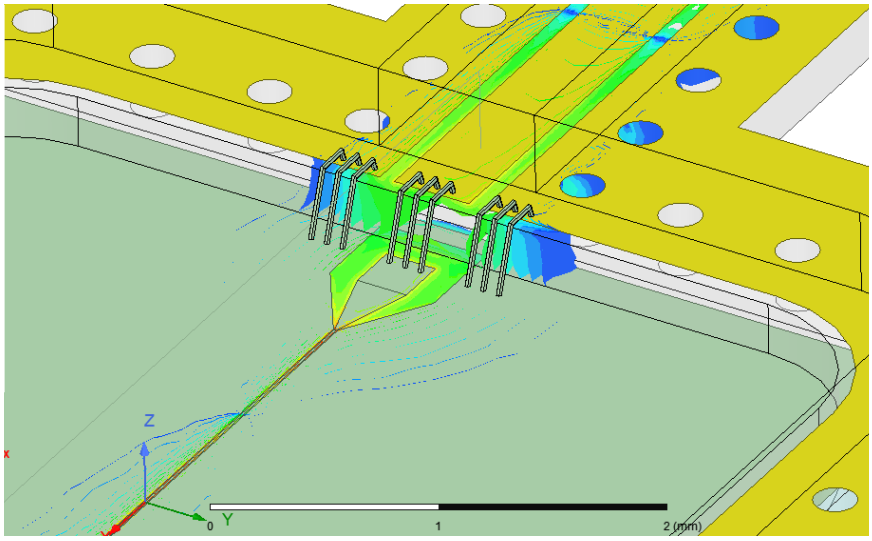

Figure 8.4: Optimized launcher geometry and wirebond configuration which achieves maximal transmission up to 27 GHz. Using a manual wirebonder, we attempt to replicate this wirebond shape, however in practice the circuit board dimensions requires slightly longer bonds.

8.3 Qubit Characterization

While monitoring transmission through the system at the low-power readout resonator frequency, simultaneously applying a second probe pulse reveals the energy spectrum of the qubit, as shown in Figure 8.5a. At low powers, we observe a deflection in transmission at the

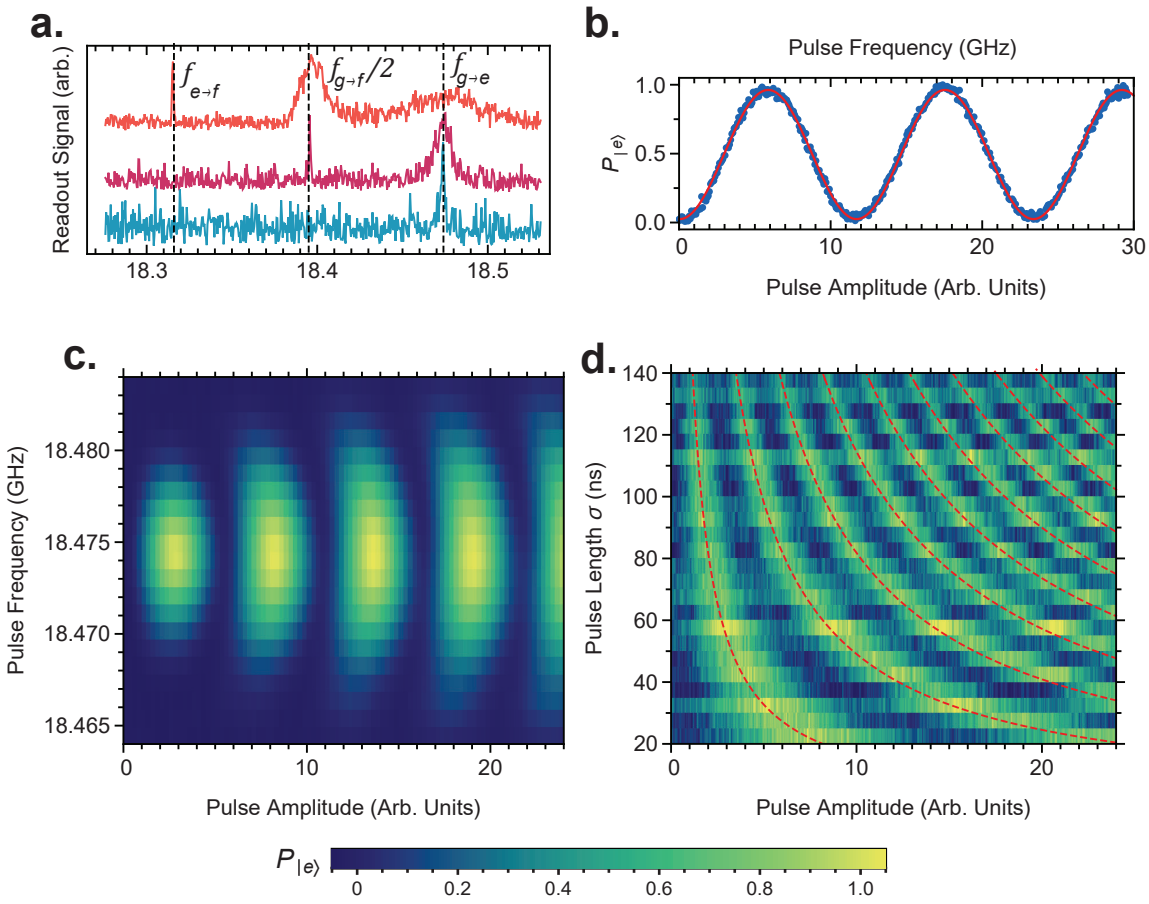

Figure 8.5: Qubit dynamics. **a)** Deflection of the readout resonator transmission signal as a function of applied qubit pulse frequency, shown for increasing qubit pulse power. At low powers (blue) a single peak is observed when the pulse is resonant with the qubit frequency ($f_{ge} = 18.474$ GHz). As power increases, the linewidth of this transition increases, and additional peaks appear from excitations into higher qubit levels through many-photon excitations ($f_{gf}/2$ etc). These features have a spacing of $\alpha/2 = (f_{ge} - f_{ef})/2$. **b)** Measured excited state probability shows Rabi oscillations when a fixed-length pulse with varying amplitude is applied at the qubit frequency. The red line is a fit to the expected sinusoidal behavior. **c)** Rabi oscillations are measured for frequencies near f_{ge}, with brighter colors corresponding to higher excited state probabilities. Away from the transition frequency, the Rabi frequency increases while the oscillation amplitude decreases and becomes power dependent. **d)** Rabi oscillations as a function of pulse amplitude and length σ, with brighter colors corresponding to higher excited state probabilities. Dashed red lines mark contours of integer π pulses where $\sigma\Omega = m\pi$.

bare qubit frequency f_{ge} as the qubit is excited, resulting in a dispersive shift of the resonator frequency. Increasing the power of the probe pulse reveals the higher energy states of the

qubit through two-photon processes $f_{gf}/2 = (f_{ge} + f_{ef})/2$ and excited state transitions f_{ef}. This allows us to measure the anharmonicity $\alpha \equiv f_{ef} - f_{ge}$, which for our qubits is typically near 200 MHz. Notably, this level spacing (which sets an upper bound on qubit gate speed) is similar to many conventional microwave qubits [21, 113, 178], but could easily be adjusted in our design by picking a different junction J_c density or capacitor size.

An applied signal on-resonance with f_{ge} will result in Rabi oscillations between the ground and excited state at a Rabi frequency Ω proportional to its amplitude. We explore this behavior by applying fixed-length ($\sigma = 25$ ns) Gaussian pulses at the qubit frequency with varying amplitude (over which we have much finer control than time) as shown in Figure 8.5b. This behavior is captured by a sinusoidal model which oscillates with pulse amplitude and length. We repeat this measurement with $\sigma = 60$ ns while varying pulse detuning from the transition $\Delta = f - f_{\mathrm{ge}}$, and summarize the results in Figure 8.5c, where brighter colors indicate the system in the excited state. At the qubit frequency, we observe a series of bright fringes; however as the detuning from the transition increases, the oscillation rate increases, while the oscillation amplitude is quickly suppressed. The bandwidth of the fringes is further reduced by the finite nature of the pulse [24, 26, 67].

To verify the time-dependence of the Rabi oscillations, we also repeat this measurement at $\Delta = 0$ while varying σ and plot the results in Figure 8.5d. Consistent with an amplitude dependnt oscillation rate, the observed fringes are evenly spaced for fixed amplitude (vertical slice) or fixed length (horizontal). Fitting the contours of the fringes to the periodic conditions $\sigma\Omega = m\pi$ allows us to extract the Rabi rate (as high as 100 MHz), and also serves as a good method for calibrating a qubit π control pulse.

8.3.1 Number Splitting

As touched on in Chapter 4, the dispersive shift χ used to measure the qubit state also work in reverse. When higher powers are used to measure the qubit readout resonator, the

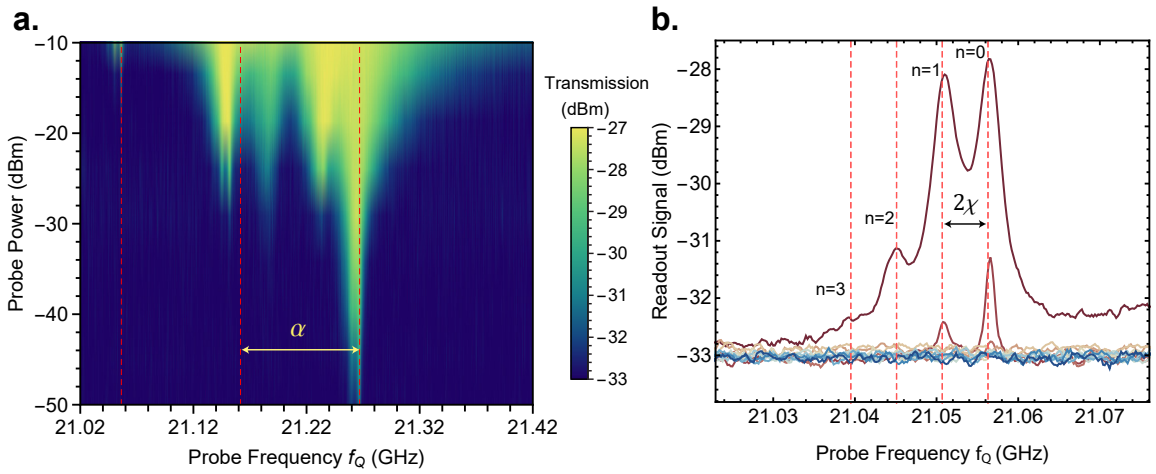

Figure 8.6: Number splitting. **a)** Deflection of the readout resonator transmission signal as a function of applied qubit probe frequency shown for increasing powers. We observe the level transitions separated by anharmonicity $\alpha/2$ similar to Figure 8.5a, however on closer inspection each peak is split into smaller features. **b)** With non-negligible readout resonator population, the sub-peaks for each transition are split by the dispersive shift 2χ.

resonator photon population shifts the qubit level transitions, as shown in Figure 8.6. In this case we are lucky to have a transition linewidth smaller than the dispersive shift χ so are able to resolve the individual transitions! This allows us to measure the state of the resonator, which in this case is likely a thermal coherent state at elevated readout powers.

8.4 Qubit Coherence Properties

We can now study qubit relaxation time and dephasing time, which dictate qubit limitations and act as sensitive probes of loss channels. We measure relaxation time by placing each qubit in its excited state and measuring it after time t: fitting the exponential decay gives the characteristic time T_1. We perform these measurements for each qubit and show averaged results as a function of qubit frequency in Figure 8.7c, finding $T_1 = 1.6\,\mu s$ for our best device. The highest-frequency qubits (for example Wafer C) are likely affected by Purcell loss, as their transitions are above their readout resonator frequencies (~ 22 GHz). To probe loss channels in detail we use the frequency-independent qubit quality factor $Q_1 = 2\pi f_q T_1$,

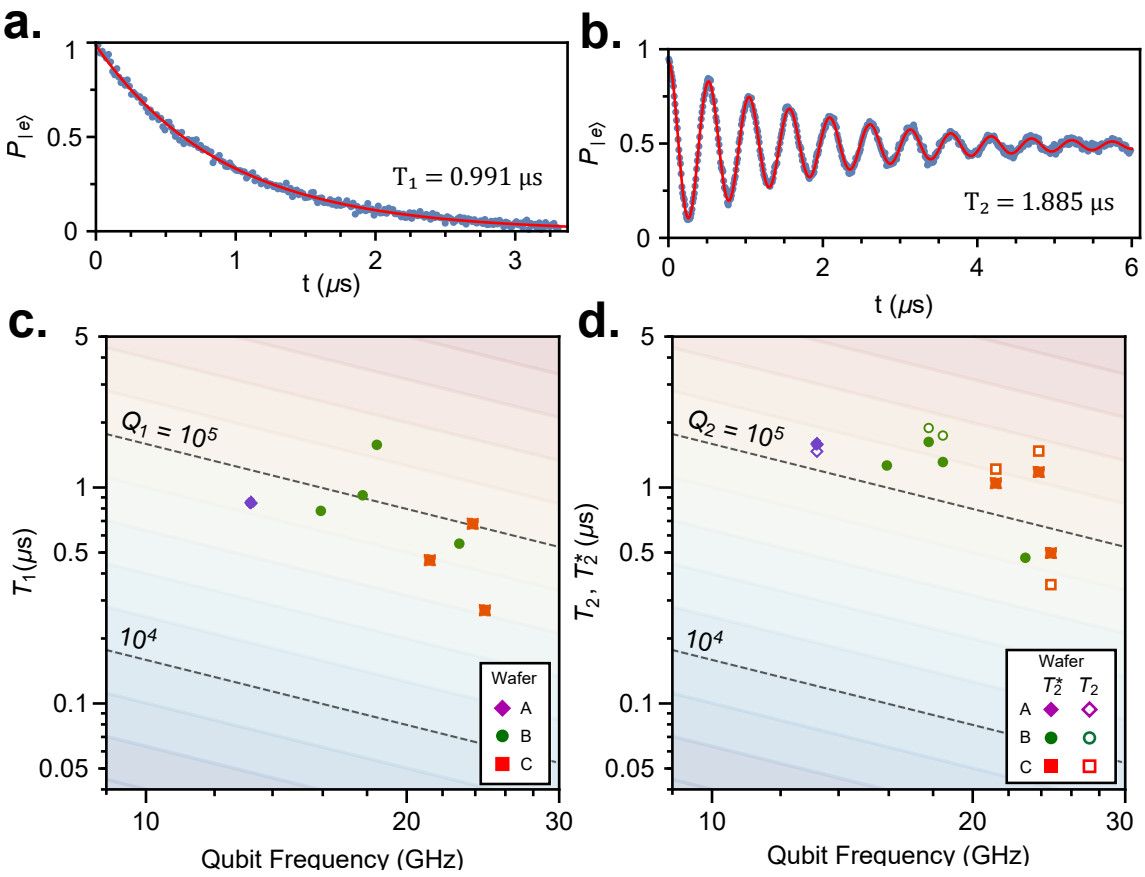

Figure 8.7: Qubit Properties. **a)** Average qubit decay time T_1 extracted by fitting the exponential decay of excited state population. **b)** T_1 plotted as a function of qubit frequency, grouped by wafer. Lines indicate qubit quality factor $Q_1 = \omega_q T_1$. We find an overall mean Q_1 of 0.792×10^5 with some wafer to wafer variation. **c)** Ramsey dephasing time T_2^* (filled points) and Hahn-echo dephasing time T_2 (hollow points) extracted by fitting the exponential decay of oscillations in (d) as a function of qubit frequency. Lines indicate dephasing quality factor. We find an average T_2^* and T_2 of 1.124 μs and 1.357 μs respectively.

which we find for our devices is just below 10^5. When rescaled to equivalent values for conventional microwave qubits, the coherence properties of our K-band qubits are comparable to transmons made with a similar junction process, which measure an average $Q_1 = 2.5 \times 10^5$ [13]. The slight decrease in our average Q_1 could be attributed to slightly higher junction and surface participation [241] arising from our smaller qubit capacitor size relative to the junction area. Quantum decoherence and material properties in this frequency range are still relatively unexplored, warranting further investigation into the nature and sources of loss.

We also perform a Ramsey experiment to measure the dephasing time T_2^*, and a Hahn-echo experiment to characterize the spin-echo dephasing time T_2 (Figure 8.7b). We find that T_2^* is consistently higher than T_1, nearing the dissipation limit $2T_1$. This suggests that pure dephasing rates are relatively low, which is expected from the low photon occupation of the higher-frequency readout resonators. The T_2 values, which add an additional π pulse to decouple the qubit from low frequency noise, are not significantly different from T_2^*. This is likely a result of the qubit decaying during the added spin-echo pulse, which is relatively long (120–240 ns) to compensate for low output power from our upconversion hardware. Since the bandwidth of these pulses is still smaller than the level spacing, our single-qubit gate speeds could be optimized to be significantly shorter [191].

8.5 Thermal Dependence of Decoherence

A key benefit of these high-frequency qubits is their thermal resilience to quasiparticles and thermal photons, so we investigate our qubits at increased operating temperatures, shown in Figure 8.8a. We observe a small decrease in T_1 with temperature above 300 mK, consistent with qubit heating from its environment [129], but importantly we don't see the drastic temperature dependence seen at these temperatures in qubits with aluminum junctions due to quasiparticle-induced loss [39, 40, 74], in line with expectations for niobium. More importantly, the measured Ramsey dephasing time shown in Figure 8.8b doesn't significantly decrease until 200 mK or so, due to the increased frequency of readout resonators in these devices compared with conventional qubits. We are able to capture this behavior with the following parameter-free model for dephasing from thermal photons in the readout resonator [43, 189]:

$$T_\phi^{-1} = \Gamma_\phi = \frac{\gamma}{2}\mathrm{Re}\left[\sqrt{\left(1+\frac{2i\chi}{\gamma}\right)^2 + \frac{8i\chi}{\gamma}n_{\mathrm{th}}} - 1\right] \tag{8.1}$$

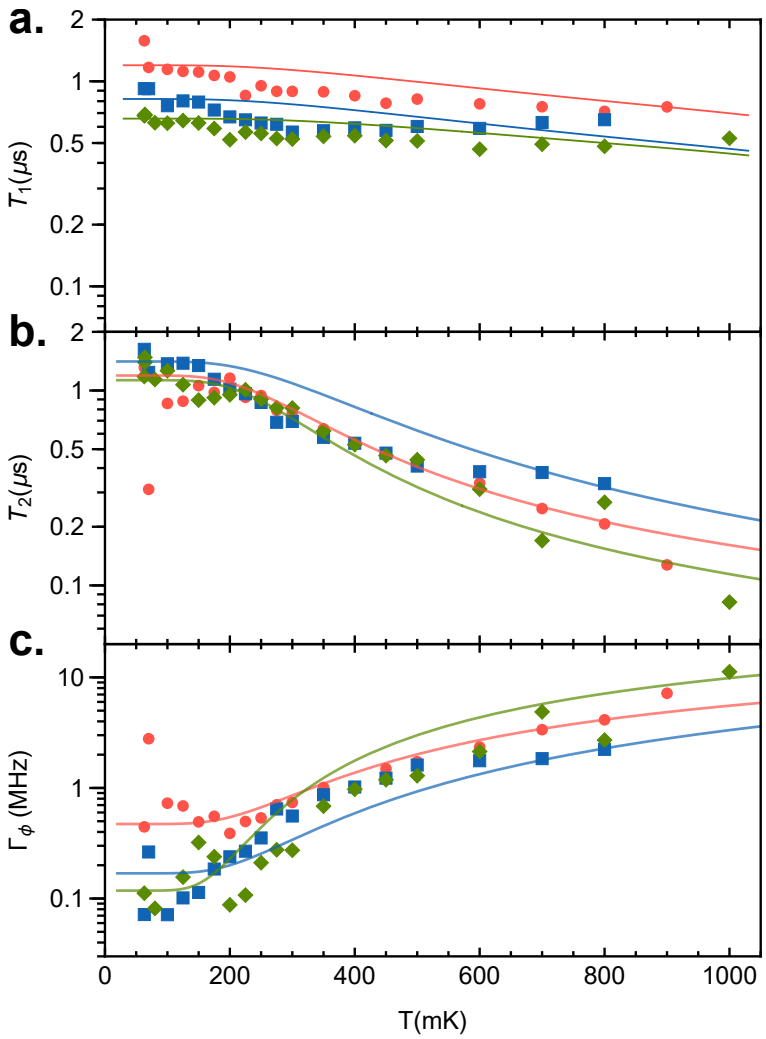

Figure 8.8: Thermal decoherence and dephasing. **a)** Decoherence time T_1 of three representative qubits measured as a function of temperature. A mild decrease is observed at higher temperature, consistent with a model including loss from increased system bath temperature (solid lines). **b)** Ramsey dephasing time T_2^* as a function of temperature. The behavior is largely captured by a parameter-free thermal dephasing model assuming a fixed T_1 (solid lines). **c)** Pure dephasing rate Γ_ϕ which has dephasing from relaxation subtracted, resulting in better agreement with the model.

Where χ is the dispersive coupling, γ is the decay rate of the readout resonator, and $n_{\text{th}} = 1/(e^{hf_R/kT}-1)$ is the resonator thermal population, set by resonator's fundamental frequency f_r. The measured dephasing rate is then given by $T_2^{-1} = T_\phi^{-1} + T_{2,0}^{-1} + 1/(2T_1)$, where T_1 and $T_{2,0}$ are measured by averaging the low-temperature values. The model, which contains

only independently-measured quantities, is overlaid for each qubit in Figure 8.8b, showing relatively good agreement. We can also go further by using the measured values of T_1 to solve for Γ_ϕ directly, as shown in Figure 8.8c, and overlay the theoretical model, which results in even better agreement. Thus, unhindered by quasiparticle loss, we have been able to probe dephasing and decoherence up to 1 K from purely thermal sources.

8.6 Conclusion

Using high-temperature niobium trilayer Josephson junctions fabricated using improved fabrication methods, we have realized transmon qubits operating as high as 24 GHz, which demonstrates good coherence properties in these high-critical-current junctions. Our qubits show comparable coherence properties to typical microwave qubits when frequency is accounted for, likely still limited by decoherence channels in fabrication defects. With our qubit decoherence unaffected by quasiparticles up to 1 K, we are able to explore the thermal dependence of dephasing in superconducting qubits, and find that our qubits can operate up to approximately 250 mK while maintaining similar performance. Our devices open the door to new opportunities for quantum experiments at elevated temperatures and paves the way to even higher frequency, higher temperature qubits.

CHAPTER 9

IMPROVED MILLIMETER-WAVE MEASUREMENTS

Extending superconducting quantum device functionality to millimeter-wave frequencies (near 100 GHz) offers new opportunities for detection and transduction[181] and access to large coupling strengths for hybrid experiments [44, 254]. Most importantly, the reduced sensitivity to thermal noise of higher-energy mm-wave photons could enable quantum experiments at liquid-helium-4 temperatures, which significantly reduces cooling complexity and power dissipation constraints, enabling new pathways for scaling up quantum computing platforms, and could facilitate direct integration with high-speed superconducting digital logic. In order to establish more robust quantum information systems, it is vital to understand and minimize decoherence in superconducting devices. Significant effort has gone towards investigating and reducing sources of loss at microwave frequencies [148], establishing that a significant remaining contribution to decoherence at the single photon level comes from two-level-systems (TLS) found in amorphous dielectric materials. However, since probing individual decoherence mechanisms requires reducing other sources of loss, the nature and limits of TLS loss contributions and their frequency dependence when scaled to millimeter-wave frequencies remain relatively unexplored.

When measured at single-photon energies, the internal quality factor (Q_i) of an on-chip resonant circuit provides insight into the maximum coherence of a quantum system formed by adding a source of nonlinearity (such as kinetic inductance [12, 215] or a high-frequency Josephson junction [13, 109]). Significant progress has been made in improving millimeter-wave resonators [12, 62, 71, 205, 215, 232], but their single-photon quality factors remain below 2-4 $\times$ 10^4: significantly lower than those measured in microwave circuits [150]. This can largely be attributed to two primary factors: first, millimeter-wave resonators frequently use substrates with low dielectric constants (such as SiO_2) to simplify high-frequency circuit design, despite their poor dielectric loss characteristics compared to crystalline silicon or

sapphire. Second, many existing millimeter-wave resonators rely on coplanar stripline or microstrip transmission line components, both of which have relatively high radiation profiles resulting in increased radiative losses, particularly at high frequencies [235].

Taking inspiration from low-loss microwave devices, a ground-shielded circuit design minimizing radiative loss offers an attractive method for increasing millimeter-wave quality factors, and better control over on-chip signal propagation and coupling. Whereas low-frequency signals can be routed to an on-chip waveguide directly through wire bonds, grounded circuits present additional design challenges in millimeter-wave bands, where signals are primarily transmitted by hollow waveguides, so require a method to efficiently direct the waveguide electromagnetic fields onto the chip. This proves to be a challenging problem, and as a result, a variety of transition structures coupling waveguides with on-chip transmission lines have been developed [94]; however, with no universal solution, transitions for specific applications are still actively studied and improved.

In this chapter, we use a transition specifically designed to measure ground-shielded millimeter-wave resonant circuits with improved control. We characterize a tapered coupling structure that efficiently confines the signal fields to an on-chip slotline waveguide, finding an insertion loss better than 0.5 dB over 14 GHz of bandwidth. We study niobium resonators patterned near the slotline and show that resonator coupling can be controlled independently without increasing radiation loss. With this novel resonator design combined with fabrication procedure, we improve on existing planar millimeter-wave devices and achieve high internal quality factors (Q_i) consistently exceeding 10^5 at single photon powers. This allows us to study the effects of oxide growth and removal on remaining loss contributions to Q_i from millimeter-wave two-level systems (TLS), and show that limits from TLS can be increased as high as 10^6.

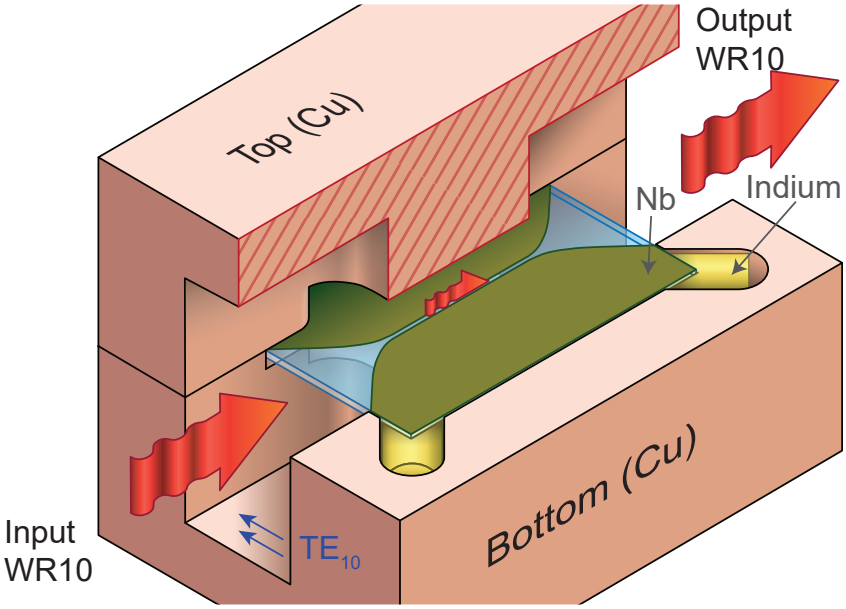

Figure 9.1: Cutaway diagram showing assembled back-to-back waveguide to on-chip slotline transition structures, with signal propagation marked in red. In the slotline region the waveguide geometry is constricted to prevent spurious propagating modes. On the chip corners, rounded channels allow indium wire (yellow) to deform, which secures the chip in place.

9.1 Tapered Waveguide Transition Design

When designing a millimeter-wave circuit to minimize loss, additional constraints apply to its coupling structure. Crystalline sapphire is an ideal substrate choice as it yields much lower dissipation than most dielectrics, but is more difficult to machine which imposes design limitations. As a result, transition designs requiring abnormally cut non-rectangular substrates [97, 234, 256], micromachining [124] or drilled holes [8] are impractical. Sapphire also presents additional challenges for high frequency circuit design due to its relatively high dielectric constant, which leads to more pronounced impedance mismatches caused by the presence of substrate in the waveguide [258]. Ensuring currents are carried by superconducting materials minimizes conduction loss. To achieve this, on-chip superconducting layers should be well separated from both waveguide and housing metal [90]: this consideration makes transitions with stripline geometries, which needs an external ground plane [253], less

ideal. While potentially offering low radiation loss, coplanar waveguide transitions [124, 163] are typically more complex, requiring multiple stages and more physical space. Finline transitions, on the other hand, consist of a single taper [72, 97, 256], making them more compact, and transform the signal into a differential mode localized on the chip surface.

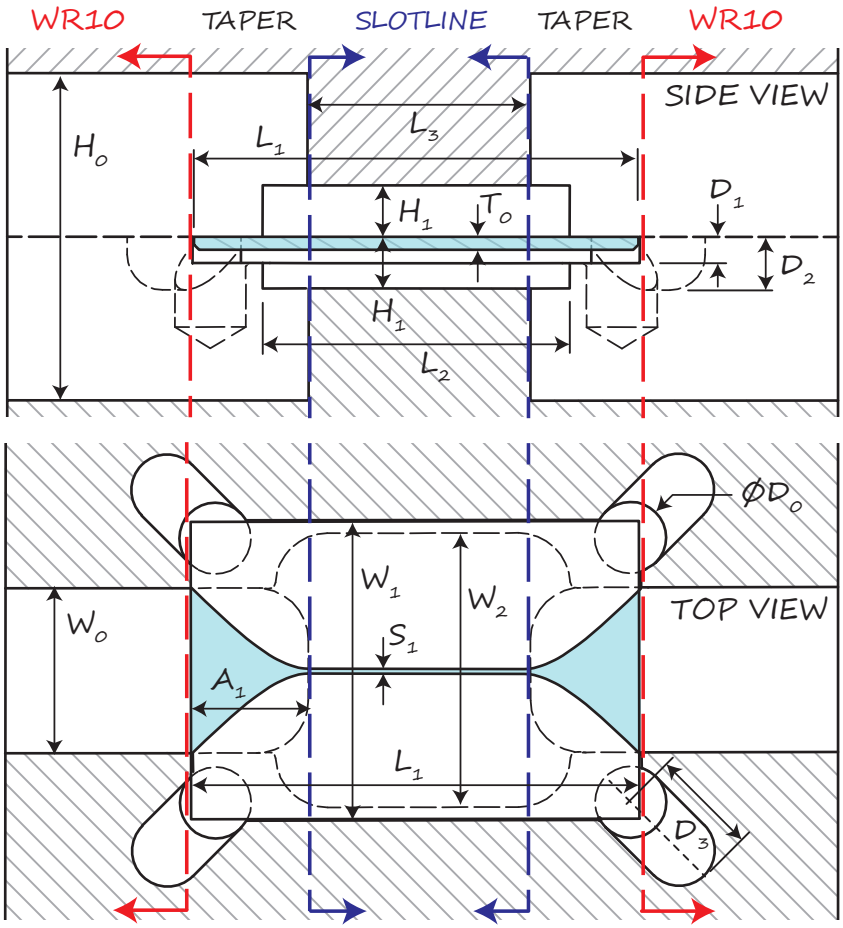

Figure 9.2: Side and top section views of the transition structure geometry, with reference planes and relevant dimensions marked. Un-metallized sections of the chip are shown in blue. The rounded corner channels are completely filled by indium.

Our transition is defined by a superconducting niobium film patterned on the top surface of a rectangular crystalline sapphire substrate centered in a rectangular waveguide. Inspired by Refs. [8, 72, 256], the geometry consists of a differential unilateral finline that tapers from the waveguide width down to a narrow slotline. A cutaway diagram of the complete structure is shown in Figure 9.1, with two transitions coupled back-to-back by a length of on-chip

Table 9.1: Optimized dimensions used for the tapered finline waveguide-slotline transition (in mm).

H_0	L_1	L_2	L_3	T_0	H_1	D_1	D_2
2.54	3.45	2.27	1.66	0.1	0.4	0.2	0.412
W_0	W_1	W_2	S_1	A_1	D_0	D_3	
1.27	2.29	2.11	0.04	0.895	0.55	0.8	

slotline useful for coupling to resonators. Where the transition terminates, the waveguide height is reduced to increase the cutoff frequency of unwanted higher-order propagating modes, while the waveguide width is broadened to increase the usable area of the chip: this allows the device area to remain suspended away from metal surfaces, which minimizes conduction loss [90]. The only direct contact with the copper enclosure occurs where the chip is clamped at its corners.

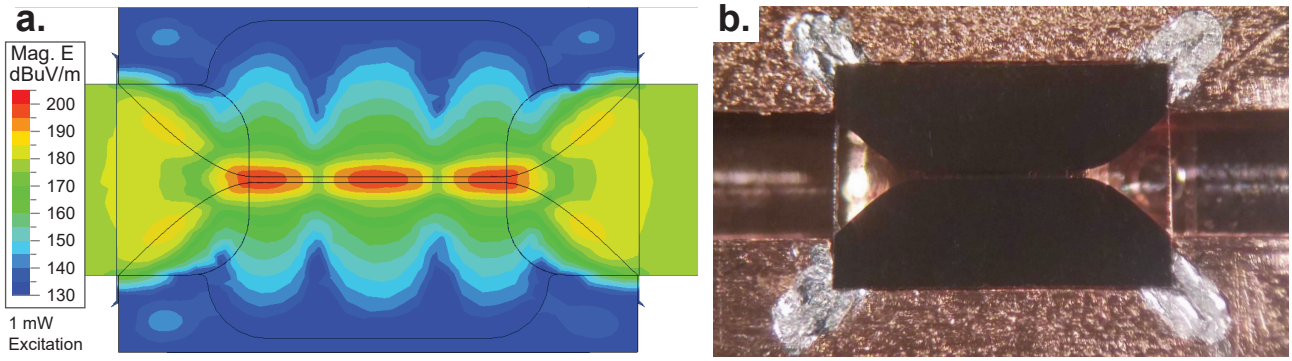

Figure 9.3: **a)** Simulated electric field distribution (log scale) viewed from the top for a wave traveling through the structure. Notably, a much greater dynamic range of electric field strengths is achieved as the signal is compressed into the slotline, as compared to just the waveguide field. **b)** Photograph of a mounted chip with back to back transition structures, with top copper block removed. The indium, visible on the corners of the chip, deforms and fully fills the designated channels, thermalizing and securing the chip.

Due to the pronounced impedance discontinuity between the sapphire chip and the waveguide [258], a broadband matching structure is difficult to achieve with an exponential or cosine Vivaldi taper contour as is typically used in finline transitions [8, 72]. Instead, we find that a curved taper shape with a nearly linear waveguide entrance can compensate for

the mismatch and can be optimized to give good performance over a section of waveguide bandwidth. The optimized contour is described by the function:

$$y(x) = \left(\frac{W_0 - S_1}{2}\right) \frac{x}{A_1} \sqrt{2 - \left(\frac{x}{A_1}\right)^2} \tag{9.1}$$

where W_0 is the smaller waveguide dimension, S_1 is the slotline width, and A_1 is the transition length. The geometry of the transition is detailed in Figure 9.2. Using finite element method simulation software[1], the above contour function and parameterized geometry dimensions are optimized for maximal transmission in the 90-100 GHz band. The resulting optimal dimensions are listed in Table 9.1. Notably, by relaxing the bandwidth optimization constraint, we achieve a taper structure less than 0.9 mm long: much smaller than the $\lambda/2$ value predicted with analytical functions [258] and more compact than many other transition structures in literature [8, 72, 91, 97, 256].

The benefit of coupling to resonant devices through an intermediate on-chip slotline transmission line is apparent when examining the magnitude of the electric field of a wave propagating through the structure, shown in Figure 9.3a. Whereas some previous implementations of millimeter-wave resonators [12, 215] interact with a uniform waveguide electric field (visible on the left and right ends of Figure 9.3a) and rely on varying resonator dipole moments to adjust coupling, our method compresses the signal to a 40 μm slotline resulting in over 50 dB of dynamic range in electric field strength across the usable area of the chip. Consequently, the dipole coupling strength for each resonator can be set by its location on the chip, without needing to adjust the resonator's dipole moment, leaving more freedom to optimize the resonator performance.

1. Ansys HFSS

9.2 Cryogenic Measurements and Calibration

The taper geometry described above is defined by chlorine reactive ion etching a 100 nm thick film of high-purity electron-beam-deposited niobium grown on a 100μm thick crystalline C-plane sapphire substrate (see Appendix A for detailed procedure). Short sections of indium wire are used to secure and thermalize the chip to the copper waveguide enclosure, visible in a photograph of the mounted chip shown in Figure 9.3b. The edges of the on-chip taper structure match the waveguide dimensions, which enables visual alignment during mounting (necessary to maximize coupling).

The assembled structure is then measured at 0.86 K in a helium-4 cryostat: using a vector network analyzer with millimeter-wave extension modules and a cryogenic low noise amplifier, we measure the complex response in transmission and reflection. Input attenuation and cryogenic isolators reduce thermal noise reaching the sample, enabling measurements in the single photon limit. These methods are detailed below.

9.2.1 Experimental Measurement Setup

All millimeter-wave characterization was performed in a custom built ^{4}He adsorption refrigerator, with a base temperature of 0.86 K, and a cycle duration of 3 hours. We generate millimeter-wave signals (75-115 GHz) at room temperature by sending microwave signals (12-19 GHz) into a frequency multiplier. The upconverted signal is sampled to establish a phase reference measurement. We convert the generated waveguide TE_{10} mode to a 1 mm diameter stainless steel and beryllium copper coaxial cable, which carries the signal to the 1 K stage of the fridge, thermalizing mechanically at each intermediate stage, then convert back to a WR-10 waveguide which leads to the device under test. The cables and waveguide-cable converters have a combined frequency-dependent loss ranging from 38.6 dB to 49.8 dB in the W-Band, dominated by the cable loss. In the case of a transmission measurement, the signal is further thermalized to 1 K by a cryogenic 20 dB attenuator, and in the case

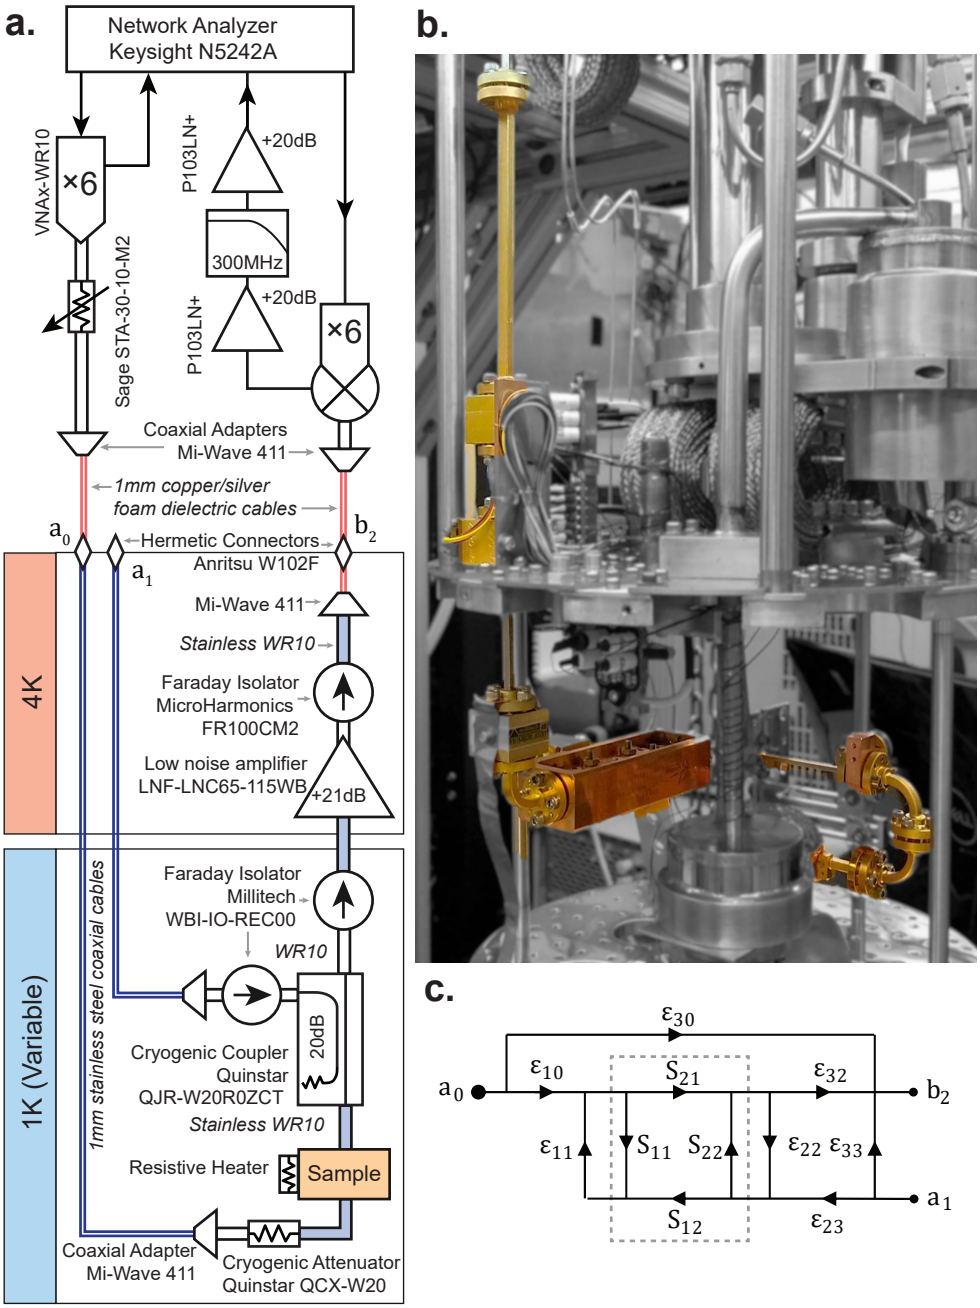

Figure 9.4: **a)** Schematic of cryogenic millimeter-wave measurement setup. Colored tabs show temperature stages inside the Helium-4 adsorption refrigerator, which reaches a base temperature of 0.86 K. **b)** A photograph highlights relevant hardware inside the fridge. **c)** Adjusted error network used for cryogenic TRL calibrations, where S_x are the respective S-parameters of the sample, and the measurement paths a_0, a_1 and b_2 are also labelled in the schematic.

of a reflection measurement, this thermalization is accomplished with a 20 dB cryogenic directional coupler with a copper body. The sample is thermally isolated from the 1 K stage of the refrigerator to allow local heating for temperature sweeps.

Wideband cryogenic millimeter-wave circulators are currently not commercially available, so instead a cryogenic directional coupler allows enables reflection measurements by allowing nearly all of the reflected and transmitted signal to pass through to a low noise amplifier. Cryogenic faraday isolators minimize retro-reflections and prevent thermal radiation from leaking in on the output side, while still allowing good transmission. Having passed outside the cryostat through custom-built hermetic adapters, the signal is downconverted and amplified for measurement. The entire setup is summarized in Figure 9.4a-b.

9.2.2 Cryogenic Calibration

The cryogenic measurements described above introduce a complex network between the sample and the measurement equipment, a calibration must be performed to obtain accurate estimates of network parameters. The necessity of attenuating components as well as active components with gain in the measurement chain make cryogenic calibrations a complex problem [38, 243, 257]. Without access to a cryogenic millimeter-wave switch that can instantly select a calibration standard in-situ, we instead rely on a standard TRL-type calibration [184] with carefully controlled sequential cooldowns with each standard. The cryogenic measurement network can summarized with a partial network of error adapters in Figure 9.4c.

The input paths a_0 and a_1 both have significant attenuation (ϵ_{10} and ϵ_{23}) to prevent room temperature noise from reaching the sample. Because of this, signals directly reflecting off the samples (which are doubly attenuated) are too faint to measure, especially when overlaid on the imperfect return loss on some of the millimeter-wave components in the chain. Instead, we use two separate input paths in order to characterize the response of the sample: when

combined, this yields an error adapter network similar to the familiar two-port TRL network [184].

Aside from attenuating terms (ϵ_{10} and ϵ_{23}) and the output amplification (ϵ_{32}) which suffice for a simple correction, we must take into account several non-ideal terms in our error network. As a consequence of highly attenuating inputs is that when normalized, the directivity terms (ϵ_{30} and ϵ_{33}) are significantly more pronounced. Additionally, the cryogenic attenuators and absorbers in our system appear have non-ideal return loss characteristics, resulting in non-negligible source match ϵ_{11} and even more pronounced directivity ϵ_{33}. Since we measure good return loss in the cryogenic faraday isolators used on the output line (see Figure 9.4a), the load match term ϵ_{22} should be significantly less pronounced.

As a result, we can simplify the procedure by neglecting ϵ_{22}, and make a further assumption that our sample structure is symmetric (in practice this can be checked by reversing the sample direction). These simplifications leave some unresolved frequency ripples on the measured spectrum (visible in Figure 5). The measured transmission and reflection response is now simply:

$$S_{21}^{\mathrm{M}} = \frac{b_2}{a_0} = \epsilon_{30} + \frac{\epsilon_{10}\epsilon_{32}S_{21}}{1-\epsilon_{11}S_{22}} \tag{9.2}$$

$$S_{22}^{\mathrm{M}} = \frac{b_2}{a_1} = \epsilon_{33} + \left(S_{22} + \frac{\epsilon_{11}S_{21}^2}{1-\epsilon_{11}S_{22}}\right)\epsilon_{23}\epsilon_{32} \tag{9.3}$$

By performing measurements of through, reflect, and line standards which have known S-parameters [184] we can solve this system of equations for relationship the error terms. Combining this with direct measurements of each input line yields true values for each error term. With the system now characterized, we can extract sample S-parameters from the transmission and reflection measurements above, giving a much more accurate picture. Since each measurement now relies on a number of calibration experiments, this imparts a linear error (with typical vector magnitude between $0.03-0.09$) on the final complex S-parameter. As the amount of measurement uncertainty now depends on the measurements themselves

(defined by the relationships in Equation B.1 - B.2), this results in S-parameter uncertainty that increases as the S-parameter decreases, apparent in the significantly higher uncertainty on our reflection measurements.

9.3 Waveguide Transition Characterization

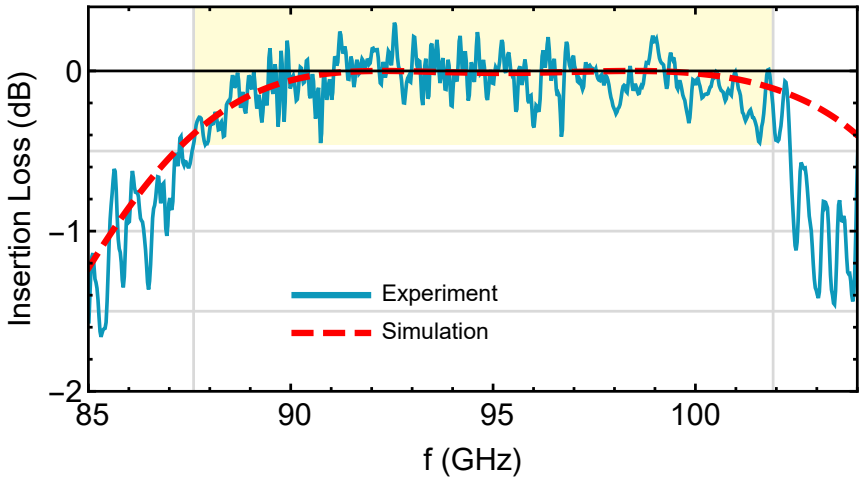

Figure 9.5: De-embedded insertion loss for two back-to-back transitions along with simulated values. In the band of interest (highlighted), we find an insertion loss better than 0.46 dB, limited by de-embedding calibration uncertainty (~ 0.3 dB).

The effectiveness of the taper transitions is tested by using transitions to convert a waveguide signal to a slotline and back, and monitoring the cryogenic transmission and reflection of the structure. With the calibration procedure described above, the effects of additional hardware on the input and output lines are removed, and de-embedded scattering parameters of the sample are recovered from cryogenic network measurements. The measurement results are summarized in Figure 9.5 and Figure 9.6. We find the transition performs best between $87.6 - 102.4$ GHz, exceeding the designed range, and define this 14.8 GHz wide frequency range as the useful operating band. Within the operating band, we find a maximum insertion loss of approximately 0.46 ± 0.35 dB (or 94.8% transmission), corresponding to a coupling efficiency of ~ 0.23 dB for a single taper structure. However these values are likely

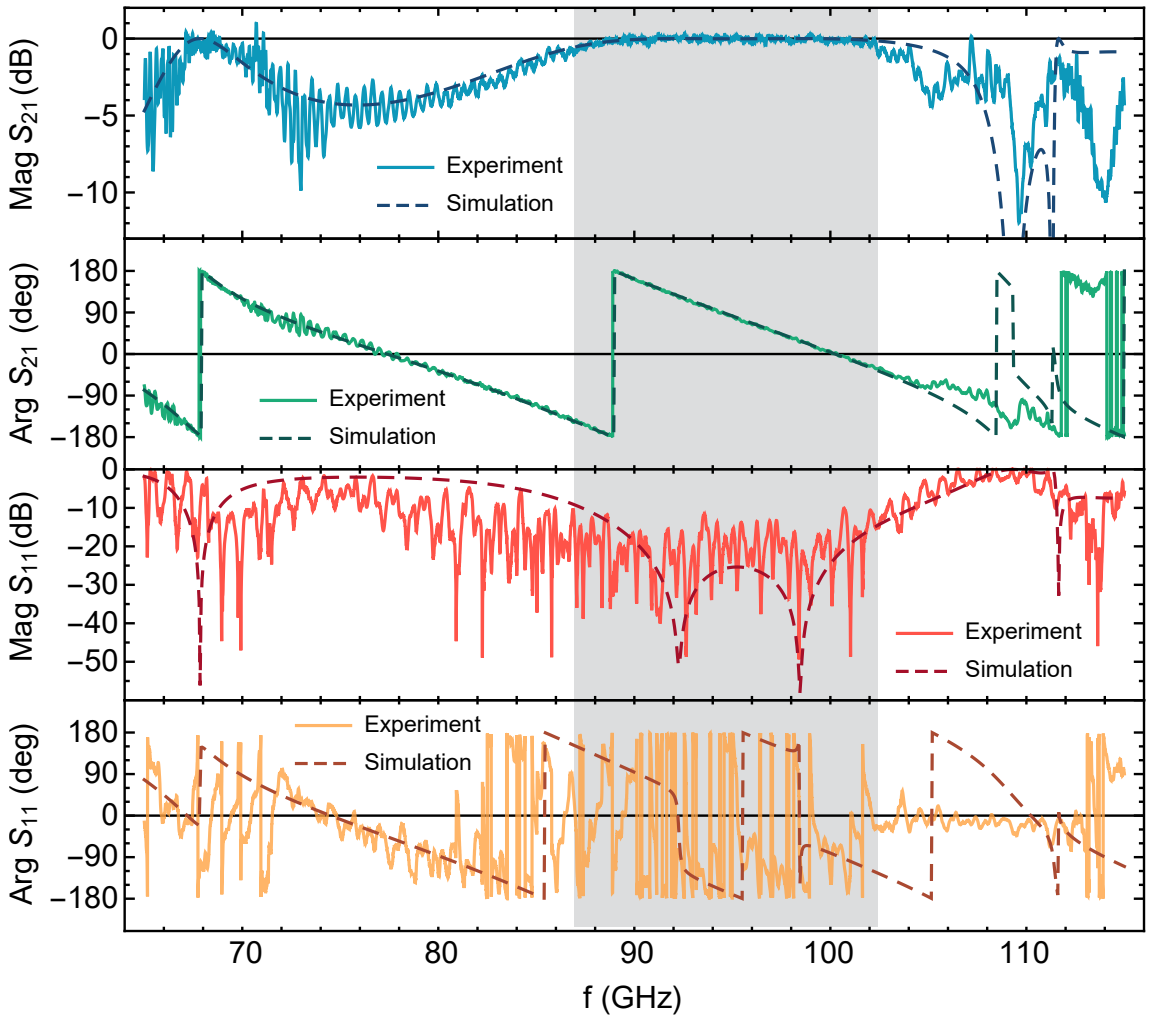

Figure 9.6: Wideband de-embedded scattering matrix parameter measurements (solid lines) for two transitions back-to-back along with respective simulations (dashed lines) showing good agreement. In the operating band marked in gray, we find a total insertion loss better than 0.5 dB, and return loss less than -13 dB.

dominated by errors introduced by calibration methods (which do not enforce passivity).

Across the W band, we find that the de-embedded transmission and reflection of the structure are in fairly good agreement with simulations, with the exception of the region near 110 GHz: in this region unwanted resonances occur in the substrate and indium mounting regions, which are difficult to predict. In the operating band, we find the return loss exceeds 13.1 ± 8.6 dB, which slightly deviates from simulation, but could be attributed to the signif-

icantly increased calibration uncertainty due to the high transparency of the structure and increased sensitivity to error terms. Combined, these measurements demonstrate a transition structure in good agreement with simulation, which in the operating band couples a signal on and off a chip with high efficiency.

9.3.1 Leakage Bypassing the Transition

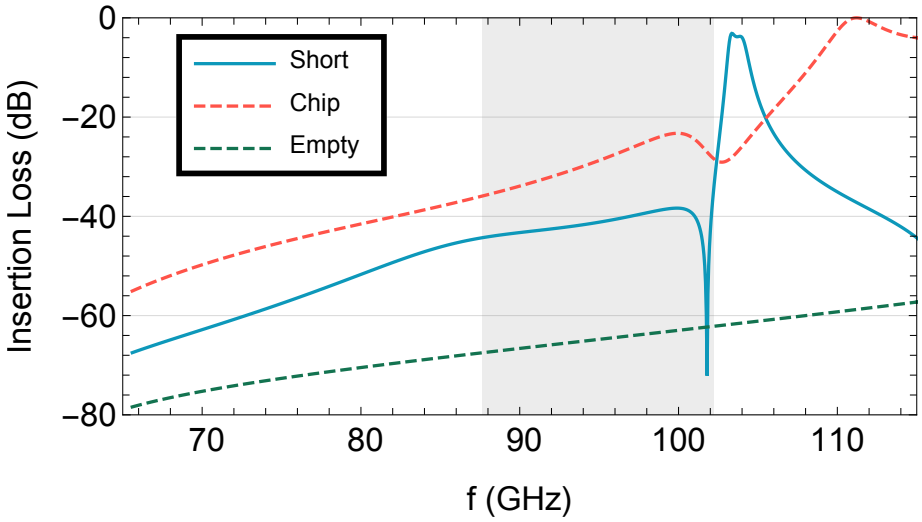

Figure 9.7: Simulated transmission parameters for: a single transition terminated by an on-chip short, a sapphire chip with no metallization, and the structure with the chip removed. The operating band is highlighted in gray.

We have thus far treated the system of two back-to-back taper transitions as a single unified network described by a set of S-parameters. However to verify the effectiveness of the transition with even greater precision, the transitions themselves could be de-embedded by defining an error sub-network inside the grey box in Figure 9.4c. This could help with getting more accurate measurements of more complex on-chip devices, but in particular the isolation term of this sub-network (equivalent to ϵ_{30} in Figure 9.4c) could shed some light on how much signal passes through the slotline and reaches the resonators as opposed to bypassing the chip entirely. With our measurement precision, the correction method

described earlier already gives fairly high uncertainty without introducing these extra error parameters, so we are unable to directly measure this leakage with good accuracy.

However, we can use simulations (which show fairly good agreement with the measured responses) to estimate the relative magnitude of the leakage. In Figure 9.7, the solid blue line shows the simulated transmission of a chip with a tapered transition on one side, and entirely covered by uniform ground plane on the other. This measurement corresponds to the *Reflect* standard used in the TRL calibration [184], and directly measures the isolation error term. Thus we find that in the band of interest for our system (highlighted in Figure 9.7), the simulated leakage is below approximately 38 dB: this value results in a total reduction between 0.05 and 0.1 dB on our measured insertion loss: significantly lower than the uncertainty from calibration discussed in the previous section.

This leakage term is small enough to ignore for resonator measurements, however sheds light on imperfections in this design. To estimate the leakage origin, we also simulate transmission of the geometry with a bare sapphire chip containing no metal, as well as transmission without the chip entirely, shown with dashed lines in Figure 9.7. From this we conclude that the bare sapphire chip supports modes which help transmit significantly more signal than just the copper enclosure geometry itself. While these spurious modes are suppressed by the surface metal (demonstrated by reduced transmission of the short) a more careful examination could help further improve the transition design.

9.4 Ground-Shielded Resonator Design

Having demonstrated a coupling structure capable of efficiently transforming a rectangular waveguide signal to and from a localized on-chip slotline, we can now design resonators decoupled from their environment while only considering local interactions. Away from the centered slotline, the chip surface is entirely covered by superconducting material acting as a ground plane, so a resonant structure patterned in this region will have its long-range

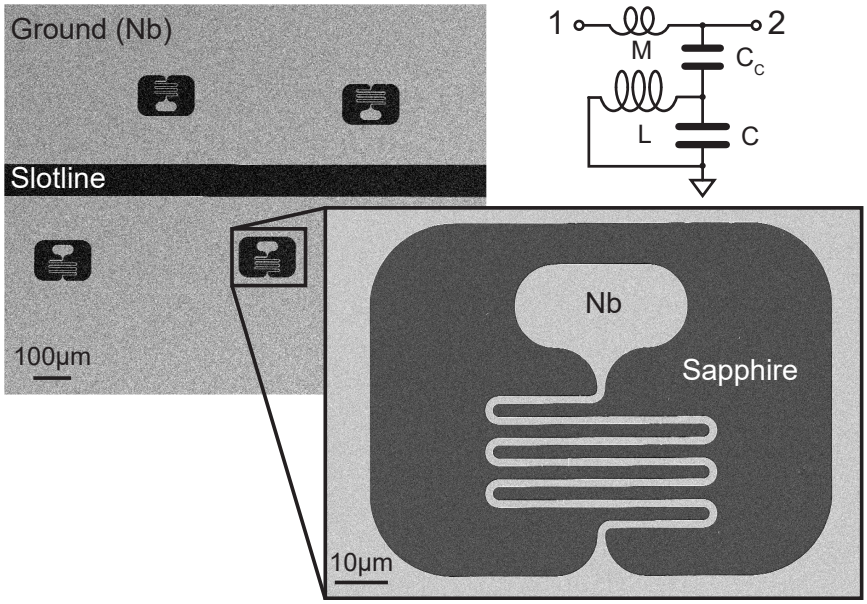

Figure 9.8: Electron micrograph of resonator geometry and coupling arrangement relative to feedline. This structure can be approximated by a simplified circuit diagram (top-right).

electric dipole interactions reduced. Our millimeter-wave resonator design is composed of a discrete capacitor island connected to the ground plane by a 2 μm-wide meandered inductor, shown in Figure 9.8. Every resonator is designed with an identical capacitor island, and its resonant frequency ω_0 is adjusted by changing the inductor length while keeping its width constant. The entire resonator has a rectangular footprint around $100 - 200$ μm per side, which is not insignificant compared to signal wavelength in the slotline ($\sim$ 1 mm): in this limit, the ground plane edges contribute significant reactive corrections. On each chip, five to six resonators are placed near the central slotline to allow interaction with the propagating signal. This differential geometry can be modelled by the single-ended LC circuit shown in Figure 9.8 including asymmetric coupling (C_C and M) [106]. This accounts for reactive contributions and impedance mismatches induced by the resonator presence.

For this resonator geometry, coupling strength can be controlled by adjusting its separation from the slotline. The interaction decreases exponentially with distance, which we demonstrate by simulating the coupling quality factor Q_e as a function of resonator sepa-

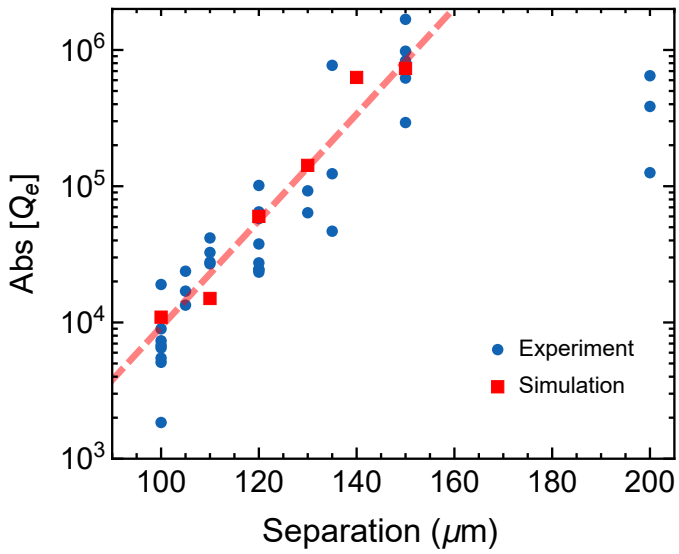

Figure 9.9: Simulated (red) resonator coupling Q_e as a function of separation from the slotline with an empirical fit (red line) used for predictions. Experimental measurements of Q_e (blue) are in reasonable agreement.

ration, plotted in red in Figure 9.9. The coupling quality factor can also be approximated empirically as a function of only the separation d by $\log_{10} Q_e = 0.06491 + 0.0390 d/\mu\text{m}$, shown as a dashed line in Figure 9.9. We find that experimental measurements of Q_e (described in the next section) follow this approximation reasonably well.

9.5 Millimeter-wave Resonator Measurements

Characterizing the complex transmission spectra of these resonators at low temperatures ($T = 0.86$ K) allows us to explore losses at millimeter-wave frequencies. Typical normalized measurements at low average photon number ($\bar{n} \approx 10$) are shown in Figure 9.10a. On resonance, we observe a dip in transmission as the resonance sweeps through a circle in the complex plane. This behavior is captured well by [106]:

$$S_{21} = 1 - \frac{Q}{Q_e^*} \frac{e^{i\phi}}{1 + 2iQ\frac{\omega - \omega_0}{\omega_0}} \tag{9.4}$$

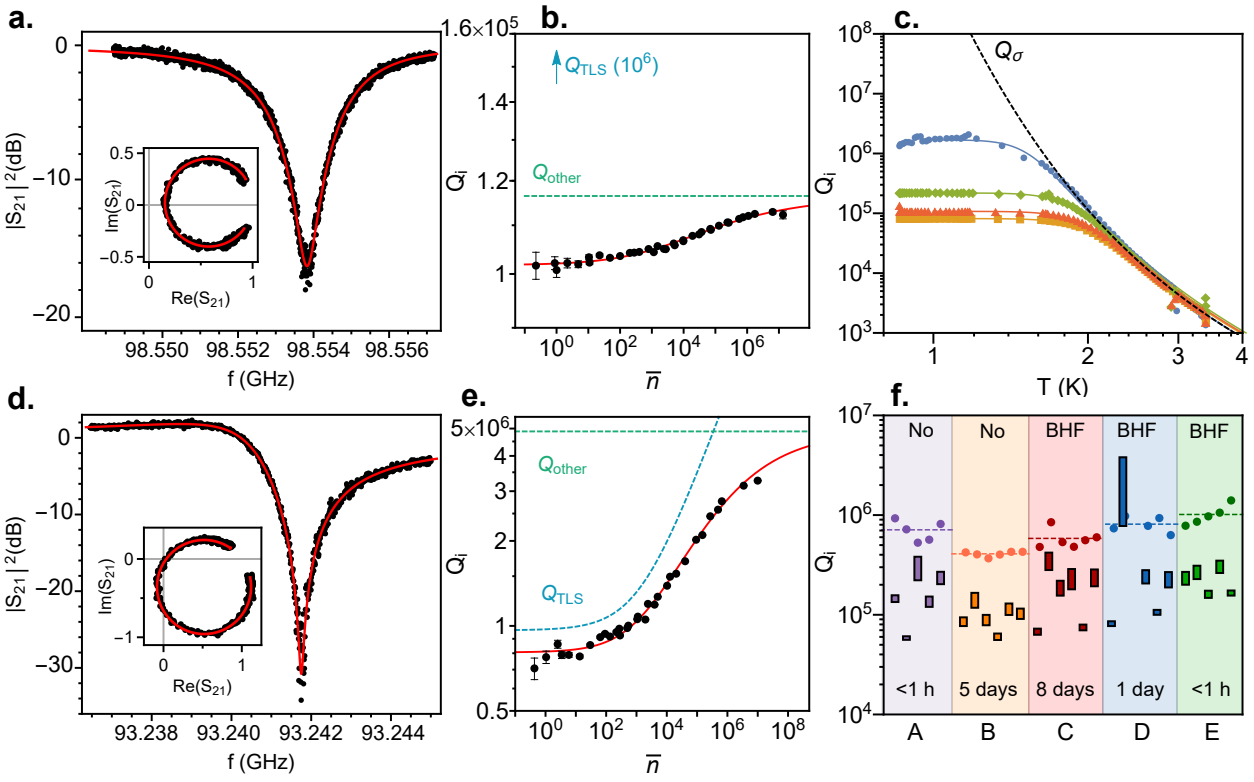

Figure 9.10: **a-b)** Complex transmission spectrum of a typical resonator, and power dependence of its internal quality factor along with fits to a model including TLS and independent loss (red). Here, Q_i is primarily limited by non-TLS loss (Q_{other}). **c)** Temperature dependence of Q_i for resonators from chip D. The black dashed line corresponds to a Bardeen-Cooper-Schrieffer (BCS) model of conductivity loss, and solid lines are respective fits to a model including conductivity and TLS loss. **d-e)** Complex transmission spectrum and Q_i power dependence of the best resonator measured in this study. For this device, Q_{TLS} is the dominant loss source. **f)** Internal quality factors for resonators in this study, grouped by etching conditions and elapsed time after fabrication. The top and bottom of the colored bars correspond to measured low-power and high-power limits of Q_i, and the points correspond to TLS induced loss $Q_{\text{TLS},0}$ with averages for each chip denoted by a dashed line.

where $Q^{-1} = Q_i^{-1} + \text{Re}[Q_e^{-1}]$ [106] and the coupling quality factor Q_e is rotated in the complex plane by $Q_e = Q_e^* e^{-i\phi}$ due to asymmetric coupling to the slotline as described in the previous section. Measuring both quadratures of the transmission spectrum to capture this asymmetry is particularly important for extracting an accurate estimate of Q_i, which is sensitive to ϕ [186].

Repeating these measurements at varying powers shows that Q_i increases with power, as

shown in Figure 9.10b. This increase can be explained by a power-dependent loss mechanism from saturating TLSs [70, 179, 195, 242] described by:

$$Q_{\mathrm{TLS}}(\bar{n}, T) = \frac{Q_{\mathrm{TLS},0}}{\tanh \frac{\hbar\omega}{kT}} \sqrt{1 + \left(\frac{\bar{n}}{n_c}\right)^{\beta} \tanh \frac{\hbar\omega}{kT}} \tag{9.5}$$

Here $Q_{\mathrm{TLS},0}$ is the inverse linear absorption from TLSs, ω is the resonant frequency, and β and n_c are parameters characterizing TLS saturation [70, 195]. Nonlinear effects (from kinetic inductance) [12] limit the power range where linear measurements can be performed, however at high powers we observe that Q_i begins to saturate, indicating the presence of other loss mechanisms.

By examining the power and temperature dependence of Q_i, we can further distinguish between sources of loss. The full behavior is captured with a model that includes TLS loss (Q_{TLS}) [70, 179, 195, 242], equilibrium quasiparticle loss (Q_σ) [145, 190, 224] and other loss mechanisms that are power and temperature-independent (Q_{other}):

$$\frac{1}{Q_i(T, \bar{n})} = \frac{1}{Q_{\mathrm{TLS}}(\bar{n}, T)} + \frac{1}{Q_\sigma(T)} + \frac{1}{Q_{\mathrm{other}}} \tag{9.6}$$

The quasiparticle loss term is parameterized by:

$$Q_\sigma(T) = Q_{\sigma_0} \frac{\sigma_2(T, T_c)}{\sigma_1(T, Tc)} \tag{9.7}$$

where σ_1 and σ_2 respectively are the real and imaginary parts of the complex surface conductance, calculated by numerically integrating the Mattis-Bardeen equations for σ_1/σ_n and σ_2/σ_n, and σ_n is the normal conductance. [145, 190, 224].

To investigate the effects of quasiparticle loss, we measure Q_i of several representative resonators as a function of temperature and show the results in Figure 9.10c along with $Q_\sigma(T)$ and fits to the full model described above. Below approximately 1.5 K we find that

Q_i is almost unaffected by Q_σ and is nearly temperature independent for most devices. As the temperature approaches a significant fraction of niobium's critical temperature (9.2 K), Q_σ rapidly becomes dominant. However, at the low temperatures relevant for quantum experiments, Q_σ exceeds measured values of Q_i by several orders of magnitude, suggesting that quasiparticle loss contributions are negligible in this regime.

Having determined that the dominant loss contributions come from power-independent loss Q_{other} and Q_{TLS}, we can neglect thermal contributions to Q_i at low temperatures. Upon inspection of the power dependence of a typical resonator in Figure 9.10b, we find the increase of Q_i from TLS saturation is relatively small, unlike what is seen in many microwave loss studies [49, 236]. Using the model above, we find $Q_{\text{TLS},0} = 0.953 \times 10^6$ while $Q_{\text{other}} = 1.17 \times 10^5$, indicating that TLSs are not the dominant source of loss. Repeating these measurements for the highest-Q resonator in this study, we find in Figure 9.10d that Equation 9.4 is able to accurately capture Q_i from the response, but unlike the other resonators, the power dependence shown in Figure 9.10e is much more pronounced. For this device, we find single-photon $Q_i = 0.827 \times 10^6$, which is comparable to state-of-the-art microwave resonators [49, 236], and the loss sources can be disentangled into $Q_{\text{TLS},0} = 1.03 \times 10^6$ and $Q_{\text{other}} = 4.18 \times 10^6$. Unlike the device above, this independent loss limit is significantly higher and the resonator loss primarily arises from Q_{TLS}, giving us a better insight into millimeter-wave TLS loss.

9.5.1 Reducing Losses With Surface Oxide Etch

When exposed to air, niobium is known to slowly evolve a lossy amorphous surface oxide layer containing dissipative sub-oxides [41] and TLSs [236]. Surface treatments are commonly used to remove this surface layer in microwave resonators to reduce loss [236]. Although TLS density has not yet been investigated in the W-band, losses from niobium sub-oxides are believed to be more pronounced at higher frequencies [41]. To study the effect of surface

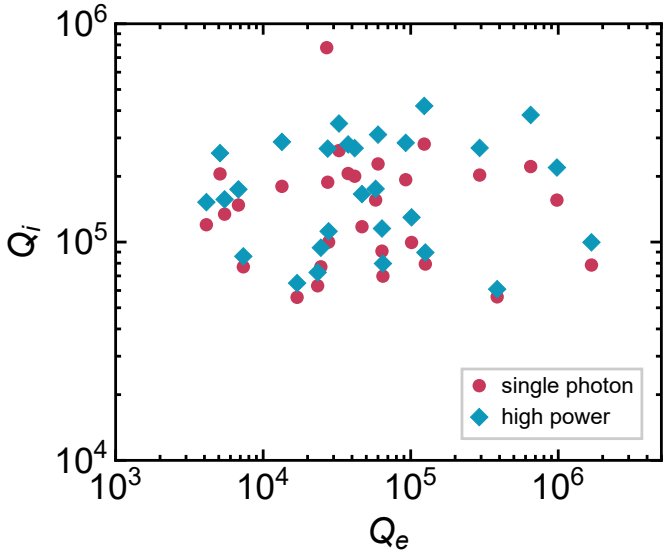

Figure 9.11: Resonator loss (internal Q) for single-photon (red) and high-powers (blue) as a function of coupling Q, showing no correlation and confirming coupling does not increase loss.

processing on Q_i, we repeat the measurements summarized in Fig. 9.10a-b for devices that underwent different aging times and etch conditions. In Fig. 9.10f, we plot the low and high power limits of measured Q_i as well as the fitted value of $Q_{\mathrm{TLS},0}$ for devices from five separate chips.

Between samples A and B we observe that 5 days of aging reduces both $Q_{\mathrm{TLS},0}$ and Q_{other}, leading to lower quality factors for resonators exposed to air for several days, which is consistent with niobium oxide regrowth that has been shown to increase loss in microwave devices [236]. This can be mitigated by selectively removing the surface layer of niobium oxide after fabrication using a buffered solution of hydrogen fluoride (BHF) [236], which we achieve by immersing samples C, D and E in a 5% BHF solution for 40 min immediately prior to mounting and measurement.

We observe that this BHF treatment can reverse the effects of air exposure in Sample C, which despite experiencing a longer air exposure of 8 days, has higher average $Q_{\mathrm{TLS},0}$ and Q_{other} than Sample B after the BHF treatment. Applying this surface treatment to samples D and E, which experienced reduced initial air exposure yield consistently lower losses than

samples A-C, including the highest-Q resonator in this study described above. Combining minimal air exposure with surface oxide removal using BHF in sample E, we are able to consistently obtain resonators with single photon internal quality factors above 1.4×10^5, and an average $Q_{\mathrm{TLS},0} = 1.04 \times 10^6$: significantly higher than previously measured for planar millimeter-wave devices [12, 63, 71, 205, 215, 232].

While these values of loss are much closer to those reported for microwave devices [236], in the power-dependent measurements above we observe that our millimeter-wave resonators are on average limited by power and temperature-independent loss Q_{other} to a much greater extent than TLS loss, which limits microwave devices [49, 236]. This loss could come from a variety of sources, including remnants of conduction loss from the copper enclosure [90], seam loss from an imperfect seal between the halves of the enclosure [33], radiation loss [39] or additional power-independent microscopic relaxation channels such as conductive loss in the niobium sub-oxides [236]. To estimate the impact of radiation losses, we can verify that our ground-shielded resonator design protects the resonance from radiative loss induced by coupling. In Figure 9.11 we plot Q_i for single-photon and high-power limits (Q_{other}) as a function of Q_e, and observe no correlation with either. Thus we can design a resonant circuit with a wide range of coupling strengths without affecting coupling to lossy or radiative channels. As a final note, an examination of all the devices in Figure 9.10f, shows that on average $Q_{\mathrm{TLS},0}$ and Q_{other} scale similarly relative to each other when affected by aging and surface treatment. This is highly suggestive that the remaining millimeter-wave source of loss is still tied to materials-induced decoherence in the superconductor surface, and warrants further studies.

9.6 Conclusion

We have demonstrated an on-chip millimeter-wave resonator design with a ten-fold improvement in loss over previous work, and leveraged this platform to investigate sources

of single-photon decoherence in the W band. Using a specifically-designed waveguide to slotline transition based on a finline taper, we present a new platform for probing on-chip devices at high frequencies. With this novel packaging design, we have improved on dipole coupling techniques previously used to address millimeter-wave resonators, and show that our coupling can be adjusted over a wide range without impacting circuit losses. Having shown that planar millimeter-wave resonators compatible with fabrication techniques can achieve performance comparable to microwave quantum circuits [49, 236], we pave the way for a millimeter-wave artificial atom that could be formed by introducing a high-frequency nonlinearity [12, 13, 109] into our design. This opens the door to a new generation of high-frequency quantum tools compatible with higher operating temperatures.

CHAPTER 10

MILLIMETER-WAVE QUBITS

The time has come to finally unite the various separate pieces we developed, and realize our overall goal of a millimeter-wave superconducting atom. In Chapter 7 we developed a high-temperature trilayer Josephson junction: a strong source of low-loss single-photon nonlinearity capable of operating at high temperatures and frequencies. We took a small step towards higher operating frequencies in Chapter 8 by scaling up our trilayer qubit design to the K band. Finally in Chapter 9 we demonstrate a millimeter-wave coupling method and ground-shielded resonator design which improves over our previous measurements described in Chapter 5. The logical next step is to combine the now proven high-frequency junctions with the upgraded millimeter-wave package.

10.1 The Millimeter Wave Qubit

After going through several iterations, our final millimeter-wave qubit design is pictured in Figure 10.1. Five to six resonators are coupled to the central slotline region from the tapered finline coupling structure described in Chapter 9. The design freedom afforded by the location-dependent coupling (provided by compressing the signal in the center of the chip) allows us to ensure that the resonators are strongly coupled to the slotline ($Q_e \sim 10^3$) while the qubits aren't directly coupled (we estimate the residual coupling is above $Q_e > 10^6$, which as we will see later does not limit qubit performance). Notably, unlike in Chapter 9, the qubit resonators are inductively coupled, as opposed to capacitively coupled[1]. For this "jellyfish" resonator design, inductive coupling scales with the overlap of current distributions between the resonant mode and the propagating slotline mode, which results in about ten times stronger coupling for the same slotline separation. This also has the benefit of making

1. Switching the dominant coupling type is achieved by simply rotating the design by 180°

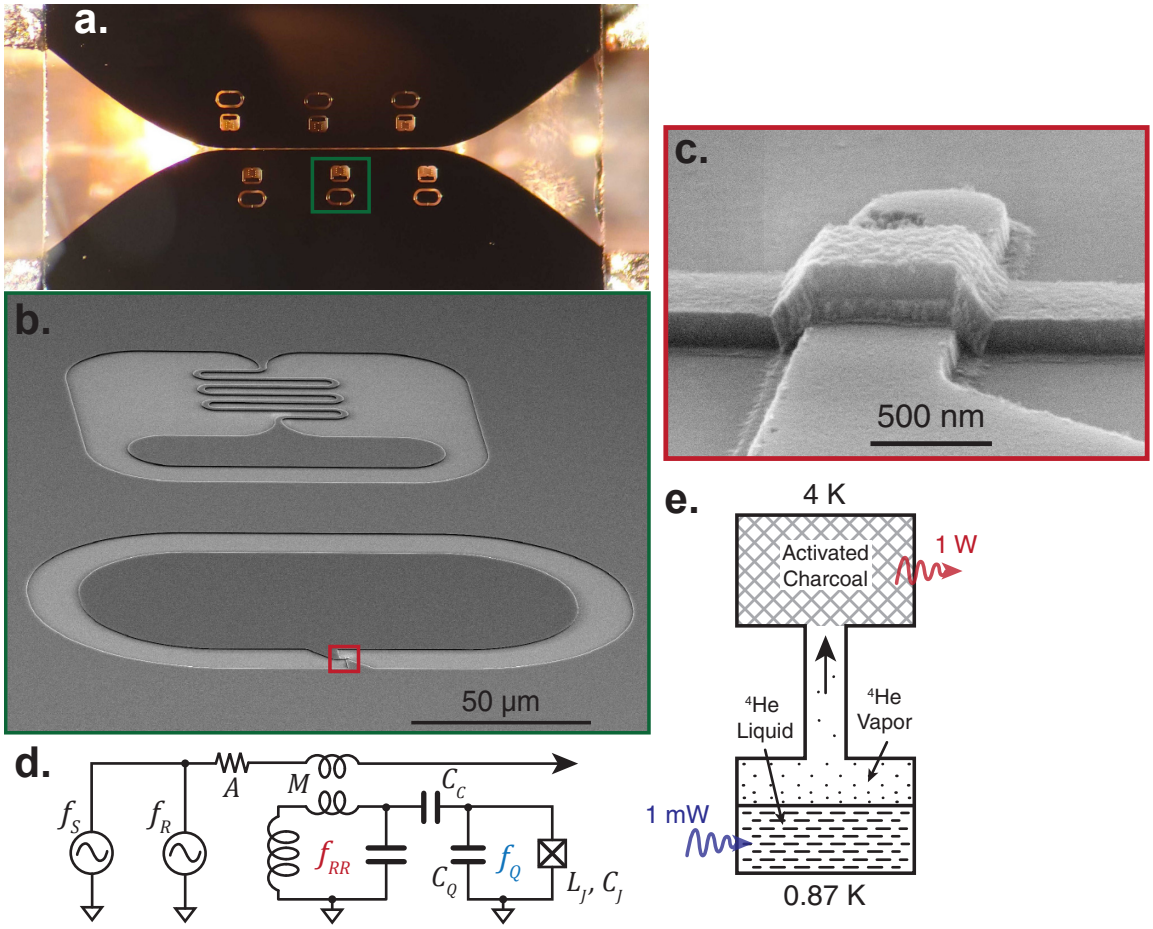

Figure 10.1: Millimeter-wave qubit geometry. **a)** Photograph of several qubits and readout resonators coupled to the tapered finline transmission structure used for measurement (see previous Chapter 9 for details). **b)** Scanning electron micrograph of a readout resonator (top) which is capacitively coupled to the qubit capacitor (bottom) **c)** A close up image of the niobium trilayer Josephson junction at the heart of the qubit. **d)** Equivalent circuit representation of the experiment along with the refrigeration mechanism shown in **e**

the capacitive end of the resonator available for coupling to the qubit capacitor.

For circuits, high frequencies are achieved primarily by reducing capacitance and inductance[2] The resulting qubit capacitor, expected to provide 38 fF (from simulations), is relatively small compared to conventional qubits (or even K band qubits). The capacitor island is simply shaped like a jellybean (or hotdog). While the electric field is nominally

2. Frequently either capacitance or inductance can be kept somewhat larger at the expense of the other, but as a rule of thumb both will need to decrease.

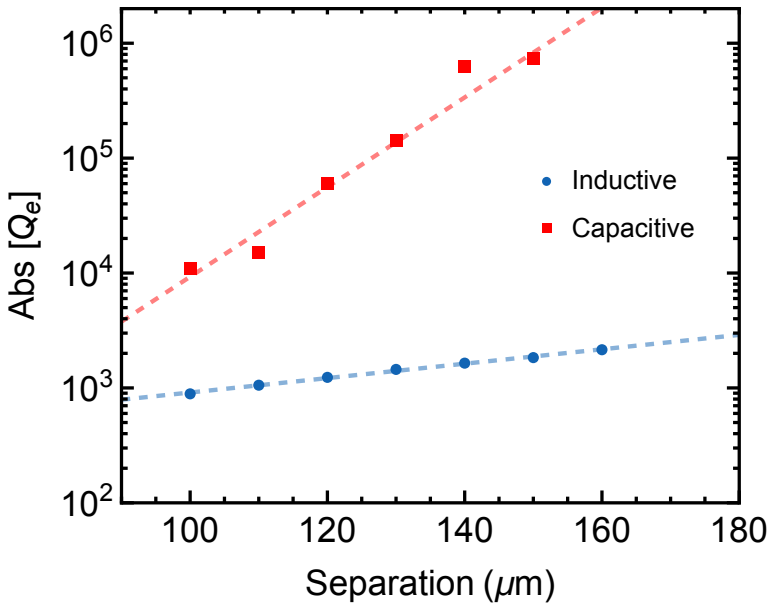

Figure 10.2: Simulated resonator coupling Qe as a function of separation from the slotline with a empirical fits (dashed lines) used for predictions. Resonators coupled capacitively (blue) similar to those measured in Chapter 9 have orders of magnitude weaker coupling which is a strong function of separation, while inductively-coupled (red) resonators (flipped by 180°) couple much more strongly.

distributed around the ground plane, the junction (shown in Figure 10.1c) is located on the opposite side of the resonator, as illustrated in Figure 10.1b. This helps promote capacitive coupling by moving inductive contributions of the ground plane away from the resonator. The entire system can be described by the circuit illustrated in Figure 10.1d. This circuit is essentially identical to the lumped-element description of the microwave and K-band qubits measured in Chapters 7 and 8.

10.1.1 Helium-4 Experiment Refrigeration

Other than the millimeter-wave frequencies and mounting method, the primary difference of this experiment is the cryogenic cooling method. Whereas the lower frequency experiments described thus far have been cooled with dilution refrigerators, the millimeter-wave experiments are all cooled by a custom-built helium-4 adsorption refrigerator, whose oper-

ating principle is summarized in Figure 10.1e. At a base temperature of 0.87 K, evaporating helium-4 provides approximately 1 mW of cooling power– which is significantly higher than a dilution refrigerator [183], and one of the major strengths of this high-frequency superconducting platform.

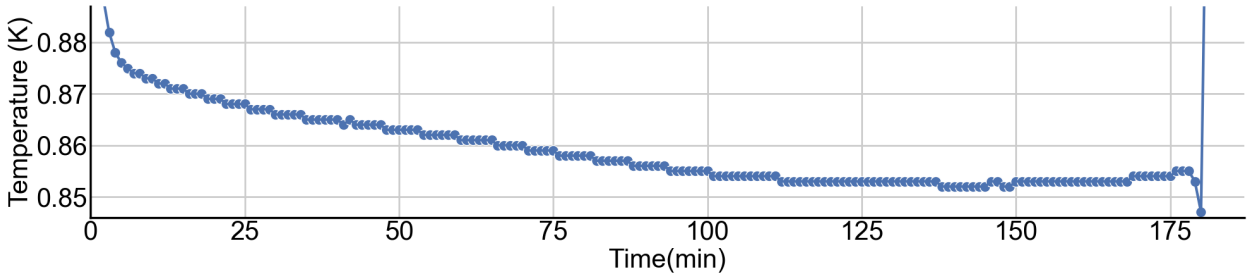

Figure 10.3: Experiment temperature during a helium cycle. For maximum stability, data is taken below 0.868 K of the lowest temperature.

The evaporation is sustained by adsorbing the helium vapor on the extremely high surface area of activated charcoal, which is thermalized at 4 K. This cryogenic setup has a significant advantage in its design and construction simplicity. However we note that a significant drawback of this refrigeration method is the limited hold time, which in our system is approximately three hours: after this time the condensed helium fully evaporates, and the system must be reset by heating the charcoal, which warms the experiment up to 4 K. The experiment temperature during one of these cycles is shown in Figure 10.3. Because of this, any millimeter-wave data acquisition is only performed during one of the 3 hour cycles. We verify this by checking whether the experiment temperature is within 10 mK of the lowest temperature before saving data. Notably this does not prevent averaged experiments longer than 3 hours so long as each individual experiment can be executed within the cycle duration.

10.2 Continuous Wave Measurements

As described earlier, several qubit–readout structures are patterned on a chip. To investigate the system, we first use single-tone millimeter-wave spectroscopy (see Chapter 9 and 5) to

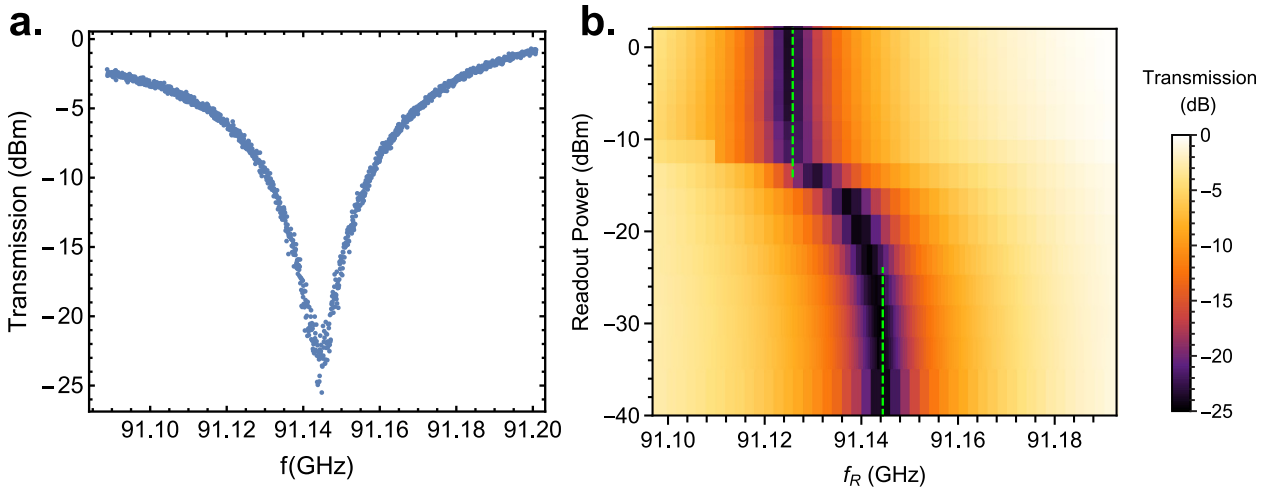

Figure 10.4: Power dependence of the readout resonator. (a) Transmission measurement of the readout resonator, taken at low power. At the resonant frequency (which is modified by the coupled qubit), we observe a dip in transmission. (b) Since the readout resonator is coupled to a highly nonlinear qubit, it will inherit some nonlinearity as well. Because of this, the readout resonator shifts down in frequency as power is increased. At sufficiently high powers, the readout resonator returns to its rest frequency. Measuring the frequency difference between these low and high power resonant frequencies (green dashed lines) allows us to determine the qubit-readout coupling g^2/Δ.

measure the complex transmission spectrum through the central slotline. Near the readout resonator resonance frequency this results in a dip in transmission as shown in Figure 10.4a, which we can fit to the side-coupled asymmetric[3] model described in Chapter 3. At low powers, the readout resonator frequency is dressed by its interactions with the qubit [25, 113, 192]. Because the readout resonator interacts with the highly nonlinear qubit, it will inherit[4] a small amount of nonlinearity itself [192]. In Figure 10.4b we plot transmission through the slotline near the resonator frequency as a function of signal power. For increasing powers, we observe the characteristic frequency shift of a Kerr-nonlinear system, confirming the resonator is coupled to something nonlinear. At sufficiently high drive strengths, sufficiently

3. Luckily the improved coupling method results in less asymmetry, however to capture the full pciture it is best to use the asymmetric model anyways.

4. As another way to think about this, the junction participates to varying degrees in every mode of the circuit [155]. Modes with higher participation will have stronger nonlinearity.

high powers in the qubit decouple it from the readout resonator, which returns[5] to its bare resonant frequency [25, 178, 192]. The frequency difference between the dressed and bare resonant frequencies is determined by the bare coupling qubit–resonator strength g [25]:

$$\delta\omega = \frac{g^2}{\Delta} = \frac{g^2}{\omega_q - \omega_{RR}} \tag{10.1}$$

This is a particularly useful measurement since it helps us determine the sign of the qubit-resonator detuning[6]. Furthermore, once the qubit transition frequency ω_q is known, we can use the frequency shift $\delta\omega$ to determine the raw qubit–resonator coupling strength g. For this coupling design, we typically find coupling strengths between 400–1500 MHz. This coupling strength is set by the coupling capacitance C_c, which depends strongly on the physical separation of the qubit and resonator capacitor islands (see Figure 10.1b).

10.2.1 Identifying Qubits with Two Tone Spectroscopy

While monitoring transmission through the system at the low-power readout resonator frequency f_R, simultaneously applying a second probe signal f_S should reveal the excitation spectrum of the coupled qubit. In practice, since the readout resonator has inherited nonlinearity [192], and there is non-negligible cross-coupling between neighboring resonators, we find that the system can be excited through several additional pathways, leading to a relatively complicated spectrum as shown in Figure 10.6. This spectrum is also complicated by our use of harmonic multipliers to generate both the readout signal f_R and probe signal f_S.

Because our heterodyne measurement setup uses a harmonically multiplied local oscillator (LO) signal, the down-mixed readout measurement is sensitive to several millimeter-wave

5. We call this a "snap".

6. As we will see in a bit, demystifying every transition for millimeter-waves is a challenge. In an earlier sample, the positive snap from a qubit above its readout resonator frequency helped us determine with certainty which features were from the qubit as opposed to harmonics.

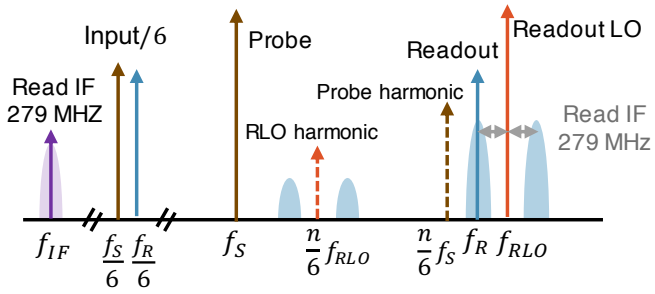

Figure 10.5: Frequency breakdown of the signals used to measure the qubit with continuous-wave spectroscopy. Microwave input signals are harmonically multiplied to the millimeter-wave band. The blue regions are down-mixed to the idler frequency measurement window (purple).

frequencies f_{Meas}, all of which can down-mix to the same idler frequency (IF) as illustrated in Figure 10.5.

$$f_{\text{Meas}}^{n,m} = \frac{n}{6} f_{RLO} \pm m f_{IF} \tag{10.2}$$

Luckily for us however, the mixer is highly power-dependent, so is predominantly sensitive the fundamental $m = 1$,$n = 6$ frequency. Additionally, since the readout signal and local oscillator signal are phase locked[7], simply averaging a few measurements together will remove most of the interference from a spurious signal getting down-mixed into the measurement band, which will inherently have a random relative phase.

The probe signal f_S also generates harmonics (n) from the multiplication process, which will have the freuqencies

$$f_{S,n} = \frac{n}{6} f_S \tag{10.3}$$

Many of the lower frequency harmonics are suppressed by the natural cutoff frequency of the rectangular waveguide (see Chapter 3), however depending on the probe frequency the $n = 4, 5, 7, 8, 9$ harmonics can be prominent. As a result, when monitoring transmission through a readout resonator, we will see a deflection signal when any harmonic of the probe

7. They are actually generated from different combinations of the same 10 MHZ signal supplied by a rubidium clock, so they have the same fundamental phase. This relative phase stability is very important since it lets us average coherently.

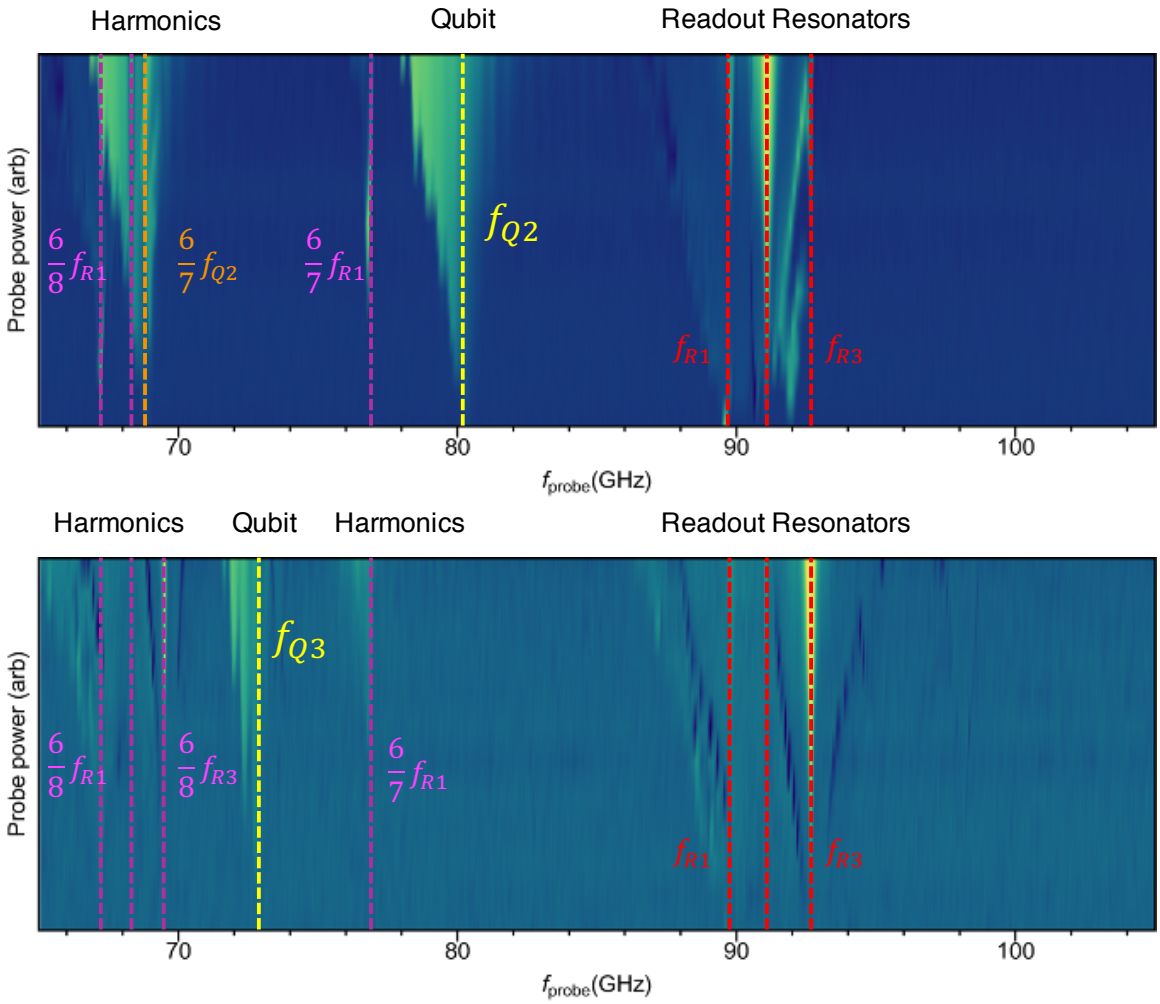

Figure 10.6: Two tone spectroscopic identification of qubit transitions. For each readout resonator, we measure transmission at its low-power resonance, while simultaneously applying a second probe signal at varying frequency and power. Here we are showing linear deflection from the idle state (lighter colors are larger deflection). When the fundamental harmonic or any other harmonic is resonant with a qubit or resonator the readout resonator measured will shift, registering a deflection. By measuring all resonators (two examples shown here) we can identify which features are harmonic excitations, and which are fundamental excitations, which locates the qubit transitions.

signal matches either the resonator frequency or the coupled qubit frequency:

$$f_S = \begin{cases} \frac{6}{n} f_{RR_i} \\ \frac{6}{n} f_{Q_j} \end{cases} \tag{10.4}$$

Allowing for cross-coupling, when measuring a given readout resonator we could also expect signals from its neighboring readout resonators RR_i, and potentially even neighboring qubits Q_j (though in practice cross-coupling $i \neq j$ between mismatched qubits and resonators is significantly weaker).

With these rules established, the complicated two tone spectrum shown in Figure 10.6 becomes a simple logical puzzle, and each individual transition can be identified. Importantly, with the resonant and harmonic excitations eliminated, we can deduce the fundamental qubit transition frequency. Combined with the detuning measurement measured in Figure 10.4, this also identifies the bare qubit–readout resonator coupling strength g.

10.2.2 Resolved Qubit Transitions

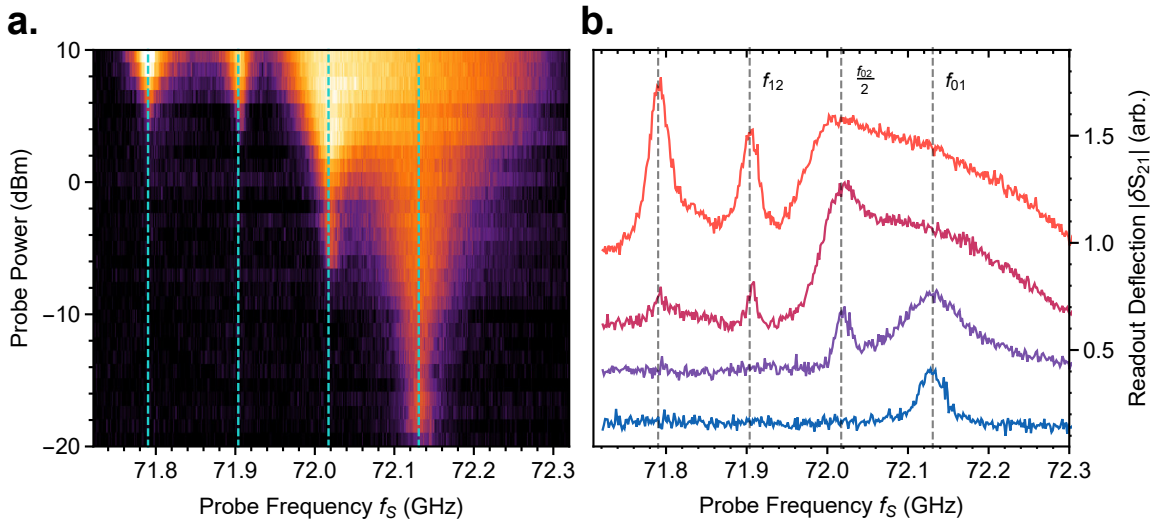

Figure 10.7: Resolved level transitions of the millimeter-wave qubit. **a)** Deflection of the readout resonator transmission signal as a function of applied qubit pulse frequency, shown for increasing qubit pulse power. Line cuts for indicative qubit drive powers are shown in **b)**. At low powers (blue) a single peak is observed when the pulse is resonant with the qubit frequency ($f_{ge} = 72.1$ GHz). With higher probe powers, the linewidth of this transition increases, and additional peaks appear corresponding to higher energy level transitions through many-photon excitations ($f_{gf}/2$ etc). These features have a spacing of $\alpha/2 = (f_{ge} - f_{ef})/2$.

Having identified the approximate qubit frequencies, we can focus on the fundamental

transitions of a promising 72 GHz qubit for which the transition linewidth appears to be lower than the transition spacing. For added sensivity we make sure to consider the vector transmission displacement $|\delta S_{21}|$, which takes advantage of both quadratures of the measured readout transmission signal

$$|\delta S_{21}| = \left|\vec{S}_{21}(n_s) - \vec{S}_{21}(0)\right| \tag{10.5}$$

We show the measured deflection magnitude for probe frequencies near the qubit transitions, measured at increasing probe powers in Figure 10.7. At low powers, we observe a deflection signal when the probe matches the bare qubit frequency f_{ge}: exciting the qubit from its rest state induces a dispersive shift in the readout resonator frequency. Increasing the power of the probe pulse reveals the higher energy states of the qubit through two-photon processes $f_{gf}/2 = (f_{ge} + f_{ef})/2$ and excited state transitions f_{ef}. This allows us to measure the anharmonicity $\alpha \equiv f_{ef} - f_{ge}$, which for the millimeter-wave qubit is 227 MHz. As we will see later on, this anharmonicity is the Achilles foot of this qubit. Notably, this level spacing (which sets an upper bound on qubit gate speed) is still similar to our K band qubit, and many conventional microwave qubits [21, 178]. This is interesting, since the qubit capacitor in this case is significantly smaller than the K band qubit, which is in turn smaller than most microwave qubits. This is a result of the relatively large Josephson junction used in the qubit (see Figure 10.1c) which evidently contributes $C_J = 57$ fF of additional capacitance.

From this measurement of anharmonicity we also conclude that our qubits have quite high junction participation ratios $p_j \approx 0.6$. Recalling from Chapter 4 that the participation ratio is determined by the qubit frequency, this in turn confirms that the junction plasma frequency is approximately 120 GHz. With a slightly smaller sized qubit capacitor and Josephson junction, we could achieve qubit anharmonicity up to 500–1000 MHz while still keeping the Josephson junction size large enough to define using optical lithography.

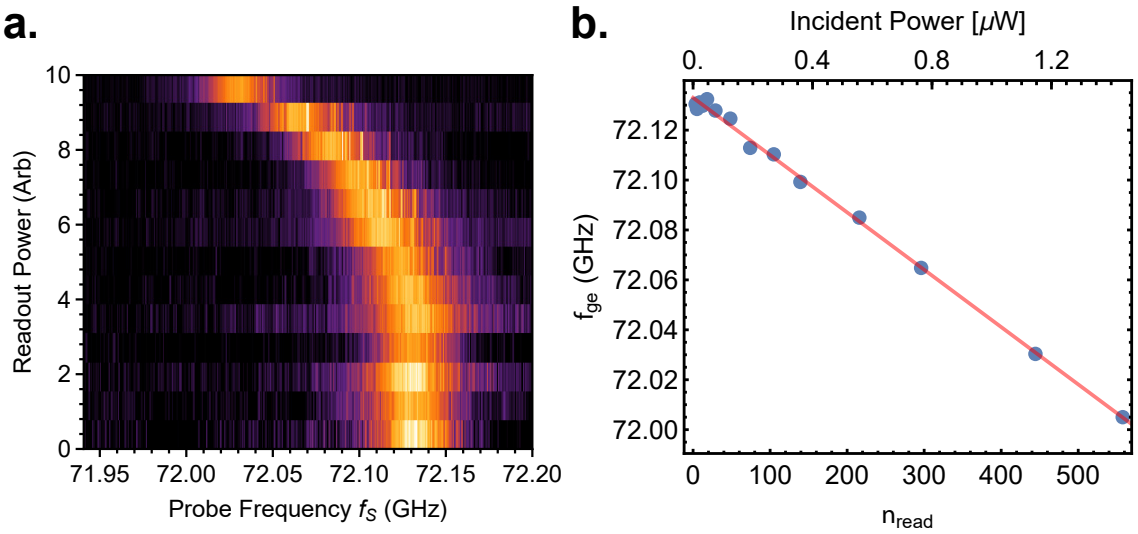

Figure 10.8: AC Stark shift. **a)**Deflection of the readout resonator transmission signal as a function of applied probe frequency, shown for increasing readout power (in dBm). Due to the resonator-qubit interactions, the qubit transition frequency shifts as the average readout resonator photon number increases. **b)** The qubit frequency shift is shown to be linear with respect to readout drive power. Since the frequency shift per photon can be calculated from qubit parameters, this allows us to obtain an accurate calibration for readout resonator photon number, and by extension the qubit photon number.

10.2.3 Drive Power Calibration via AC Stark Shift

We can still learn more from the dynamics of the two-tone qubit transition spectroscopy measured in Figure 10.7. In preparation, we will need an accurate estimate of the incident power arriving at the qubit. Luckily, the qubit itself provides an excellent method for calibrating power. Since we have established the qubit frequency, bare coupling strength, and resonator detuning, along with the resonator properties (linewidth, frequency), we can take advantage of the dispersive shift between them [113, 200], recalling that the qubit frequency is ac Stark shifted by the number of photons in the readout resonator:

$$f_{ge}(n_r) = f_{ge,0} - \chi n_r \tag{10.6}$$

To investigate this, we repeat the two-tone experiment as a function of probe frequency, for a low probe power (-18 dBm with respect[8] to the powers in Figure 10.7). The results are summarized in Figure 10.8. As expected, we observe that as the readout power increases the qubit transition f_{ge} decreases. Transforming the readout power into resonator photon number we verify that the freuqency shift is linear, and based on the dispersive shift calculated from known values, we can solve for the effective signal attenuation. Interestingly, the effective attenuation calculated with this method agrees with the values calculated using transmission measurements calibrated with a power meter[9]. With the exact shift calculated, we arrive at Figure 10.8b, in which the incident power and readout photon number are calibrated based on the qubit response. From this measurement combined with the frequency-dependent expression for photon number (see Chapter 3) we can also easily establish the number of qubit photons n_s for a particular probe drive strength.

10.2.4 Qubit Properties From Transition Linewidth

The probe signal used in Figure 10.7 to excite the qubit also induces dephasing dependent on the power of the applied power. The induced dephasing rate can be found by measuring the linewidth of the transition [201], which has an approximately gaussian profile at low powers. In Figure 10.9 we show the measured qubit transitions for two probe powers, along with gaussian fits capturing the linewidth of the transition σ. For higher probe powers, we also incorporate higher order transitions[10] to ensure the model accurately captures the linewidth of the first transition f_{ge}.

The linewidth of the transition is expected to scale with applied power [201] which is proportional to $n_s\omega_{\text{vac}}^2$. The vacuum Rabi frequency is simply g, so the linewidth will have

8. The exact probe power doesn't mean much since it goes through several manual attenuators.

9. The calculated and calibrated attenuation is much higher than we designed! I blame the stainless steel coaxial input line. Stay tuned for ScoobyDoo and the mystery of where the extra attenuation comes from!

10. The linewidth of the two-photon transition can also be modelled in a similar fashion, however will be half that of single-photon transitions.

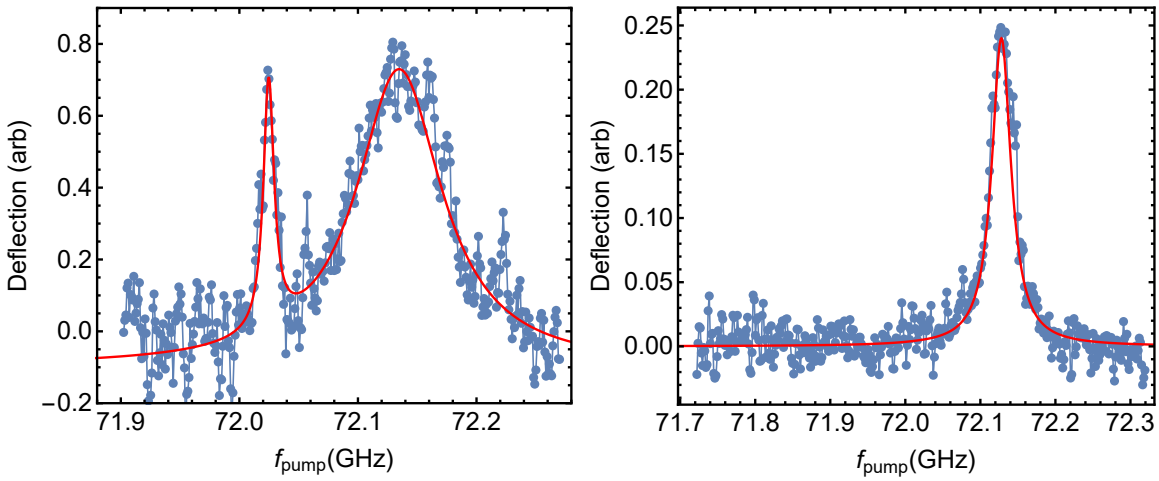

Figure 10.9: The feature corresponding to the ground–excited state transition in the readout resonator deflection measurements has a gaussian profile, allowing us to fit the linewidth of the transition. Measurements on the left and right are shown for drive strengths corresponding to 2.12×10^{-4} and 3.02×10^{-5} qubit photons respectively. For sufficiently high powers, the edges of the transition overlap with the next two-photon transition, so the model includes both of these transitions.

the following power dependence:

$$2\pi\sigma = \frac{1}{T_2'} = \sqrt{\frac{1}{T_2^2} + n_s g^2 \frac{T_1}{T_2}} \tag{10.7}$$

In this expression, T_2 is the dephasing rate of the qubit and T_1 is its relaxation rate. In Figure 10.10 we plot σ^2 as a function of the calibrated qubit photon number n_s, finding that the linewidth has the expected linear relationship. The intercept and slope allow us to estimate a dephasing rate of $T_2 = 20.89$ ns and relaxation time of $T_1 = 47.32$ ns respectively. It's worth mentioning that with this method, the dephasing time is only calculated from fitting the linewidth, however the relaxation time requires a sequence of additional measurements, making the relaxation time estimate significantly less accurate. As we will see soon, there are more reliable methods of measuring dephasing and relaxation times. So as not to overestimate the qubit performance, we will save the discussion of coherence properties till then.

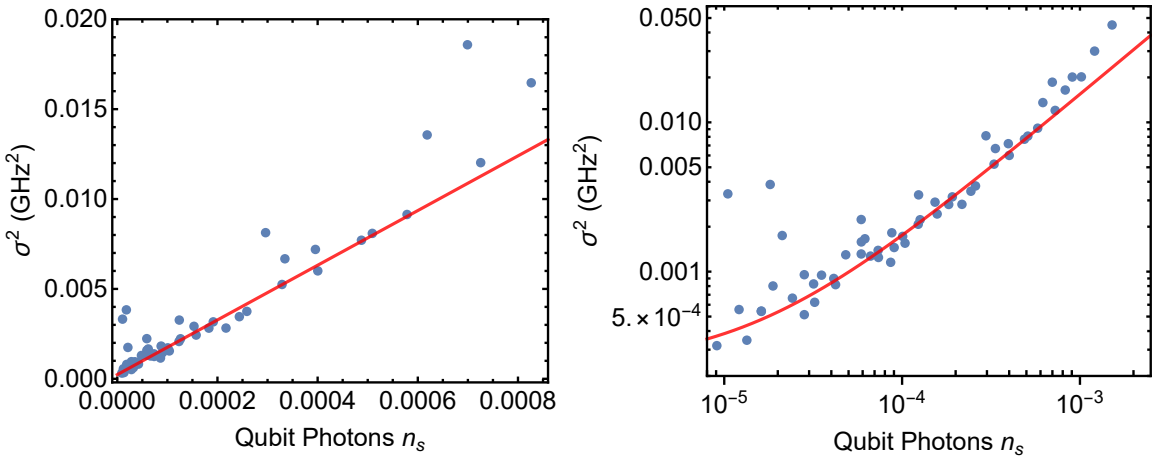

Figure 10.10: The fitted g–e transition linewidths σ are plotted as a function of drive photon number. The square of the linewidth is determined by the qubit dephasing time, and increases linearly with drive power, with a rate set by the qubit decoherence time. Having calibrated the photon number relative to drive power, a linear fit can extract a dephasing time of 21 ns and with lesser accuracy a decoherence time of 47 ns.

Table 10.1: Parameter summary for the millimeter-wave qubit (in MHz)

f_{RR}	γ_T	g^2/Δ	f_{ge}	Δ	α	g	χ
91,151.3	84.281	-19.44	72,137	-19,014.3	-228	607.979	-0.230342

10.3 Pulsed Qubit Characterization

We now turn to the time domain to demonstrate coherent control of the millimeter-wave qubit, and to obtain more accurate measurements of its coherent properties. Having estimated qubit decoherence times on the order of 20 ns using two-tone spectroscopy, we are presented with the challenge of generating sufficiently fast pulses to measure the qubit before it decoheres.

10.3.1 Sub Nanosecond Pulse Measurements

In order to take advantage of the phase-referenced quadrature measurements we have been using to measure the continuous-wave complex-valued transmission transmission through the fridge, we would like to largely keep readout measurement method the same. At the same

time, measuring the qubit state with a continuous signal will induce additional dephasing during state preparation and measurement operations. To address this we instead perform short duration pulsed measurements of the readout resonator dispersive shift, while only making minor modifications to the measurement hardware (namely the network analyzer, which can generate several phase-locked signals with different frequencies). This is accomplished with the addition of a microwave switch[11] to the microwave readout signal (F1) prior to multiplication, as shown in Figure 10.11. This effectively chops the millimeter-wave readout signal depending on a TTL signal, which can be generated from the same phase reference as the pulse frequency: as a result we will obtain a millimeter-wave pulse[12] with a consistent global phase. In practice, the pulse timing resolution of the network analyzer (10 ns) and the rise time of the pin diode microwave switch (around 6 ns) means that we can practically generate 20 ns readout pulses (which are shortened to a full-width half-maximum of approximately 14 ns). While sufficient to be able to resolve qubit dynamics, this pulse length is unfortunately comparable to the qubit dephasing time, so this measurement method is a good place for future improvement.

The upconverted readout pulse is sampled[13] internally in the VNAx Tx module using a copy of the same LO signal (F2) used for down-mixing the amplified measurement signal that comes out of the fridge. By measuring the relative phase of the signal that passed through the fridge with respect to the measured referenced signal, we are able to remove the effects of global phase drift and fluctuations.

Millimeter-wave multipliers and mixers are inherently nonlinear devices [184] so while the output power can be controlled by varying the input power, this scaling is nonlinear, and

11. Based on the Hatridge lab design

12. You can convince yourself that taking the 6th power of a square pulse produces a square-ish 6th harmonic.

13. The directional coupler orientation is portrayed accurately: since the reference port mixer in our module was not functioning correctly, we are technically sampling the reflected signal from the voltage controlled attenuator, which is significant.

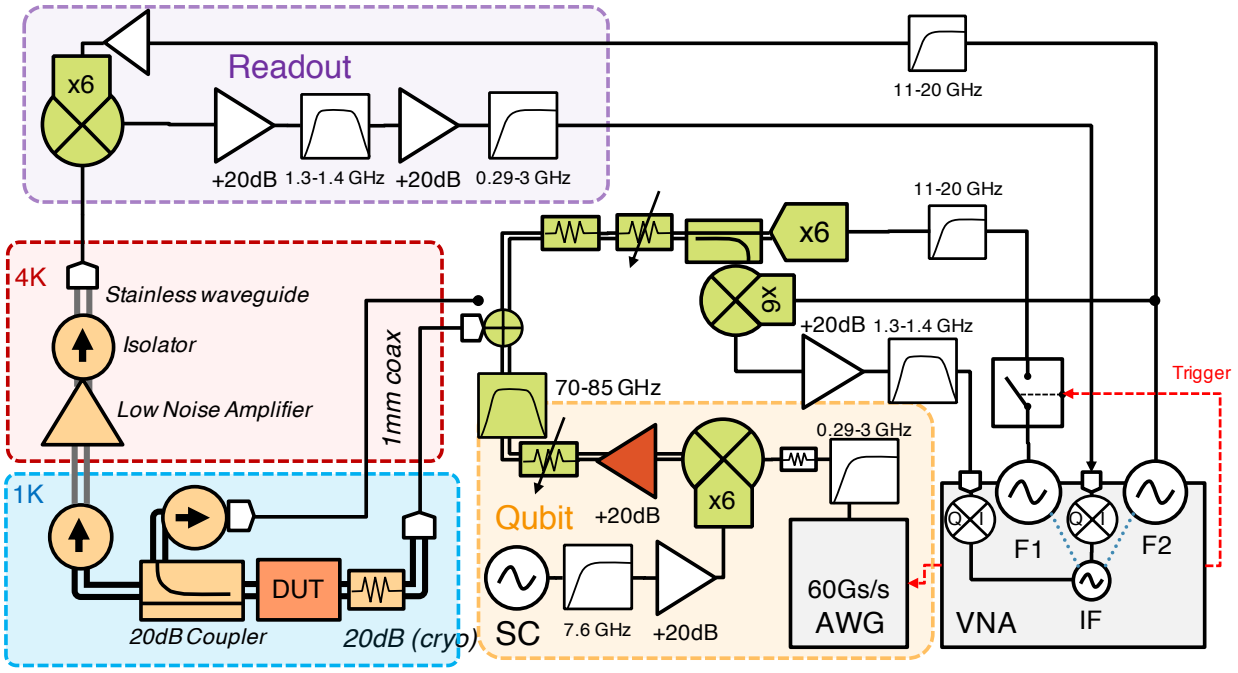

Figure 10.11: Millimeter-wave pulsed measurement setup. Double lines indicate waveguide regions. Red and blue regions highlight they cryogenic sections of the experiment. Components highlighted in green, red and orange are millimeter-wave components.

affects the harmonic power in a different way, making the resulting noise power generated by the system unpredictable. A much more reliable method is to use fixed input powers while varying attenuation on the millimeter-wave side, which should give more consistent results. To achieve the large dynamic range of powers needed for our experiments with some degree of automation, we use a combination of manual vane attenuators, input power and voltage-controlled attenuators to adjust the millimeter-wave power. To minimize noise added to the signal entering the fridge, fixed or manual attenuators are placed last in the chain.

The qubit pulse is generated in a heterodyne fashion, similar to typical microwave experiments[14]. Triggered by a phase-locked trigger signal from the network analyzer, a 65 GS/s analog waveform generator (AWG Tektronix 70001a)[15] generates a microwave pulse with

14. See heterodyne pulse generation, for example the readout channel in Chapter 7

15. This device is typically used to generate all of the pulses in a microwave qubit experiment, but we are just using it to generate the pulse image for upconversion!

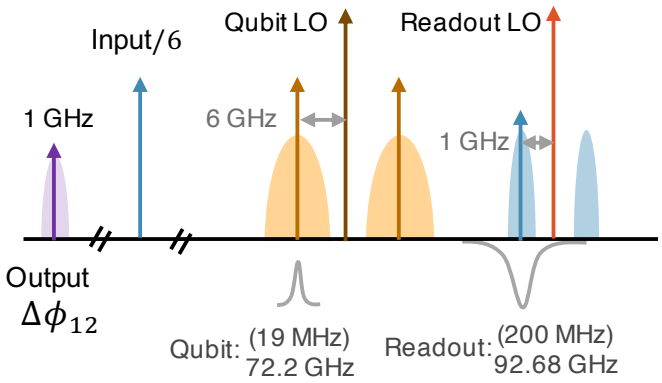

Figure 10.12: Frequency breakdown of the signals used for pulsed qubit measurements.

sub-ns resolution, which is mixed with a secondary LO (not phase-locked with the readout signals)[16] to generate a millimeter-wave pulse. The millimeter-wave mixer has limited conversion efficiency, so the qubit pulses are further amplified with a low-noise amplifier. Unfortunately this amplifier has a fairly low gain-compression point, so as a result the AWG amplitude does not directly translate to the millimeter-wave pulse amplitude. A vane attenuator and filter help block LO leakage and image signals, which help reduce the noise seen by the qubit to a small extent, however more filtering is likely necessary.

The addition of the heterodyne qubit pulse generation to the millimeter-wave input increases the number of frequencies at play, which complicates the frequency spectrum in our experiments. We summarize the frequencies used for pulse generation and readout, along with the bandwidth and frequencies of the qubit and readout resonator in Figure 10.12. Since the Fourier bandwidth of short millimeter-wave pulses can get quite high, and each multiplied signal generates additional images, it is important to consider the frequencies at play in the millimeter-wave system in order to avoid frequency collisions.

16. Since the qubit frequency is different than the readout frequency and we aren't doing fancy qubit–resonator operations it is okay to have random relative phase.

10.3.2 Pulse Timing

Having determined the qubit frequency with continuous-wave two-tone spectroscopy, we re-verify the qubit transitions by upconverting a continuous signal instead of a pulse at the qubit frequency using the measurement setup described above, described in C. This process helps verify that sufficient power is available[17] at the qubit frequency, and helps avoid unwanted frequency collisions.

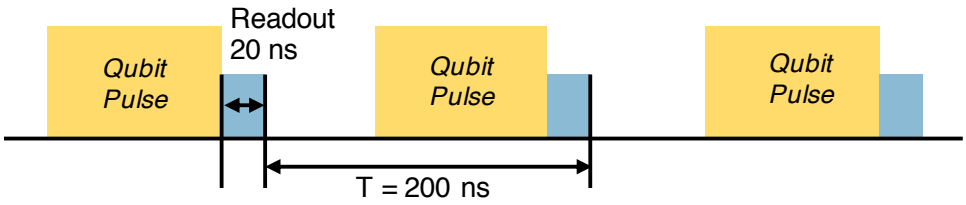

Figure 10.13: Millimeter-wave pulse timing.

With frequencies established, we can begin to explore the time domain with pulsed experiments. Since the millimeter-wave qubit coherence times are below 50 ns[18] we do not need to wait for long for the qubit to reset in between experiments. We choose an experiment repetition rate of 200 ns, with timings summarized in Figure 10.13. A timing calibration[19] synchronizes the qubit pulse region (orange) with the readout pulse (nominally 20 ns long). Within the synchronized qubit pulse region (typically 120 ns left available), we can program in the exact timings of qubit control pulses. The data acquisition, which is relatively slow, takes place for nearly 200 μs: during this time, the readout measurements of each individual experiment are integrated coherently (along with noise in between readout pulses). For this reason, we make sure to minimize the overall experiment repetition rate, since this directly increases measurement fidelity.

17. This helped us realize that we needed an amplifier on the qubit pulse side.

18. :(

19. Accomplished by matching the qubit frequencies with the readout frequencies, and measuring the downmixed output with an oscilloscope.

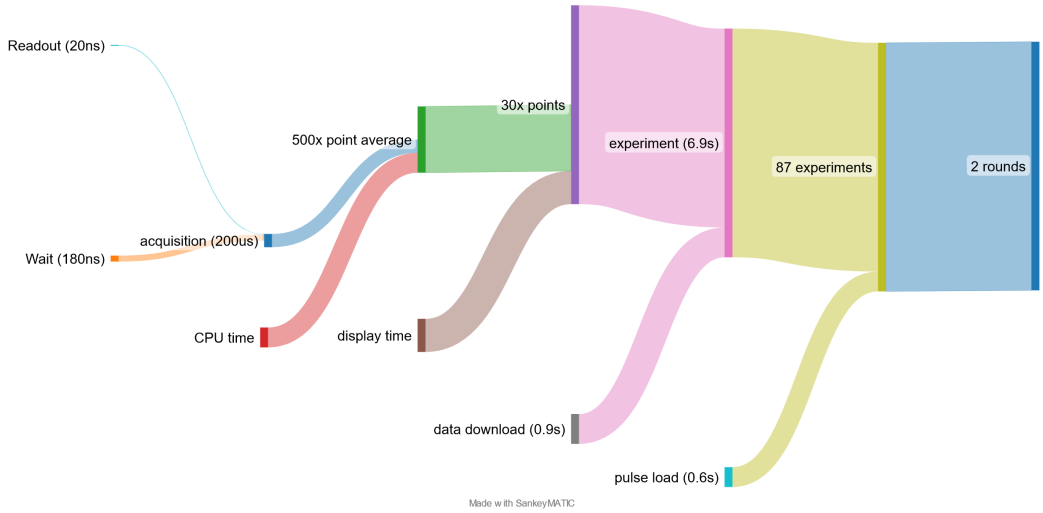

Figure 10.14: Time breakdown of a millimeter-wave experiment (not to scale). Most of the time is spent processing data and communicating with instruments.

The timing breakdown of an example pulse experiment consisting of a sweep of 87 points is shown in Figure 10.14. During the acquisition window, approximately 1000 experiments are integrated[20]. In descending order of speed, 500 point averages and 30 sweep averages are performed on the network analyzer. Including two start-to-finish averages (which help decouple slow fluctuations), each point has effectively been averaged 50 million times. The reason for this high number of averages is the extremely low dispersive shift relative to the resonator linewidth, which results in very low fidelity for an individual experiment.

10.3.3 Flat Top Pulse Envelope

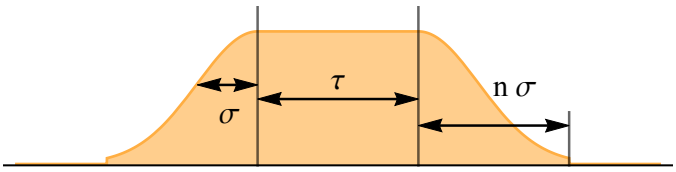

Figure 10.15: Millimeter-wave pulse envelope shape.

For qubit gate operations, we select a flat top pulse with gaussian edges, as illustrated in

20. Due to the dead time in acquisition, each experiment acquisition has a fidelity below 10 percent.

Figure 10.15: this pulse shape compromises between the constant rabi rotation rate during the flat part of the pulse (length τ) while reducing the Fourier bandwidth of the pulse with Gaussian edges (with Gaussian width σ). For keep the overall pulse length practical, the pulse envelope is truncated after n Gaussian widths, making the total pulse length

$$t_{\text{pulse}} = \tau + 2n\sigma \tag{10.8}$$

For most of the pulses in the experiments that follow, we will use $\sigma = 1.5$ ns and $n = 2.5$, resulting in overall pulse times of $\tau + 7.5$ ns. While this means that even a $\tau = 0$ pulse will take a significant percentage of the qubit coherence time, the qubit anharmonicity limits how fast the qubit control pulses can be, as we will see later on.

10.3.4 Frequency Dependence of Rabi Oscillations

When a fixed-length pulse, the qubit will also oscillate between the ground and excited state depending on the pulse amplitude: this Rabi oscillation rate Ω decreases to a minimum as the pulse is brought on resonance with the qubit transition frequency f_{ge}. In Figure 10.16, we measure the readout signal as a function of AWG pulse amplitude[21], for varying pulse frequencies. At the qubit frequency, we observe a series of dark fringes; however as the detuning from the transition increases, the oscillation rate increases, while the oscillation contrast is quickly suppressed. The bandwidth of the fringes is affected by the finite nature of the pulse [24, 26, 67], as well as mild variation in pulse amplitude at varying frequency. From Figure 10.16a we observe that the expected fringe behavior is only resolvable above f_{ge}. Similarly to how higher level transitions appear at frequencies evenly spaced by $\alpha/2$ in Figure 10.7, we observe Rabi oscillations driven between higher levels at lower frequencies. These oscillations require much more power than oscillations between the ground and excited

21. The actual millimeter-wave pulse amplitude is gain-compressed.

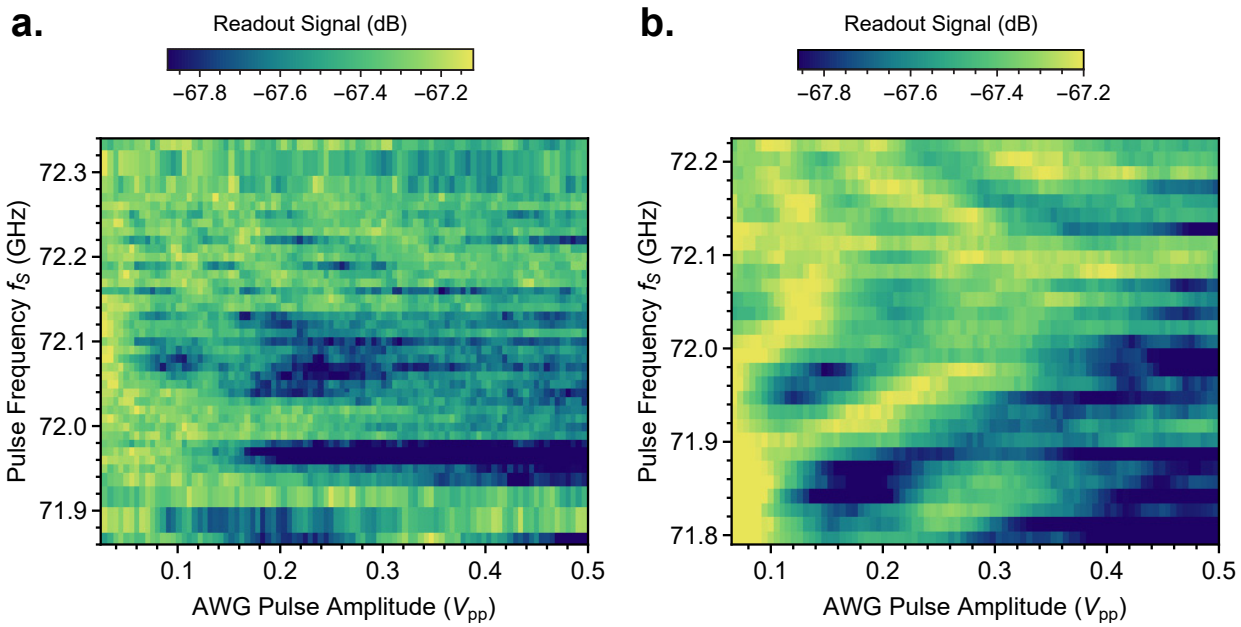

Figure 10.16: Frequency dependence of Rabi oscillation experiments. **a)** Inverse deflection of the readout signal with relative to its rest state with respect to original pulse amplitude, for varying pulse frequencies. Darker colors are higher deflection magnitudes corresponding to the excited states of the qubit. Near the qubit resonance (72.1 GHz) the oscillation fringes slow and increase in contrast. Below the qubit frequency, oscillations between higher order transitions (g–f and e–f) are also visible, complicating the spectrum. Horizontal stripes are a result of interleaved experiments with different averaging conditions. **b** The same experiment is repeated with shorter pulses ($\tau = 0$) less attenuation, which increases the overall oscillation rate. This broadens the features in the frequency domain, causing the various transition oscillations to overlap and interfere. Because the qubit is now oscillating between superpositions of the first four energy levels the initial fringe contrast is reduced.

states, however complicate the frequency spectrum of the Rabi fringes.

Increasing the amplitude (achieved by reducing the attenuation with the manual vane attenuator) allows for shorter qubit pulses, which broadens the fringe pattern so much that the fringes interfere with the higher level qubit transitions, as observed in Figure 10.16b. Interestingly enough in both of these cases, the center of the fringes is located below the fundamental qubit frequency measured with two-tone spectroscopy. This suggests that the pulse measurement setup may be populating the readout resonator with unwanted noise.

Fixed-length pulses with non-uniform pulse envelopes result in a nontrivial final qubit state, since the adiabiticity of the qubit evolution varies during the pulse duration [24, 26, 67].

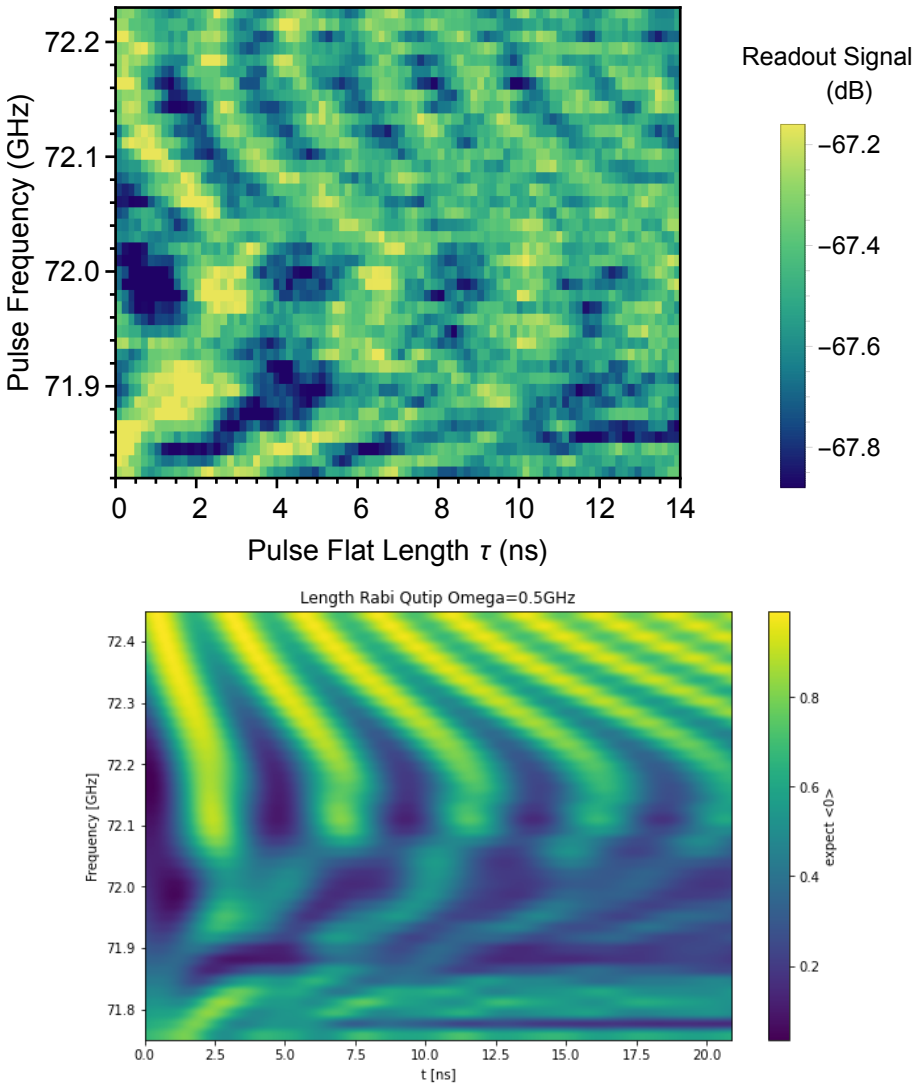

Figure 10.17: **a)** Rabi oscillations as a function of pulse length and pulse frequency for a large amplitude pulse. As the pulse detuning increases, the Rabi oscillation frequency should increase. However in this experiment the pulse amplitude is high enough such that the Rabi fringes for higher level transitions interfere with each other, resulting in fringes that merge as the frequency drops below the qubit frequency. **b)** A master equation simulation of the qubit shows qualitative similarity in fringe patterns.

The behavior is still sinusoidal with respect to pulse length and amplitude Ω, and can be written explicitly for a finite gaussian pulse truncated at $\pm n\sigma$:

$$P_e = \frac{\Omega^2 e^{-n^2}}{\Omega^2 e^{-n^2} + \Delta^2} \sin^2\left(\frac{\sigma}{2}\int_{-n}^{n} \partial\tau \sqrt{\Omega^2 e^{-\tau^2} + \Delta^2}\right) \tag{10.9}$$

Where Δ is the detuning from the transition frequency. On resonance ($\Delta = 0$) we recover the familiar expression

$$P_e = \sin^2(\pi\sigma\Omega) \tag{10.10}$$

From this behavior we conclude that the bandwidth of Rabi oscillations is reduced for real pulse envelopes. To remedy this, we can instead more closely emulate the ideal Rabi behavior by keeping the pulse amplitude (and thus Rabi rate) constant while adjusting the evolution with the pulse length, which should help increase the bandwidth of the fringes. To reduce the impact on pulse bandwidth the gaussian width of the pulse edges is held constant while the pulse flat length τ is swept. However since the qubit is expected to dephase for long pulses, we attempt to compensate by increasing the Rabi rate with high pulse amplitudes. The resulting experiments are shown in Figure 10.17. While the fringe bandwidth has increased, we find that this results in interference between higher qubit energy levels. A simulation of the master equation can qualitatively describe the behavior, but this ultimately proves that the qubit pulses are fundamentally limited by the qubit anharmonicity.

10.3.5 Optimized Rabi Oscillations

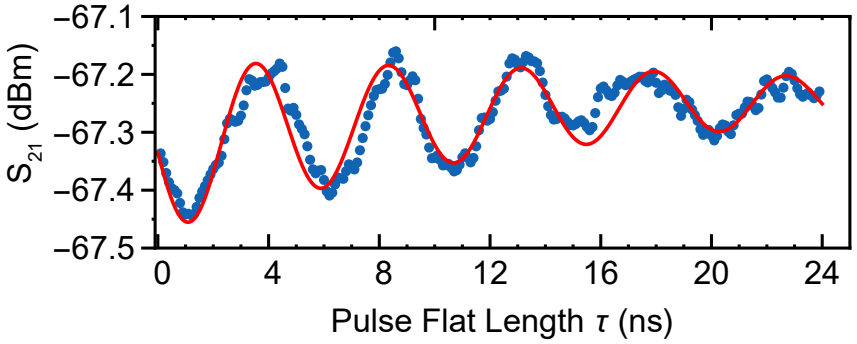

Figure 10.18: Rabi oscillations. When a pulse is applied at the qubit frequency, we expect it to oscillate between its ground and excited states. The readout resonator transmission is measured for various applied qubit pulse lengths τ. Here the pulse edges have a gaussian profile with $\sigma = 1.5$ ns. The red line is a fit to a decaying sinusoid, yielding a Rabi frequency of 208.6 MHz.

Having established a maximum Rabi rate, we select an attenuation which results in Rabi oscillations just below the anharmonicity: while off-resonance this still may still results in some interference with higher-level transitions, at and above f_{ge} we can still expect relatively well understood interactions that don't involve the higher energy levels. In 10.18 we show an optimized experiment chosen for an optimal frequency on-resonance with the qubit, where the qubit oscillates between its ground and excited state as a function of the pulse flat length τ. Due to the finite length of the pulse edges, $\tau = 0$ already corresponds to nonzero qubit rotation. For longer pulses we observe that the oscillations decay towards the ground state of the qubit, which is consistent with qubit decoherence during the length of the pulse, which are comparable to the expected qubit coherence times.

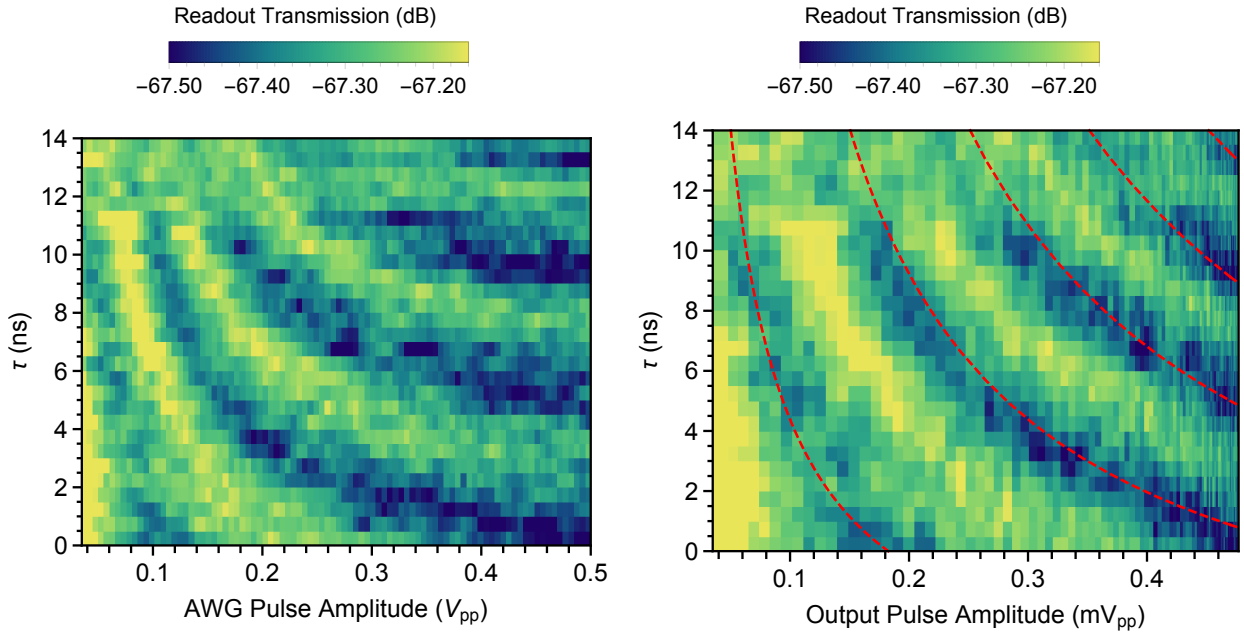

Figure 10.19: The final qubit state following a resonant pulse should depend on the product of pulse length and pulse amplitude. To verify this we perform a sequence of Rabi experiments while varying pulse length τ and pulse amplitude. To limit the effects of pulse bandwidth, the gaussian edges of the pulse are kept to $\sigma = 1.5$ ns. The raw data for this experiment is shown on the left, where qualitatively we observe that as pulse amplitude increases, so does the Rabi oscillation rate, evidenced by the reduced vertical fringe spacing. Since the pulse gain is compressed by an amplifier, it is more useful to examine this data in terms of the calibrated millimeter-wave pulse amplitude seen by the qubit, as shown on the right. From this we can overlay the expected contours where the qubit is in its excited state.

To verify the behavior of the Rabi oscillations, we also repeat this measurement while varying pulse amplitude and plot the results in Figure 10.19. Consistent with Equation 10.10, the observed fringes are evenly spaced for fixed amplitude (vertical slice) or fixed length (horizontal) when the pulse amplitude is corrected for gain compression. Fitting the contours of the fringes to the periodic conditions $\sigma\Omega = m\pi$ allows us to extract the Rabi rate (as high as 208 MHz), and also serves as a good method for calibrating a qubit π control pulse.

10.3.6 Millimeter-wave Qubit Coherence in the Time Domain

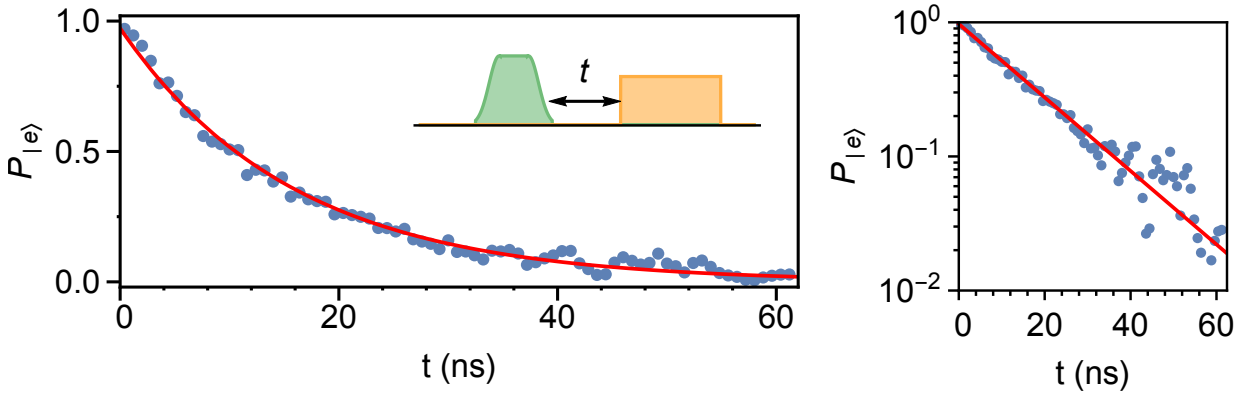

Figure 10.20: Qubit coherence properties. A relaxation experiment (illustrated in the inset) measuring qubit decay time T_1 by fitting the exponential decay of excited state population following a time delay t after a π pulse. The same data is shown on a logarithmic scale on the right, confirming the exponential decay. From this measurement we find $T_1 = 15.849$ ns.

Using the Rabi experiment above as a calibration, we select parameters for a relatively short ($\tau = 2$ ns) π pulse. Finally with control over the qubit, we can study the millimeter-wave qubit relaxation and dephasing in the time domain, which are the best method for benchmarking the qubit performance. We measure relaxation time by placing each qubit in its excited state with the aforementioned π pulse and measuring it after time t. Performing this experiment with negative t allows us to calibrate the exact pulse timing. The results are shown in Figure 10.20: fitting the exponential decay gives the characteristic relaxation

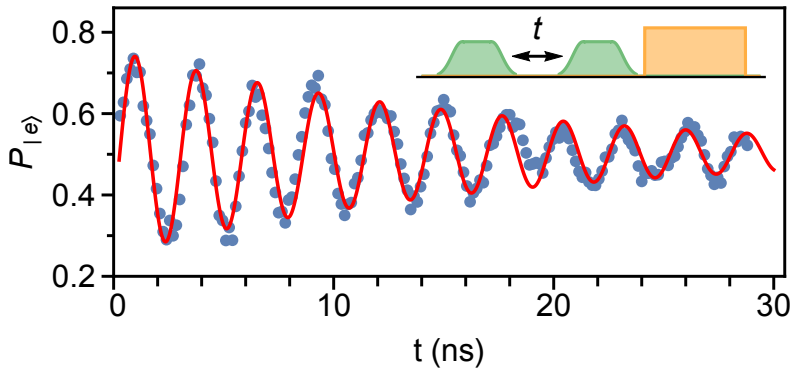

Figure 10.21: Ramsey experiment (illustrated in inset) measuring qubit dephasing time T_2 by fitting the decay of oscillations. The second $\pi/2$ pulse has its phase advanced by 320 MHz. We extract a Ramsey dephasing time of 17.466 ns.

time T_1. By checking the results on a logarithmic scale, we verify that the decay is indeed exponential.

By reducing its amplitude while keeping the pulse envelope fixed, we obtain a $\pi/2$ pulse corresponding to a half-rotation between the ground and excited state. We perform a Ramsey experiment[22] by measuring the qubit state after applying two $\pi/2$ pulses separated by time t. The phase of the second $\pi/2$ pulse is advanced at a fixed frequency detuning (320 MHz): sweeping through phase simultaneously measures both the coherence of the qubit as well as its frequency. The results in an oscillation between the ground and excited state, at the detuned frequency in the qubit frame, which decays with the characteristic qubit dephasing time T_2. From this we extract a oscillation frequency of 359 MHz, indicating the qubit frequency is 39 MHz away (which sounds high but is fairly close considering how wide the control pulses are), and a dephasing time of 17.466 ns. Notably due to the length of a single pulse, this oscillation begins with reduced contrast, since the qubit has already partially decayed, consistent with the length of control pulses relative to the coherence time.

22. Due to our terrible readout fidelity, this experiment took more than a day of data acquisition!

Table 10.2: Millimeter-wave Qubit Coherence Properties

T_1	T_2	Q_1	Q_2
15.849 ns	17.466 ns	7.183×10^3	7.916×10^3

10.4 Qubit Performance Properties in Context

The qubit coherence properties measured with time-domain spectroscopy are summarized in Table 10.2. The coherence properties are lower in the time domain than previously estimated: the measured dephasing time is consistent but slightly lower than the estimated value from continuous wave linewidth experiments, while the relaxation time is signifcantly lower than the linewidth estimate. This highlights the limited accuracy of the linewidth method for estimating coherence times in an enivronment like ours where incident noise is not well controlled. Interestingly enough this qubit's quality factors are relatively low compared to the microwave or K band devices.

In hindsight[23] our qubit performance is primarily limited in two ways. First, from subsequent[24] measurements of failed HDPCVD K-Band qubits, we now know that the junction preparation method used to make the millimeter-wave qubits (which also relies on HDPCVD spacer deposition) generates a significant source of loss, consistent with junction quality factors between 10^3–10^4. Luckily, this will be an easy future fix, since we have now demonstrated a low-loss junction fabrication method with the K-Band qubits. Second, we have repeatedly seen that the millimeter-wave qubit is limited by its low anharmonicity. Since this is mainly a design decision, a future iteration of this device utilizing smaller capacitors could easily provide anharmonicities above 500 MHz with no changes to the fabrication process[25]. With

23. Hindsight is 20-20!

24. The millimeter-wave qubits were fabricated months before we figured out how to optimize the high-J_C PECVD junction process.

25. We actually measured qubits with anharmonicities as high as 560 MHz, however they were incredibly lossy due to a bad batch of junction residues.

higher resolution junction lithography (such as electron beam lithography) we could achieve higher precision junction areas and increase the anharmonicity limit even further.

Importantly, this suggests that we are not seeing fundamental limits of coherence at millimeter-waves but rather have only begun optimizing these quantum systems in a previously unexplored regime. On a positive note, for the first time we have realized an artificial supercoducting atom operating at liquid helium-4 temperatures. Our demonstration, akin to a proof of concept device, has provided us with the knowledge about superconducting quantum devices in a new freuqency range, putting us in a perfect place to begin translating the many microwave quantum experiments to the millimeter-wave range.

CHAPTER 11

CONCLUSIONS AND OUTLOOK

Through a circuitous journey exploring many aspects of superconducting quantum devices, we have taken a few first steps in scaling up microwave qubits to the millimeter-wave range. Throughout all of our devices, we see a repeated theme of increasing qubit resilience against its environment, reinforcing the benefits of higher frequencies. Our niobium trilayer microwave qubits were unhindered by quasiparticles above 160 mK, which currently limit state-of-the-art aluminum microwave qubits. Furthermore we showed that these qubits were capable of operating in the presence of significantly stronger magnetic fields than possible with aluminum junctions. Scaled up to the K band, we saw qubits able to operate as high as 250 mK while maintaining their low temperature coherence properties. Finally we demonstrated coherent control over a millimeter-wave 72 GHz qubit capable of operating at 0.87 K, which is high enough to utilize a significantly less complex refrigeration mechanism. Aside from the order of magnitude increase in maximum possible operation speed, the cooling method alone has huge implications for scaling up quantum systems or integrating with more complex experiments.

11.1 Coherence Properties

For each round of qubits operating at increasing ferquencies, we explored their coherence properties in depth. Taking a step back, we can take a moment to compare their performance across the various technologies. In Figure 11.1 we combine all qubit relaxation times for the devices discussed here, grouped by frequency range and by wafer type. We see that qubit quality factor $Q_1 = \omega_q T_1$ is an invaluable tool for comparing qubits with vastly different frequencies. Likewise in Figure 11.2 we combine all qubit dephasing times, grouped by frequency range and by wafer type.

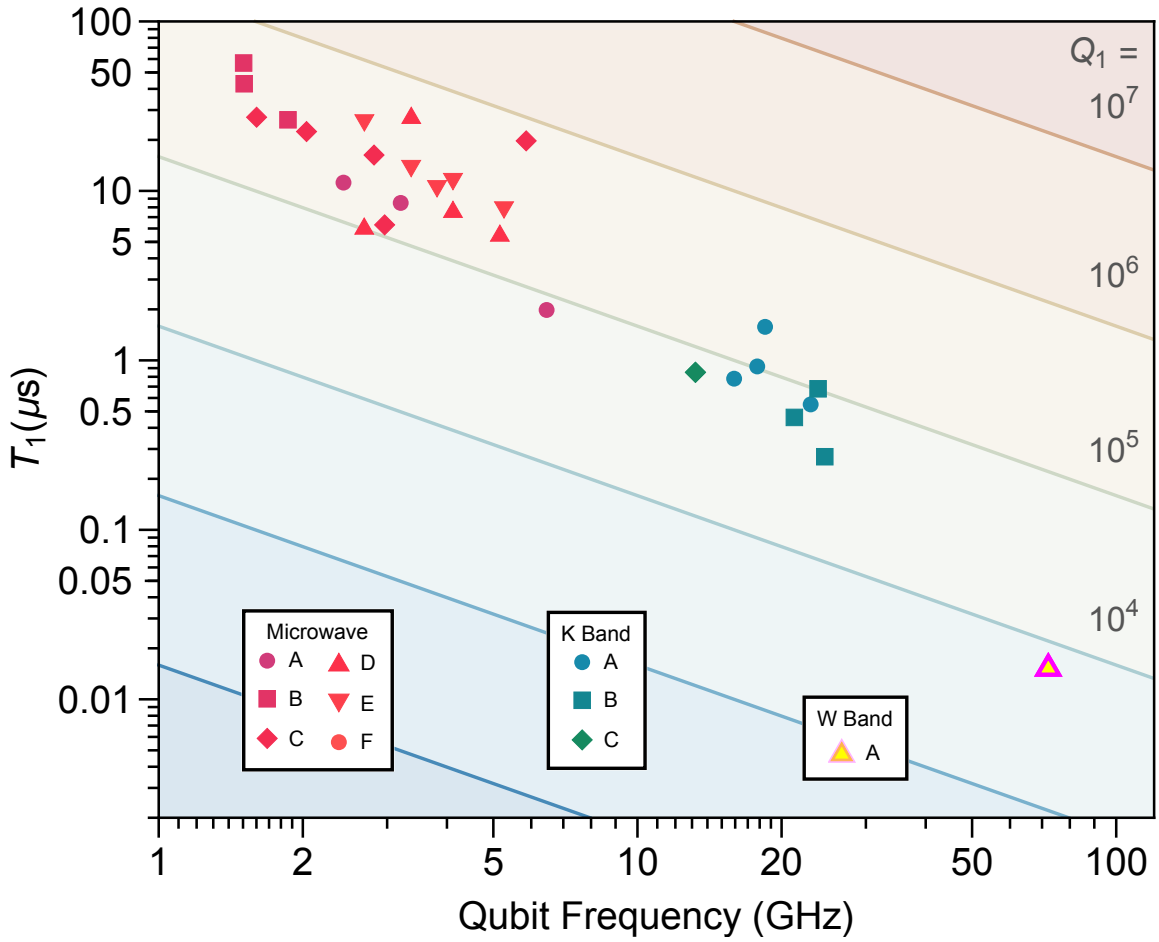

Figure 11.1: Qubit relaxation times T_1 summarized as a function of qubit frequency for the niobium trilayer devices studied in this work. Lines indicate qubit quality factor $Q_1 = \omega_q T_1$.

For all of our devices, we frequently find that as with early iterations of microwave qubits [148], an uphill battle must be fought against losses and decoherence. Given the depth of existing background, one might be tempted to conclude that the higher-frequency devices we explored may be inherently more lossy- however we note that the higher frequency devices we began to explore are the first of their kind, and have not yet experienced decades of refinement and investigation. Indeed this is an exciting development for the quantum community as a whole, as these high frequency devices offer a new lens through which decoherence channels can be investigated.

In particular we found that a potentially useful tool for disentangling loss channels in any

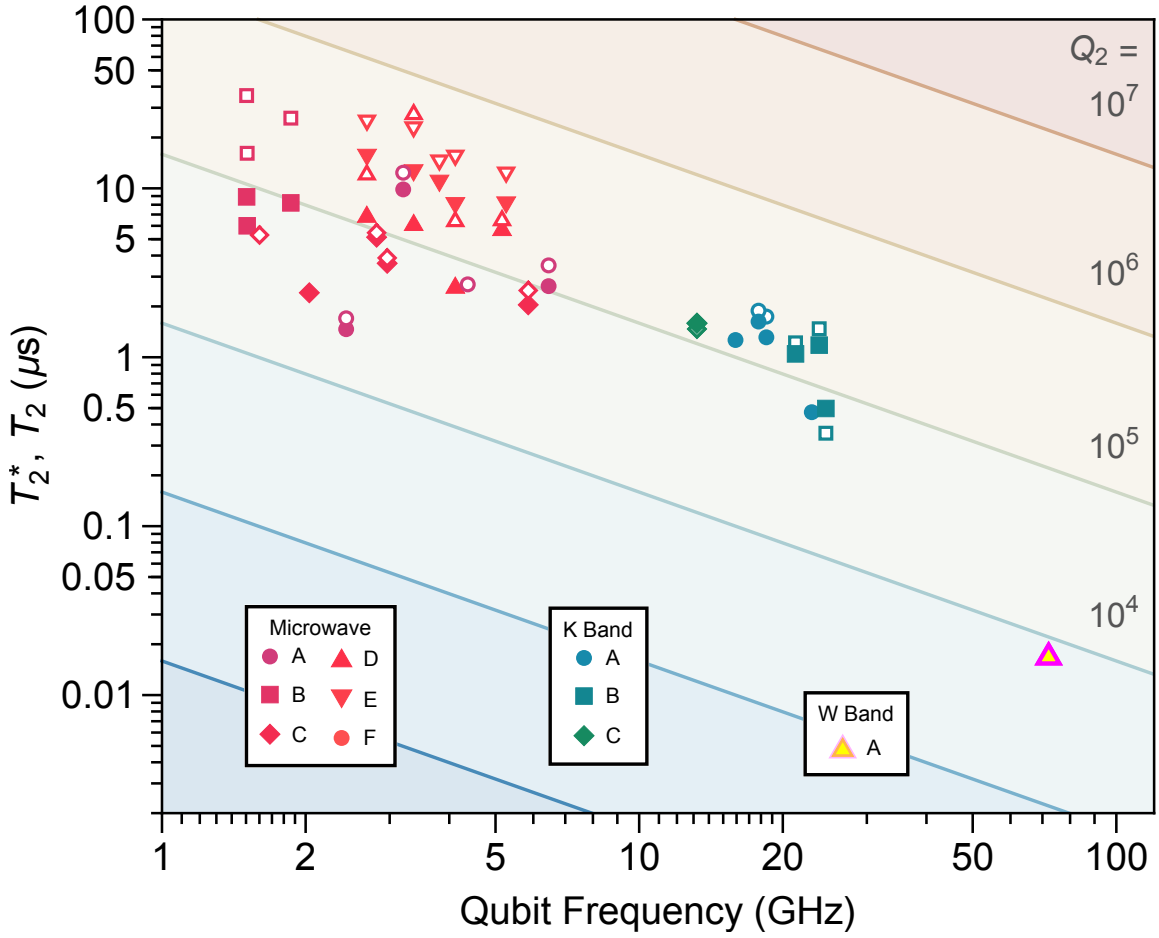

Figure 11.2: Ramsey dephasing time T_2^* (filled points) and Hahn-echo dephasing time T_2 (hollow points) summarized as a function of qubit frequency for the niobium trilayer devices studied in this work. Lines indicate quality factor $Q_2 = \omega_q T_2$.

kind of Josephson junction-based qubits is the junction participation p_j, a measure of how much energy is stored in the junction oxide relative to the rest of the circuit. In Figure 11.3 we compare the qubit quality factors for all of the devices discussed in this work along with conventional qubit examples from literature and from our lab. In fact, this has illustrated the source of some improvements in conventional microwave devices, which have achieved extremely long coherence times by diluting the stored energy away from the junction, and into the rest of the circuit, which has in parallel seen significant materials optimization to further reduce loss[1]. However this brings to light a potential avenue for further microwave

1. In the past seven years we have seen the gold standard for quality factor increase from 10^6 to 10^7!

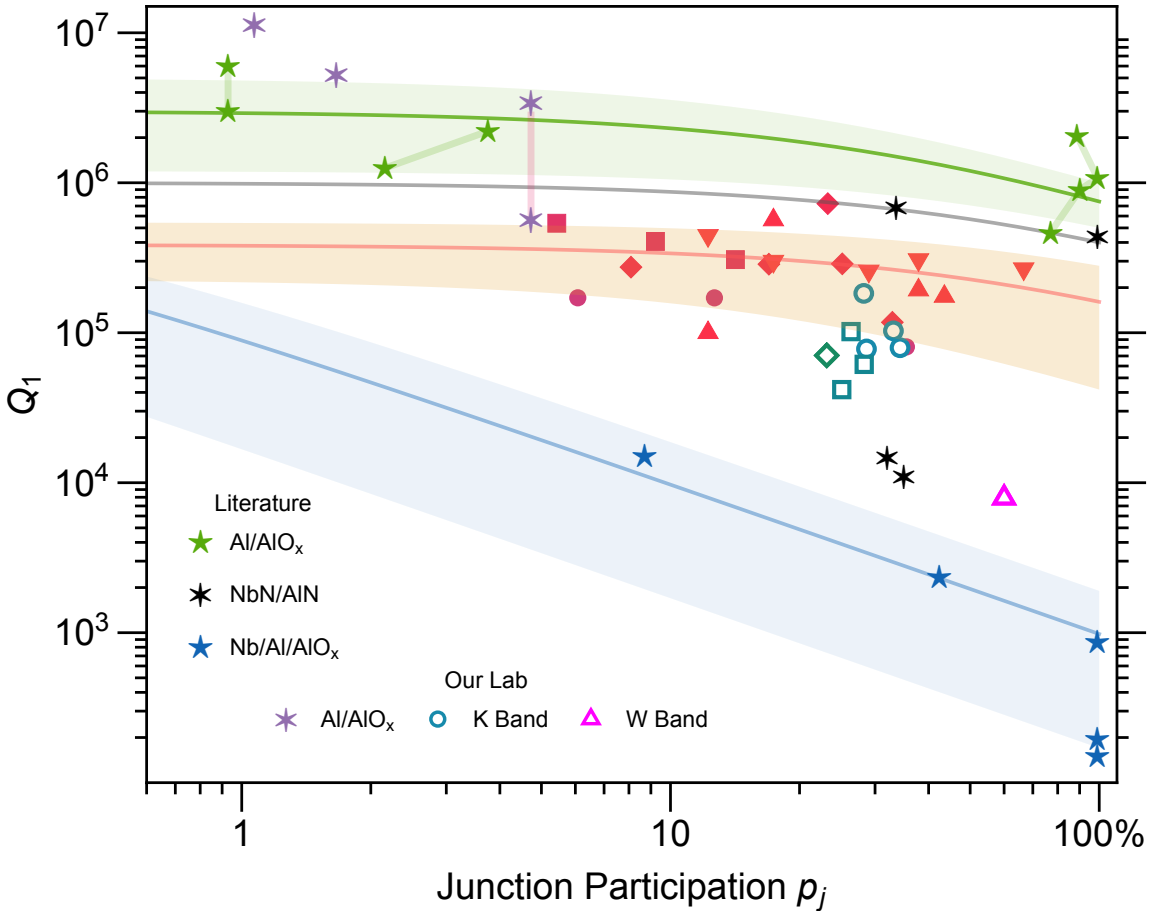

Figure 11.3: Qubit quality factors as a function of their junction participation ratio plotted for our high frequency qubits (hollow points) our microwave trilayer qubits (reds) aluminum junction qubits from our lab (purple) and selected qubits in literature (blue, black, green). Lines and shaded confidence regions show $Q_1^{-1} = p_J/Q_J + p_0/Q_0$ as a guide to the eye.

qubit improvements focusing on losses specifically in the junction.

In our case, we find that the relatively low-loss operation of our high-temperature and high-frequency qubits is made possible by the improvements made to the standard niobium trilayer Josephson junction. Evidently by taking steps to reduce the presence of lossy materials we have been able to increase the junction quality factor by a factor of 10 to 100! And by no means are these junctions ideal- we still have a long way to go even to improve niobium trilayer junctions, and many avenues available to explore the properties of other high-temperature Josephson junctions.

11.2 Applications

Having investigated how microwave qubits can be scaled up to higher frequencies we can't directly conclude that higher-frequencies are a direct upgrade over the microwave bands for realizing quantum information bits. Aside from being a challenge to realize experimentally, we have seen evidence that fighting sources of loss is an uphill battle with respect to device frequency, and unfortunately decoherence channels don't appear to go away. (This frequency dependence in and of itself may be an interesting phenomenon to investigate- for example perhaps we can use this new qubit platform to interrogate the frequency-dependence of material-based two level system density). It is our hope however that these first attempts at higher frequency devices will pave the way for subsequent improvements which will only help improve their usefulness for quantum operations specifically.

That said, these novel qubits have demonstrated their strength in terms of environmental resilience. This in particular makes these qubits an attractive platform for experiments making use of qubits in challenging experimental environments [44]. In particular more resilient qubits could be beneficial for quantum detection experiments [56] and could be directly integrated for high-frequency quantum detection experiments in particular [41], which would be revolutionary for millimeter-wave astronomy.

11.3 Future Directions

We have done the first few steps in characterizing these high-frequency qubits, and while it would be very useful to improve their performance, there are already a number of experiments that would be interesting to replicate on this high-frequency platform. We have mentioned previously that frequency is a kind of orthogonal upgrade that could be applied to any microwave qubit design. To this end it would be particularly interesting to scale up promising microwave designs, such as Kerr cat qubits [80] (or any bosonic qubit encoding really), or

fluxonium qubits [209, 266], which have recently proven to simultaneously support fast gate operations and long coherence times. It follows that the circuit parameters for these devices could be scaled up directly, which could potentially make them an attractive qubit platform to investigate.

It would be fascinating to realize a coupled microwave–millimeter-wave system proposed in Ref. [181]: whether with kinetic inductance or with a niobium trilayer Josephson junction, this could be useful for transduction experiments (whether millimeter or optical). However a more fun application might be to use this coupled system to artificially cool a microwave mode (Similar perhaps to the bath engineering done in Ref. [233]).

A particularly exciting direction would be to directly integrate our millimeter-wave qubit with Rydberg neutral atoms [119], which have atomic transitions in the millimeter-wave range. As the millimeter-wave and optical energy difference is significantly reduced, this could be a promising platform for realizing long distance communication [244] or for demonstrating optical readout of a millimeter-wave qubit [156, 240]. Regardless we look forward to the future of higher-frequency[2] superconducting quantum devices.

2. Hopefully more millimeter-waves!

REFERENCES

[7] Fairview .047 low loss semi-rigid coax cable with tin plated copper outer conductor and microporous ptfe 76.5 pct vop dielectric, October 2023. [Online; accessed 14. Oct. 2023]. (Cited on page 47).

[8] Mahmoud Al Henawy and Martin Schneider. Rectangular waveguide to coplanar stripline transition based on a unilateral finline. In *2012 6th European Conference on Antennas and Propagation (EUCAP)*, pages 405–409, March 2012. doi:10.1109/EuCAP.2012.6206336. (Cited on pages 198, 199, 200, and 201).

[9] M. Virginia P. Altoé, Archan Banerjee, Cassidy Berk, Ahmed Hajr, Adam Schwartzberg, Chengyu Song, Mohammed Alghadeer, Shaul Aloni, Michael J. Elowson, John Mark Kreikebaum, Ed K. Wong, Sinéad M. Griffin, Saleem Rao, Alexander Weber-Bargioni, Andrew M. Minor, David I. Santiago, Stefano Cabrini, Irfan Siddiqi, and D. Frank Ogletree. Localization and mitigation of loss in niobium superconducting circuits. *PRX Quantum*, 3(2):020312, April 2022. ISSN 2691-3399. doi:10.1103/PRXQuantum.3.020312. (Cited on pages 163, 178, 179, and 180).

[10] Vinay Ambegaokar and Alexis Baratoff. Tunneling between superconductors. *Physical Review Letters*, 10(11):486, 1963. (Cited on pages 79, 82, 152, and 164).

[11] S Anders, M Schmelz, L Fritzsch, R Stolz, V Zakosarenko, T Schönau, and HG Meyer. Sub-micrometer-sized, cross-type nb–alox–nb tunnel junctions with low parasitic capacitance. *Superconductor Science and Technology*, 22(6):064012, 2009. (Cited on page 152).

[12] Alexander Anferov, Aziza Suleymanzade, Andrew Oriani, Jonathan Simon, and David I. Schuster. Millimeter-Wave Four-Wave Mixing via Kinetic Inductance for Quantum Devices. *Physical Review Applied*, 13(2):024056, February 2020. doi:10.1103/PhysRevApplied.13.024056. (Cited on pages 143, 182, 196, 201, 213, 216, and 217).

[13] Alexander Anferov, Kan-Heng Lee, Fang Zhao, Jonathan Simon, and David I Schuster. Improved coherence in optically-defined niobium trilayer junction qubits. *arXiv preprint arXiv:2306.05883*, June 2023. doi:10.48550/arXiv.2306.05883. (Cited on pages 184, 192, 196, and 217).

[14] Anthony J Annunziata, Daniel F Santavicca, Luigi Frunzio, Gianluigi Catelani, Michael J Rooks, Aviad Frydman, and Daniel E Prober. Tunable superconducting nanoinductors. *Nanotechnology*, 21(44):445202, 2010. (Cited on page 17).

[15] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney,

Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574:505–510, October 2019. ISSN 1476-4687. doi:10.1038/s41586-019-1666-5. (Cited on page 2).

[16] Neil W Ashcroft and N David Mermin. *Solid state physics*. Cengage Learning, 2022. (Cited on pages xxx, 8, 9, 21, and 26).

[17] O V Astafiev, L B Ioffe, S Kafanov, Yu A Pashkin, K Yu Arutyunov, D Shahar, O Cohen, and J S Tsai. Coherent quantum phase slip. *Nature*, 484(7394):355, 2012. (Cited on page 116).

[18] Mustafa Badaroglu. More moore. In *2021 IEEE International Roadmap for Devices and Systems Outbriefs*, page 30. Ieee. doi:10.1109/irds54852.2021.00010. (Cited on page 1).

[19] Inder Jit Bahl. *Lumped elements for RF and microwave circuits*. Artech house, 2003. (Cited on pages 66, 121, 124, and 187).

[20] A Barannik, NT Cherpak, MA Tanatar, S Vitusevich, V Skresanov, PC Canfield, and Ruslan Prozorov. Millimeter-wave surface impedance of optimally-doped ba (fe 1- x co x) 2 as 2 single crystals. *Physical Review B*, 87(1):014506, 2013. (Cited on page 8).

[21] Rami Barends, Julian Kelly, Anthony Megrant, Daniel Sank, Evan Jeffrey, Yu Chen, Yi Yin, Ben Chiaro, Josh Mutus, Charles Neill, et al. Coherent josephson qubit suitable for scalable quantum integrated circuits. *Physical review letters*, 111(8):080502, 2013. (Cited on pages 168, 172, 186, 190, and 227).

[22] Melissa R Beebe, Douglas B Beringer, Matthew C Burton, Kaida Yang, and R Alejandra Lukaszew. Stoichiometry and thickness dependence of superconducting properties of niobium nitride thin films. *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, 34(2):021510, 2016. (Cited on pages 97, 100, and 109).

[23] M T Bell, I A Sadovskyy, L B Ioffe, A Yu Kitaev, and M E Gershenson. Quantum superinductor with tunable nonlinearity. *Physical review letters*, 109(13):137003, 2012. (Cited on page 116).

[24] PR Berman, Lixin Yan, Keng-Hwee Chiam, and Ruwang Sung. Nonadiabatic transitions in a two-level quantum system: Pulse-shape dependence of the transition probability for a two-level atom driven by a pulsed radiation field. *Physical Review A*, 57 (1):79, 1998. (Cited on pages 94, 190, 237, and 238).

[25] Lev S Bishop, Eran Ginossar, and SM Girvin. Response of the strongly driven jaynes-cummings oscillator. *Physical review letters*, 105(10):100505, 2010. (Cited on pages 222 and 223).

[26] Iavor I. Boradjiev and Nikolay V. Vitanov. Control of qubits by shaped pulses of finite duration. *Phys. Rev. A*, 88(1):013402, July 2013. ISSN 2469-9934. doi:10.1103/PhysRevA.88.013402. (Cited on pages 92, 94, 190, 237, and 238).

[27] Sangita Bose, Pratap Raychaudhuri, Rajarshi Banerjee, Parinda Vasa, and Pushan Ayyub. Mechanism of the size dependence of the superconducting transition of nanostructured nb. *Physical review letters*, 95(14):147003, 2005. (Cited on pages 155 and 156).

[28] Sangita Bose, Rajarshi Banerjee, Arda Genc, Pratap Raychaudhuri, Hamish L Fraser, and Pushan Ayyub. Size induced metal–insulator transition in nanostructured niobium thin films: intra-granular and inter-granular contributions. *Journal of Physics: Condensed Matter*, 18(19):4553, 2006. (Cited on pages 155 and 156).

[29] Maurizio Bozzi, Anthimos Georgiadis, and Kaijie Wu. Review of substrate-integrated waveguide circuits and antennas. *IET Microwaves, Antennas & Propagation*, 5(8): 909–920, 2011. (Cited on page 97).

[30] AI Braginski, J Talvacchio, MA Janocko, and JR Gavaler. Crystalline oxide tunnel barriers formed by thermal oxidation of aluminum overlayers on superconductor surfaces. *Journal of applied physics*, 60(6):2058–2064, 1986. (Cited on page 145).

[31] Alex I Braginski. Superconductor electronics: Status and outlook. *Journal of superconductivity and novel magnetism*, 32:23–44, 2019. (Cited on page 142).

[32] Alex I Braginski. Superconductor electronics: status and outlook. *Journal of Superconductivity and Novel Magnetism*, 32(1):23–44, 2019. (Cited on page 97).

[33] T. Brecht, M. Reagor, Y. Chu, W. Pfaff, C. Wang, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf. Demonstration of superconducting micromachined cavities. *Applied Physics Letters*, 107(19):192603, November 2015. doi:10.1063/1.4935541. (Cited on pages 66 and 216).

[34] Ari David Brown, Emily M Barrentine, Samuel H Moseley, Omid Noroozian, and Thomas Stevenson. The impact of standard semiconductor fabrication processes on polycrystalline nb thin-film surfaces. *IEEE Transactions on Applied Superconductivity*, 27(4):1–5, 2016. (Cited on page 97).

[35] M Brune, S Haroche, V Lefevre, JM Raimond, and N Zagury. Quantum nondemolition measurement of small photon numbers by rydberg-atom phase-sensitive detection. *Physical Review Letters*, 65(8):976, 1990. (Cited on page 2).

[36] Colin D Bruzewicz, John Chiaverini, Robert McConnell, and Jeremy M Sage. Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews*, 6(2), 2019. (Cited on page 2).

[37] Bruce Bumble, Andy Fung, Anu B Kaul, Alan W Kleinsasser, George L Kerber, Paul Bunyk, and Eric Ladizinsky. Submicrometer nb/al−alo_x/nb integrated circuit fabrication process for quantum computing applications. *IEEE transactions on applied superconductivity*, 19(3):226–229, 2009. (Cited on pages 144, 151, 164, and 166).

[38] Giuseppe Cataldo, Edward J. Wollack, Emily M. Barrentine, Ari D. Brown, S. Harvey Moseley, and Kongpop U-Yen. Analysis and calibration techniques for superconducting resonators. *Rev. Sci. Instrum.*, 86(1), January 2015. ISSN 0034-6748. doi:10.1063/1.4904972. (Cited on page 204).

[39] G. Catelani, R. J. Schoelkopf, M. H. Devoret, and L. I. Glazman. Relaxation and frequency shifts induced by quasiparticles in superconducting qubits. *Physical Review B*, 84(6):064517, August 2011. doi:10.1103/PhysRevB.84.064517. (Cited on pages 66, 181, 193, and 216).

[40] Gianluigi Catelani. Parity switching and decoherence by quasiparticles in single-junction transmons. *Physical Review B*, 89(9):094522, 2014. (Cited on pages 181 and 193).

[41] C L Chang, P A R Ade, Z Ahmed, S W Allen, K Arnold, J E Austermann, A N Bender, L E Bleem, B A Benson, J E Carlstrom, et al. Low loss superconducting microstrip development at argonne national lab. *IEEE Transactions on Applied Superconductivity*, 25(3):1–5, June 2014. doi:10.1109/tasc.2014.2369231. (Cited on pages 97, 214, and 250).

[42] WH Chang. The inductance of a superconducting strip transmission line. *Journal of Applied Physics*, 50(12):8129–8134, 1979. (Cited on page 15).

[43] A. A. Clerk and D. Wahyu Utami. Using a qubit to measure photon-number statistics of a driven thermal oscillator. *Physical Review A*, 75(4):042302, April 2007. ISSN 2469-9934. doi:10.1103/PhysRevA.75.042302. (Cited on pages 4, 183, and 193).

[44] AA Clerk, KW Lehnert, P Bertet, JR Petta, and Y Nakamura. Hybrid quantum systems with circuit quantum electrodynamics. *Nature Physics*, 16(3):257–267, March 2020. doi:10.1038/s41567-020-0797-9. (Cited on pages 5, 196, and 250).

[45] John W Coburn and Harold F Winters. Plasma etching—a discussion of mechanisms. *Journal of vacuum Science and Technology*, 16(2):391–403, 1979. (Cited on pages 149 and 157).

[46] Neil G Connelly and William E Geiger. Chemical redox agents for organometallic chemistry. *Chemical Reviews*, 96(2):877–910, 1996. (Cited on pages 160 and 161).

[47] Thomas Connolly, Pavel D Kurilovich, Spencer Diamond, Heekun Nho, Charlotte GL Bøttcher, Leonid I Glazman, Valla Fatemi, and Michel H Devoret. Coexistence of nonequilibrium density and equilibrium energy distribution of quasiparticles in a superconducting qubit. *arXiv preprint arXiv:2302.12330*, 2023. (Cited on pages 118 and 143).

[48] Leon N Cooper. Bound electron pairs in a degenerate fermi gas. *Physical Review*, 104 (4):1189, 1956. (Cited on pages 8, 17, and 28).

[49] Kevin D. Crowley, Russell A. McLellan, Aveek Dutta, Nana Shumiya, Alexander P. M. Place, Xuan Hoang Le, Youqi Gang, Trisha Madhavan, Nishaad Khedkar, Yiming Cady Feng, Esha A. Umbarkar, Xin Gui, Lila V. H. Rodgers, Yichen Jia, Mayer M. Feldman, Stephen A. Lyon, Mingzhao Liu, Robert J. Cava, Andrew A. Houck, and Nathalie P. de Leon. Disentangling Losses in Tantalum Superconducting Circuits. January 2023. doi:10.48550/arXiv.2301.07848. (Cited on pages 66, 214, 216, and 217).

[50] HM Cundy and AP Rollet. Mathematical models, rev. ed, 1961. (Cited on pages 127 and 131).

[51] A Danowitz, K Kelley, J Mao, JP Stevenson, M Horowitz, and CPU DB. Recording microprocessor history. *ACM Queue Magazine*, 10(4), 2002. (Cited on pages xi and 1).

[52] Andrew Danowitz, Kyle Kelley, James Mao, JP Stevenson, M Horowitz, Omid Azizi, JS Brunhaver, Ron Ho, Stephen Richardson, Ofer Shacham, and Alex Solomatnikov. Cpu db. Technical report, Stanford VSLI Group, 2015. (Cited on pages xi and 1).

[53] PGf de Gennes. Boundary effects in superconductors. *Reviews of Modern Physics*, 36 (1):225, 1964. (Cited on pages 23, 24, and 25).

[54] Oliver Dial, Douglas T. McClure, Stefano Poletto, G. A. Keefe, Mary Beth Rothwell, Jay M. Gambetta, David W. Abraham, Jerry M. Chow, and Matthias Steffen. Bulk and surface loss in superconducting transmon qubits. *Supercond. Sci. Technol.*, 29(4): 044001, March 2016. ISSN 0953-2048. doi:10.1088/0953-2048/29/4/044001. (Cited on pages 66 and 178).

[55] David P DiVincenzo and Daniel Loss. Quantum information is physical. *Superlattices and Microstructures*, 23(3-4):419–432, 1998. (Cited on page 3).

[56] Akash V. Dixit, Srivatsan Chakram, Kevin He, Ankur Agrawal, Ravi K. Naik, David I. Schuster, and Aaron Chou. Searching for dark matter with a superconducting qubit. *Phys. Rev. Lett.*, 126(14):141302, April 2021. ISSN 1079-7114. doi:10.1103/PhysRevLett.126.141302. (Cited on page 250).

[57] Simon Doyle, Philip Mauskopf, J Naylon, Adrian Porch, and C Duncombe. Lumped element kinetic inductance detectors. *Journal of Low Temperature Physics*, 151:530–536, 2008. (Cited on pages 15, 16, and 28).

[58] E F C Driessen, P C J J Coumou, R R Tromp, P J De Visser, and T M Klapwijk. Strongly disordered tin and nbtin s-wave superconductors probed by microwave electrodynamics. *Physical review letters*, 109(10):107003, 2012. (Cited on page 100).

[59] A Dunsworth, R Barends, Yu Chen, Zijun Chen, B Chiaro, A Fowler, B Foxen, E Jeffrey, J Kelly, PV Klimov, et al. A method for building low loss multi-layer wiring for superconducting microwave devices. *Applied Physics Letters*, 112(6):063502, 2018. (Cited on pages 150 and 156).

[60] SK Dutta, H Xu, AJ Berkley, RC Ramos, MA Gubrud, JR Anderson, CJ Lobb, and FC Wellstood. Determination of relaxation time of a josephson junction qubit. *Physical Review B*, 70(14):140502, 2004. (Cited on pages 142, 143, and 173).

[61] Christopher Eichler and Andreas Wallraff. Controlling the dynamic range of a josephson parametric amplifier. *EPJ Quantum Technology*, 1(1):2, 2014. (Cited on pages 73, 74, 110, 114, and 115).

[62] A Endo, C Sfiligoj, S J C Yates, J J A Baselmans, D J Thoen, S M H Javadzadeh, P P Van der Werf, A M Baryshev, and T M Klapwijk. On-chip filter bank spectroscopy at 600–700 ghz using nbtin superconducting resonators. *Applied Physics Letters*, 103(3): 032601, 2013. (Cited on pages 97 and 196).

[63] A. Endo, C. Sfiligoj, S. J. C. Yates, J. J. A. Baselmans, D. J. Thoen, S. M. H. Javadzadeh, P. P. van der Werf, A. M. Baryshev, and T. M. Klapwijk. On-chip filter bank spectroscopy at 600–700 GHz using NbTiN superconducting resonators. *Applied Physics Letters*, 103(3):032601, July 2013. doi:10.1063/1.4813816. (Cited on page 216).

[64] Byeong Ho Eom, Peter K Day, Henry G LeDuc, and Jonas Zmuidzinas. A wideband, low-noise superconducting amplifier with high dynamic range. *Nature Physics*, 8(8): 623, 2012. (Cited on page 116).

[65] Farzad Faramarzi, Peter Day, Jacob Glasby, Sasha Sypkens, Marco Colangelo, Ralph Chamberlin, Mohammad Mirhosseini, Kevin Schmidt, Karl K. Berggren, and Philip Mauskopf. Initial design of a w-band superconducting kinetic inductance qubit. *IEEE Transactions on Applied Superconductivity*, 31(5):ArticleSequenceNumber:1800605, March 2021. doi:10.1109/tasc.2021.3065304. (Cited on pages 118, 132, and 133).

[66] DK Finnemore, TF Stromberg, and CA Swenson. Superconducting properties of high-purity niobium. *Physical Review*, 149(1):231, 1966. (Cited on pages 142 and 155).

[67] Kevin A Fischer, Lukas Hanschke, Malte Kremser, Jonathan J Finley, Kai Müller, and Jelena Vučković. Pulsed rabi oscillations in quantum two-level systems: beyond the area theorem. *Quantum Science and Technology*, 3(1):014006, 2017. (Cited on pages 94, 190, 237, and 238).

[68] José A Flores-Livas, Lilia Boeri, Antonio Sanna, Gianni Profeta, Ryotaro Arita, and Mikhail Eremets. A perspective on conventional high-temperature superconductors at high pressure: Methods and materials. *Physics Reports*, 856:1–78, 2020. (Cited on pages xi, 7, and 8).

[69] Denise C. Ford, Lance D. Cooley, and David N. Seidman. Suppression of hydride precipitates in niobium superconducting radio-frequency cavities. *Supercond. Sci. Technol.*, 26(10):105003, August 2013. ISSN 0953-2048. doi:10.1088/0953-2048/26/10/105003. (Cited on pages 161 and 162).

[70] Jiansong Gao. *The physics of superconducting microwave resonators*. PhD thesis, California Institute of Technology, 2008. (Cited on pages 3, 114, 120, and 213).

[71] Jiansong Gao, Anastasios Vayonakis, Omid Noroozian, Jonas Zmuidzinas, Peter K Day, and Henry G Leduc. Measurement of loss in superconducting microstrip at millimeter-wave frequencies. In *AIP Conference Proceedings*, volume 1185, pages 164–167. Aip, December 2009. doi:10.1063/1.3292306. (Cited on pages 97, 196, and 216).

[72] Malte Giese, Tim Meinhardt, and Arne F. Jacob. Compact wideband single-ended and differential microstrip-to-waveguide transitions at W-band. In *2015 IEEE MTT-S International Microwave Symposium*, pages 1–4, May 2015. doi:10.1109/mwsym.2015.7167143. (Cited on pages 199, 200, and 201).

[73] Vitaly L Ginzburg, Vitaly Lazarevich Ginzburg, and LD Landau. *On the theory of superconductivity*. Springer, 2009. (Cited on pages 18 and 21).

[74] Leonid Glazman and Gianluigi Catelani. Bogoliubov quasiparticles in superconducting qubits. *SciPost Physics Lecture Notes*, page 031, 2021. (Cited on page 193).

[75] H Gokan, S Esho, and Y Ohnishi. Dry etch resistance of organic materials. *Journal of the electrochemical Society*, 130(1):143, 1983. (Cited on page 158).

[76] Alexandre Avraamovitch Golubov, M Yu Kupriyanov, and E Il'Ichev. The current-phase relation in josephson junctions. *Reviews of modern physics*, 76(2):411, 2004. (Cited on page 152).

[77] M Göppl, A Fragner, M Baur, R Bianchetti, S Filipp, JM Fink, PJ Leek, G Puebla, L Steffen, and Andreas Wallraff. Coplanar waveguide resonators for circuit quantum electrodynamics. *Journal of Applied Physics*, 104(11):113904, 2008. (Cited on pages 53 and 121).

[78] David J Griffiths. Introduction to electrodynamics, 2005. (Cited on pages 41 and 121).

[79] A Grimm, S Jebari, D Hazra, F Blanchet, F Gustavo, JL Thomassin, and M Hofheinz. A self-aligned nano-fabrication process for vertical nbn–mgo–nbn josephson junctions. *Superconductor Science and Technology*, 30(10):105002, 2017. (Cited on page 97).

[80] A. Grimm, N. E. Frattini, S. Puri, S. O. Mundhada, S. Touzard, M. Mirrahimi, S. M. Girvin, S. Shankar, and M. H. Devoret. Stabilization and operation of a kerr-cat qubit. *Nature*, 584:205–209, August 2020. ISSN 1476-4687. doi:10.1038/s41586-020-2587-z. (Cited on page 250).

[81] Leif Grönberg, Mikko Kiviranta, Visa Vesterinen, Janne Lehtinen, Slawomir Simbierowicz, Juho Luomahaara, Mika Prunnila, and Juha Hassel. Side-wall spacer passivated sub-μm josephson junction fabrication process. *Superconductor Science and Technology*, 30(12):125016, 2017. (Cited on pages 143, 146, 151, 152, 167, and 184).

[82] S. Hähnle, N. v. Marrewijk, A. Endo, K. Karatsu, D. J. Thoen, V. Murugesan, and J. J. A. Baselmans. Suppression of radiation loss in high kinetic inductance superconducting co-planar waveguides. *Applied Physics Letters*, 116(18), May 2020. ISSN 0003-6951. doi:10.1063/5.0005047. (Cited on page 119).

[83] S Hailey-Dunsheath, PS Barry, CM Bradford, G Chattopadhyay, P Day, S Doyle, M Hollister, A Kovacs, H G LeDuc, N Llombart, et al. Optical measurements of superspec: A millimeter-wave on-chip spectrometer. *Journal of Low Temperature Physics*, 176(5-6):841–847, 2014. (Cited on pages 3 and 97).

[84] Lian-Chen Han, Yu Xu, Jin Lin, Fu-Sheng Chen, Shao-Wei Li, Cheng Guo, Na Li, Dong-Dong Li, Yu-Huai Li, Ming Gong, et al. Active reset of superconducting qubits using the electronics based on rf switches. *AIP Advances*, 13(9), 2023. (Cited on pages 4 and 183).

[85] William J Otter Stephen M Hanham and Nick M Ridler. 100 ghz ultra-high q-factor photonic crystal resonators. *Nature*, 425:944–947, 2003. (Cited on page 97).

[86] Hila Hashemi, Alejandro W Rodriguez, J D Joannopoulos, Marin Soljačić, and Steven G Johnson. Nonlinear harmonic generation and devices in doubly resonant kerr cavities. *Physical Review A*, 79(1):013812, 2009. (Cited on page 113).

[87] D Hazra, S Jebari, R Albert, F Blanchet, A Grimm, C Chapelier, and M Hofheinz. Microwave response and electrical transport studies of disordered s wave superconductor: Nbn thin films. *arXiv preprint arXiv:1806.03935*, 2018. (Cited on pages 108 and 109).

[88] David Hilbert. *Ueber die stetige Abbildung einer Linie auf ein Flchenstck*. URL https://www.digizeitschriften.de/id/235181684_0038|log40. (Cited on page 131).

[89] Wei Hong, Zhi Hao Jiang, Chao Yu, Debin Hou, Haiming Wang, Chong Guo, Yun Hu, Le Kuai, Yingrui Yu, Zhengbo Jiang, et al. The role of millimeter-wave technologies in 5g/6g wireless communications. *IEEE Journal of Microwaves*, 1(1):101–122, 2021. (Cited on page 183).

[90] Sihao Huang, Benjamin Lienhard, Greg Calusine, Antti Vepsäläinen, Jochen Braumüller, David K Kim, Alexander J Melville, Bethany M Niedzielski, Jonilyn L Yoder, Bharath Kannan, et al. Microwave package design for su-

perconducting quantum processors. *PRX Quantum*, 2(2):020306, April 2021. doi:10.1103/PRXQuantum.2.020306. (Cited on pages 186, 188, 198, 200, and 216).

[91] Cheng-Fu Hung, An-Shyi Liu, Chih-Hung Chien, C.-L. Wang, and Ruey-Beei Wu. Bandwidth enhancement on waveguide transition to conductor backed CPW with high dielectric constant substrate. *IEEE Microwave and Wireless Components Letters*, 15 (2):128–130, February 2005. doi:10.1109/lmwc.2004.842862. (Cited on page 201).

[92] A. Imamolu, H. Schmidt, G. Woods, and M. Deutsch. Strongly interacting photons in a nonlinear cavity. *Physical Review Letters*, 79(8):1467–1470, August 1997. ISSN 1079-7114. doi:10.1103/PhysRevLett.79.1467. (Cited on page 119).

[93] Yachin Ivry, Chung-Soo Kim, Andrew E Dane, Domenico De Fazio, Adam N McCaughan, Kristen A Sunter, Qingyuan Zhao, and Karl K Berggren. Universal scaling of the critical temperature for thin films near the superconducting-to-insulating transition. *Physical Review B*, 90(21):214515, 2014. (Cited on pages 97 and 100).

[94] A.K. Jastrzebski, J.G.M. Yip, and Daiqing Li. Waveguide packaging at mm-wave frequencies. In *14th International Conference on Microwaves, Radar and Wireless Communications. MIKON - 2002. Conference Proceedings (IEEE Cat.No.02EX562)*, volume 1, pages 121–135 vol.1, May 2002. doi:10.1109/mikon.2002.1017820. (Cited on page 197).

[95] Blaise Jeanneret and SP Benz. Application of the josephson effect in electrical metrology. *The European Physical Journal Special Topics*, 172(1):181–206, 2009. (Cited on page 142).

[96] Wentao Jiang, Christopher J Sarabalis, Yanni D Dahmani, Rishi N Patel, Felix M Mayor, Timothy P McKenna, Raphaël Van Laer, and Amir H Safavi-Naeini. Efficient bidirectional piezo-optomechanical transduction between microwave and optical frequency. *Nature communications*, 11(1):1166, 2020. (Cited on page 5).

[97] Sun Jing, Liang Fa-guo, Han Li-hua, Sun Xiao-ying, and ZhengYan-qiu. Waveguide-to-Microstrip Antipodal Finline Transition at W Band. In *2013 Third International Conference on Instrumentation, Measurement, Computer, Communication and Control*, pages 510–513, September 2013. doi:10.1109/imccc.2013.116. (Cited on pages 198, 199, and 201).

[98] N. M. Jisrawi, M. W. Ruckman, T. R. Thurston, G. Reisfeld, M. Weinert, M. Strongin, and M. Gurvitch. Reversible depression in the T_c of thin Nb films due to enhanced hydrogen adsorption. *Phys. Rev. B*, 58(10):6585–6590, September 1998. ISSN 2469-9969. doi:10.1103/PhysRevB.58.6585. (Cited on page 161).

[99] Chaitali Joshi, Wenyuan Chen, Henry G. LeDuc, Peter K. Day, and Mohammad Mirhosseini. Strong kinetic-inductance kerr nonlinearity with titanium nitride nanowires. *Physical Review Applied*, 18(6):064088, December 2022. ISSN 2331-7019.

doi:10.1103/PhysRevApplied.18.064088. (Cited on pages 69, 118, 121, 122, 132, 134, and 135).

[100] Ch Kaiser, JM Meckbach, KS Ilin, J Lisenfeld, R Schäfer, AV Ustinov, and M Siegel. Aluminum hard mask technique for the fabrication of high quality submicron nb/al–alox/nb josephson junctions. *Superconductor science and technology*, 24(3):035005, 2010. (Cited on pages 142, 143, and 173).

[101] Anand Kamlapure, Mintu Mondal, Madhavi Chand, Archana Mishra, John Jesudasan, Vivas Bagwe, L Benfatto, Vikram Tripathi, and Pratap Raychaudhuri. Measurement of magnetic penetration depth and superconducting energy gap in very thin epitaxial nbn films. *Applied Physics Letters*, 96(7):072509, 2010. (Cited on pages 97 and 100).

[102] Fei Ke, Olivia Chen, Yanzhi Wang, and Nobuyuki Yoshikawa. Demonstration of a 47.8 ghz high-speed fft processor using single-flux-quantum technology. *IEEE Transactions on Applied Superconductivity*, 31(5):1–5, 2021. (Cited on page 2).

[103] Anthony R Kerr, S-K Pan, Stéphane MX Claude, Philip Dindo, Arthur W Lichtenberger, and Eugene F Lauria. Development of the alma-north america sideband-separating sis mixers. In *2013 IEEE MTT-S International Microwave Symposium Digest (MTT)*, pages 1–4. Ieee, 2013. (Cited on page 97).

[104] John Kerr. Xl. a new relation between electricity and light: Dielectrified media birefringent. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 50(332):337–348, 1875. (Cited on page 70).

[105] Khaled Khalaf, Vojkan Vidojkovic, Piet Wambacq, and John R. Long. *Data Transmission at Millimeter Waves*. Springer, Berlin, Germany, 2015. ISBN 978-3-662-46938-5. (Cited on page 2).

[106] M S Khalil, M J A Stoutimore, F C Wellstood, and K D Osborn. An analysis method for asymmetric resonator transmission applied to superconducting devices. *Journal of Applied Physics*, 111(5):054510, March 2012. doi:10.1063/1.3692073. (Cited on pages 59, 63, 64, 75, 105, 106, 174, 210, 211, and 212).

[107] Raphaël Khan, F Massel, and T T Heikkilä. Cross-kerr nonlinearity in optomechanical systems. *Physical Review A*, 91(4):043822, 2015. (Cited on page 116).

[108] Aditya Shreyas Kher. *Superconducting nonlinear kinetic inductance devices*. PhD thesis, California Institute of Technology, 2017. (Cited on pages 109 and 112).

[109] Sunmi Kim, Hirotaka Terai, Taro Yamashita, Wei Qiu, Tomoko Fuse, Fumiki Yoshihara, Sahel Ashhab, Kunihiro Inomata, and Kouichi Semba. Enhanced coherence of all-nitride superconducting qubits epitaxially grown on silicon substrate. *Communications Materials*, 2(1):1–7, December 2021. doi:10.1038/s43246-021-00204-4. (Cited on pages 77, 143, 166, 173, 180, 196, and 217).

[110] J Kitaygorsky, I Komissarov, A Jukna, D Pan, O Minaeva, N Kaurova, A Divochiy, A Korneev, M Tarkhov, B Voronov, et al. Dark counts in nanostructured nbn superconducting single-photon detectors and bridges. *IEEE Transactions on Applied Superconductivity*, 17(2):275–278, 2007. (Cited on page 106).

[111] Morten Kjaergaard, Mollie E Schwartz, Jochen Braumüller, Philip Krantz, Joel I-J Wang, Simon Gustavsson, and William D Oliver. Superconducting qubits: Current state of play. *Annual Review of Condensed Matter Physics*, 11:369–395, 2020. doi:10.1146/annurev-conmatphys-031119-050605. (Cited on pages 2, 84, 118, 142, and 172).

[112] Alan W Kleinsasser, Ronald E Miller, and William H Mallison. Dependence of critical current density on oxygen exposure in nb-alo/sub x/-nb tunnel junctions. *IEEE transactions on applied superconductivity*, 5(1):26–30, 1995. (Cited on page 152).

[113] Jens Koch, M Yu Terri, Jay Gambetta, Andrew A Houck, David I Schuster, Johannes Majer, Alexandre Blais, Michel H Devoret, Steven M Girvin, and Robert J Schoelkopf. Charge-insensitive qubit design derived from the cooper pair box. *Physical Review A*, 76(4):042319, 2007. (Cited on pages 85, 86, 88, 91, 121, 172, 186, 188, 190, 222, and 228).

[114] U Kongpop, Ari D Brown, Samuel H Moseley, Omid Noroozian, Edward J Wollack, et al. A cryogenic waveguide mount for microstrip circuit and material characterization. *IEEE Transactions on Applied Superconductivity*, 27(4):1–4, 2016. (Cited on page 97).

[115] NB Kopnin. Theory of superconductivity. *Helsinki University of Technology*, 2006. (Cited on pages 18, 20, 21, 26, and 27).

[116] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver. A quantum engineer's guide to superconducting qubits. *Applied Physics Reviews*, 6(2), June 2019. doi:10.1063/1.5089550. (Cited on pages 84 and 86).

[117] S. Krinner, S. Storz, P. Kurpiers, P. Magnard, J. Heinsoo, R. Keller, J. Lütolf, C. Eichler, and A. Wallraff. Engineering cryogenic setups for 100-qubit scale superconducting circuit systems. *EPJ Quantum Technology*, 6(1):1–29, December 2019. ISSN 2196-0763. doi:10.1140/epjqt/s40507-019-0072-0. (Cited on page 4).

[118] Stefan Kuhr, Sébastien Gleyzes, Christine Guerlin, Julien Bernu, U Busk Hoff, Samuel Deléglise, Stefano Osnaghi, Michel Brune, J-M Raimond, Serge Haroche, et al. Ultra-high finesse fabry-pérot superconducting resonator. *Applied Physics Letters*, 90(16): 164101, 2007. (Cited on page 97).

[119] Aishwarya Kumar, Aziza Suleymanzade, Mark Stone, Lavanya Taneja, Alexander Anferov, David I Schuster, and Jonathan Simon. Quantum-enabled millimetre wave to optical transduction using neutral atoms. *Nature*, 615(7953):614–619, 2023. (Cited on pages 5, 44, 182, and 251).

[120] Ray Kurzweil. *The age of spiritual machines: When computers exceed human intelligence*. Penguin, 2000. (Cited on page 1).

[121] Prem K. Kythe. *Transmission Lines and*. Routledge Handbooks Online, March 2019. ISBN 978-1-13874847-7. doi:10.1201/9781315180236-22. (Cited on page 50).

[122] Wai M Lee. Method for removing etching residue using a hydroxylamine-containing composition, January 9 1996. US Patent 5,482,566. (Cited on pages 161 and 185).

[123] Wai Mun Lee, Katy Ip, Xuan-Dung Dinh, and David John Maloney. Semiconductor process residue removal composition and process, November 30 2004. US Patent 6,825,156. (Cited on pages 161 and 185).

[124] Yongshik Lee, J.P. Becker, J.R. East, and L.P.B. Katehi. Fully micromachined finite-ground coplanar line-to-waveguide transitions for W-band applications. *IEEE Transactions on Microwave Theory and Techniques*, 52(3):1001–1007, March 2004. doi:10.1109/tmtt.2004.823580. (Cited on pages 198 and 199).

[125] Antonio Leo, G Grimaldi, Roberta Citro, Angela Nigro, Sandro Pace, and RP Huebener. Quasiparticle scattering time in niobium superconducting films. *Physical Review B*, 84(1):014536, 2011. (Cited on pages 5 and 21).

[126] Edward Leonard Jr, Matthew A Beck, J Nelson, Brad G Christensen, Ted Thorbeck, Caleb Howington, Alexander Opremcak, Ivan V Pechenezhskiy, Kenneth Dodge, Nicholas P Dupuis, et al. Digital coherent control of a superconducting qubit. *Physical Review Applied*, 11(1):014009, 2019. (Cited on pages 4, 143, and 182).

[127] K.K. Likharev and V.K. Semenov. Rsfq logic/memory family: a new josephson-junction technology for sub-terahertz-clock-frequency digital systems. *IEEE Transactions on Applied Superconductivity*, 1(1):3–28, 1991. doi:10.1109/77.80745. (Cited on page 142).

[128] Andrew Lingenfelter, David Roberts, and A. A. Clerk. Unconditional fock state generation using arbitrarily weak photonic nonlinearities. *Science Advances*, 7(48), November 2021. ISSN 2375-2548. doi:10.1126/sciadv.abj1916. (Cited on pages 119, 132, and 136).

[129] J Lisenfeld, A Lukashenko, M Ansmann, JM Martinis, and AV Ustinov. Temperature dependence of coherent oscillations in josephson phase qubits. *Physical review letters*, 99(17):170504, 2007. (Cited on pages 142, 143, 173, 181, and 193).

[130] Jürgen Lisenfeld, Alexander Bilmes, and Alexey V. Ustinov. Enhancing the coherence of superconducting quantum bits with electric fields. *npj Quantum Inf.*, 9(8):1–5, January 2023. ISSN 2056-6387. doi:10.1038/s41534-023-00678-9. (Cited on page 67).

[131] Chuan-Hong Liu, Andrew Ballard, David Olaya, Daniel R Schmidt, John Biesecker, Tammy Lucas, Joel Ullom, Shravan Patel, Owen Rafferty, Alexander Opremcak, et al. Single flux quantum-based digital control of superconducting qubits in a multi-chip module. *arXiv preprint arXiv:2301.05696*, 2023. (Cited on pages 4, 143, and 182).

[132] Evgeniya H Lock, Peng Xu, Timothy Kohler, Lizmarie Camacho, Joseph Prestigiacomo, Yaniv J Rosen, and Kevin D Osborn. Using surface engineering to modulate superconducting coplanar microwave resonator performance. *IEEE Transactions on Applied Superconductivity*, 29(6):1–8, 2019. (Cited on page 288).

[133] Ryan Lott. 28gbps microstrip with pepper jack cheese as substrate, October 2016. [Online; accessed 12. Oct. 2023]. (Cited on page 48).

[134] AN Lykov. The superconducting mixed state of artificial microstructures. *Advances in Physics*, 42(3):263–342, 1993. (Cited on page 154).

[135] P. Magnard, P. Kurpiers, B. Royer, T. Walter, J.-C. Besse, S. Gasparinetti, M. Pechal, J. Heinsoo, S. Storz, A. Blais, and A. Wallraff. Fast and unconditional all-microwave reset of a superconducting qubit. *Physical Review Letters*, 121(6):060502, August 2018. ISSN 1079-7114. doi:10.1103/PhysRevLett.121.060502. (Cited on pages 4 and 183).

[136] Paul Magnard, Simon Storz, Philipp Kurpiers, Josua Schär, Fabian Marxer, Janis Lütolf, Theo Walter, J-C Besse, Mihai Gabureac, Kevin Reuer, et al. Microwave quantum link between superconducting circuits housed in spatially separated cryogenic systems. *Physical Review Letters*, 125(26):260502, 2020. (Cited on page 4).

[137] N. Maleeva, L. Grünhaupt, T. Klein, F. Levy-Bertrand, O. Dupre, M. Calvo, F. Valenti, P. Winkel, F. Friedrich, W. Wernsdorfer, A. V. Ustinov, H. Rotzinger, A. Monfardini, M. V. Fistul, and I. M. Pop. Circuit quantum electrodynamics of granular aluminum resonators. *Nature Communications*, 9(3889):1–7, September 2018. ISSN 2041-1723. doi:10.1038/s41467-018-06386-9. (Cited on page 118).

[138] Nataliya Maleeva, Lukas Grünhaupt, T Klein, F Levy-Bertrand, O Dupre, M Calvo, F Valenti, P Winkel, F Friedrich, W Wernsdorfer, et al. Circuit quantum electrodynamics of granular aluminum resonators. *Nature communications*, 9(1):3889, 2018. (Cited on pages 109 and 110).

[139] Maciej Malinowski. How fast are quantum computers (part 2: clock speeds). *Reading the quantum*, December 2022. URL https://m-malinowski.github.io/2022/12/04/how-fast-are-quantum-computers-part-2.html. (Cited on page 2).

[140] M. Malnou, M. R. Vissers, J. D. Wheeler, J. Aumentado, J. Hubmayr, J. N. Ullom, and J. Gao. Three-wave mixing kinetic inductance traveling-wave amplifier with near-quantum-limited noise performance. *PRX Quantum*, 2(1):010302, January 2021. ISSN 2691-3399. doi:10.1103/PRXQuantum.2.010302. (Cited on pages 119 and 142).

[141] HJ Mamin, E Huang, S Carnevale, CT Rettner, N Arellano, MH Sherwood, C Kurter, B Trimm, M Sandberg, RM Shelby, et al. Merged-element transmons: Design and qubit performance. *Physical Review Applied*, 16(2):024023, 2021. (Cited on pages 84, 172, 173, and 182).

[142] Benoit B. Mandelbrot. *Fractals: Form, Chance and Dimension*. W.H.Freeman & Company, January 1977. ISBN 978-0-71670473-7. (Cited on page 131).

[143] John M Martinis, Sae Nam, Joe Aumentado, and C Urbina. Rabi oscillations in a large josephson-junction qubit. *Physical review letters*, 89(11):117901, 2002. (Cited on pages 142, 143, and 173).

[144] John M Martinis, M Ansmann, and J Aumentado. Energy decay in superconducting josephson-junction qubits from nonequilibrium quasiparticle excitations. *Physical review letters*, 103(9):097002, 2009. (Cited on pages 5, 21, 31, 118, and 183).

[145] Daniel C Mattis and John Bardeen. Theory of the anomalous skin effect in normal and superconducting metals. *Physical Review*, 111(2):412, July 1958. doi:10.1103/PhysRev.111.412. (Cited on pages 31, 66, 108, 118, 175, 181, and 213).

[146] Benjamin A Mazin. *Microwave kinetic inductance detectors*. PhD thesis, California Institute of Technology, 2005. (Cited on page 114).

[147] R McDermott and MG Vavilov. Accurate qubit control with single flux quantum pulses. *Physical Review Applied*, 2(1):014007, 2014. (Cited on pages 4, 143, and 182).

[148] Robert McDermott. Materials origins of decoherence in superconducting qubits. *IEEE Transactions on Applied Superconductivity*, 19(1):2–13, February 2009. doi:10.1109/tasc.2008.2012255. (Cited on pages 142, 196, and 247).

[149] Alexander McDonald and Aashish A Clerk. Exact solutions of interacting dissipative systems via weak symmetries. *Physical Review Letters*, 128(3):033602, 2022. (Cited on page 139).

[150] C. R. H. McRae, H. Wang, J. Gao, M. R. Vissers, T. Brecht, A. Dunsworth, D. P. Pappas, and J. Mutus. Materials loss measurements using superconducting microwave resonators. *Review of Scientific Instruments*, 91(9):091101, September 2020. doi:10.1063/5.0017378. (Cited on page 196).

[151] C A Mears, Qing Hu, P L Richards, A H Worsham, D E Prober, and A V Räisänen. Quantum-limited heterodyne detection of millimeter waves using superconducting tantalum tunnel junctions. *Applied physics letters*, 57(23):2487–2489, 1990. (Cited on page 97).

[152] Owen Medeiros, Marco Colangelo, Ilya Charaev, and Karl K Berggren. Measuring thickness in thin nbn films for superconducting devices. *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, 37(4):041501, 2019. (Cited on pages 99 and 106).

[153] Anthony Megrant, Charles Neill, Rami Barends, Ben Chiaro, Yu Chen, Ludwig Feigl, Julian Kelly, Erik Lucero, Matteo Mariantoni, Peter J J O'Malley, et al. Planar superconducting resonators with internal quality factors above one million. *Applied Physics Letters*, 100(11):113510, 2012. (Cited on pages 59, 63, 64, 106, and 285).

[154] Robert Meservey and Paul M Tedrow. Measurements of the kinetic inductance of superconducting linear structures. *Journal of Applied Physics*, 40(5):2028–2034, 1969. (Cited on page 15).

[155] Zlatko K Minev, Zaki Leghtas, Shantanu O Mundhada, Lysander Christakis, Ioan M Pop, and Michel H Devoret. Energy-participation quantization of josephson circuits. *npj Quantum Information*, 7(1):131, 2021. (Cited on pages 65 and 222).

[156] Mohammad Mirhosseini, Alp Sipahigil, Mahmoud Kalaee, and Oskar Painter. Superconducting qubit to optical photon transduction. *Nature*, 588(7839):599–603, 2020. (Cited on pages 5 and 251).

[157] CJ Mogab, AC Adams, and DL Flamm. Plasma etching of si and sio2—the effect of oxygen additions to cf4 plasmas. *Journal of applied physics*, 49(7):3796–3803, 1978. (Cited on page 158).

[158] Baharak Mohajer-Iravani and Omar M. Ramahi. Wideband circuit model for planar ebg structures. *IEEE Transactions on Advanced Packaging*, 33(1):169–179, 2010. (Cited on page 37).

[159] Mintu Mondal, Anand Kamlapure, Madhavi Chand, Garima Saraswat, Sanjeev Kumar, John Jesudasan, L Benfatto, Vikram Tripathi, and Pratap Raychaudhuri. Phase fluctuations in a strongly disordered s-wave nbn superconductor close to the metal-insulator transition. *Physical review letters*, 106(4):047001, 2011. (Cited on pages 99 and 101).

[160] J E Mooij. Percolation, localization, and superconductivity. In *Nato Asi Series*, page 325. Plenum Press New York, 1984. (Cited on pages 100 and 106).

[161] J E Mooij and Yu V Nazarov. Superconducting nanowires as quantum phase-slip junctions. *Nature Physics*, 2(3):169, 2006. (Cited on page 116).

[162] Shin'ichi Morohashi and Shinya Hasuo. Experimental investigations and analysis for high-quality nb/al-alo x/nb josephson junctions. *Journal of applied physics*, 61(10): 4835–4849, 1987. (Cited on pages 82, 151, 152, 166, and 178).

[163] V.S. Mottonen and A.V. Raisanen. Novel wide-band coplanar waveguide-to-rectangular waveguide transition. *IEEE Transactions on Microwave Theory and Techniques*, 52(8):1836–1842, August 2004. doi:10.1109/tmtt.2004.831580. (Cited on page 199).

[164] R Movshovich, B Yurke, P G Kaminsky, A D Smith, A H Silver, R W Simon, and M V Schneider. Observation of zero-point noise squeezing via a josephson-parametric amplifier. *Physical review letters*, 65(12):1419, 1990. (Cited on page 116).

[165] Shyam P Muraka, Moshe Eizenberg, and Ashok K Sinha. *Interlayer dielectrics for semiconductor technologies*, volume 1. Elsevier, 2003. (Cited on page 66).

[166] Yasunobu Nakamura, Hirotaka Terai, Kunihiro Inomata, Tsuyoshi Yamamoto, Wei Qiu, and Zhen Wang. Superconducting qubits consisting of epitaxially grown nbn/aln/nbn josephson junctions. *Applied Physics Letters*, 99(21):212502, 2011. (Cited on pages 77, 143, and 173).

[167] Sang Boo Nam. Theory of electromagnetic properties of superconducting and normal systems. i. *Physical Review*, 156(2):470, 1967. (Cited on page 31).

[168] Michael A Nielsen and Isaac L Chuang. Quantum computation and quantum information. *Phys. Today*, 54(2):60, 2001. (Cited on page 91).

[169] David Niepce, Jonathan Burnett, and Jonas Bylander. High kinetic inductance nb n nanowire superinductors. *Physical Review Applied*, 11(4):044014, 2019. (Cited on pages 97, 99, 100, 106, 116, and 119).

[170] Simon E Nigg, Hanhee Paik, Brian Vlastakis, Gerhard Kirchmair, Shyam Shankar, Luigi Frunzio, MH Devoret, RJ Schoelkopf, and SM Girvin. Black-box superconducting circuit quantization. *Physical Review Letters*, 108(24):240502, 2012. (Cited on page 69).

[171] Yong Niu, Yong Li, Depeng Jin, Li Su, and Athanasios V Vasilakos. A survey of millimeter wave communications (mmwave) for 5g: opportunities and challenges. *Wireless networks*, 21(8):2657–2676, 2015. (Cited on page 97).

[172] Omid Noroozian, Emily Barrentine, Ari Brown, Giuseppe Cataldo, Negar Ehsan, Wen-Ting Hsieh, Thomas Stevenson, Kongpop U-yen, Edward Wollack, and Samuel Harvey Moseley. μ-spec: An efficient compact integrated spectrometer for submillimeter astrophysics. In *26TH International Symposium on Space Terahertz Technology*, 2015. (Cited on pages 3 and 97).

[173] V Novotny and PPM Meincke. Single superconducting energy gap in pure niobium. *Journal of Low Temperature Physics*, 18:147–157, 1975. (Cited on page 151).

[174] Ibrahim Nsanzineza. *Vortices and Quasiparticles in Superconducting Microwave Resonators*. PhD thesis, Syracuse University, 2016. (Cited on pages 21 and 106).

[175] David Olaya, Manuel Castellanos-Beltran, Javier F Pulecio, John Biesecker, Soroush Khadem, Theodore Lewitt, Peter Hopkins, Paul Dresselhaus, and Samuel Benz. Planarized process for single-flux-quantum circuits with self-shunted nb/nb$_x$si$_{1-x}$/nb josephson junctions. *IEEE Transactions on Applied Superconductivity*, 2019. (Cited on page 97).

[176] Aaron D O'Connell, M Ansmann, Radoslaw C Bialczak, Max Hofheinz, Nadav Katz, Erik Lucero, C McKenney, Matthew Neeley, Haohua Wang, Eva Maria Weig, et al. Microwave dielectric loss at single photon energies and millikelvin temperatures. *Applied Physics Letters*, 92(11), 2008. (Cited on pages 66, 179, and 285).

[177] Hanhee Paik, SK Dutta, RM Lewis, TA Palomaki, BK Cooper, RC Ramos, H Xu, AJ Dragt, JR Anderson, CJ Lobb, et al. Decoherence in dc squid phase qubits. *Physical Review B*, 77(21):214510, 2008. (Cited on pages 142, 143, and 173).

[178] Hanhee Paik, David I Schuster, Lev S Bishop, Gerhard Kirchmair, Gianluigi Catelani, Adam P Sears, BR Johnson, MJ Reagor, Luigi Frunzio, Leonid I Glazman, et al. Observation of high coherence in josephson junction qubits measured in a three-dimensional circuit qed architecture. *Physical Review Letters*, 107(24):240501, 2011. (Cited on pages 85, 143, 190, 223, and 227).

[179] David P. Pappas, Michael R. Vissers, David S. Wisbey, Jeffrey S. Kline, and Jiansong Gao. Two Level System Loss in Superconducting Microwave Resonators. *IEEE Transactions on Applied Superconductivity*, 21(3):871–874, June 2011. doi:10.1109/tasc.2010.2097578. (Cited on pages 67 and 213).

[180] G Peano. [a]sur une courbe qui remplit toute une aire plaine, [o] math, 1890. (Cited on page 131).

[181] Marek Pechal and Amir H Safavi-Naeini. Millimeter-wave interconnects for microwave-frequency quantum machines. *Physical Review A*, 96(4):042305, October 2017. doi:10.1103/PhysRevA.96.042305. (Cited on pages 4, 116, 122, 132, 196, and 251).

[182] Alexander PM Place, Lila VH Rodgers, Pranav Mundada, Basil M Smitham, Mattias Fitzpatrick, Zhaoqi Leng, Anjali Premkumar, Jacob Bryon, Andrei Vrajitoarea, Sara Sussman, et al. New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds. *Nature communications*, 12(1):1–6, 2021. (Cited on pages 172 and 173).

[183] Frank Pobell. *Matter and Methods at Low Temperatures*. Springer, Berlin, Germany, 2007. ISBN 978-3-540-46360-3. URL https://link.springer.com/book/10.1007/978-3-540-46360-3. (Cited on pages 101 and 221).

[184] David M Pozar. *Microwave engineering*. John Wiley & Sons, November 2009. (Cited on pages 45, 46, 47, 48, 55, 57, 60, 204, 205, 209, 232, 277, 278, and 280).

[185] Anjali Premkumar, Conan Weiland, Sooyeon Hwang, Berthold Jäck, Alexander PM Place, Iradwikanari Waluyo, Adrian Hunt, Valentina Bisogni, Jonathan Pelliciari, Andi Barbour, et al. Microscopic relaxation channels in materials for superconducting qubits. *Communications Materials*, 2(1):1–9, 2021. (Cited on pages 155, 156, 161, 163, 172, 173, and 185).

[186] S. Probst, F. B. Song, P. A. Bushev, A. V. Ustinov, and M. Weides. Efficient and robust analysis of complex scattering data under noise in microwave resonators. *Review of Scientific Instruments*, 86(2):024706, February 2015. ISSN 0034-6748. doi:10.1063/1.4907935. (Cited on pages 62 and 212).

[187] S. E. Rasmussen, K. S. Christensen, S. P. Pedersen, L. B. Kristensen, T. Bækkegaard, N. J. S. Loft, and N. T. Zinner. Superconducting circuit companion—an introduction with worked examples. *PRX Quantum*, 2(4):040204, December 2021. ISSN 2691-3399. doi:10.1103/PRXQuantum.2.040204. (Cited on pages 69, 78, and 86).

[188] Alexander P. Read, Benjamin J. Chapman, Chan U. Lei, Jacob C. Curtis, Suhas Ganjam, Lev Krayzman, Luigi Frunzio, and Robert J. Schoelkopf. Precision measurement of the microwave dielectric loss of sapphire in the quantum regime with parts-per-billion sensitivity. *Phys. Rev. Appl.*, 19(3):034064, March 2023. ISSN 2331-7019. doi:10.1103/PhysRevApplied.19.034064. (Cited on pages 178, 179, and 285).

[189] Matthew Reagor, Wolfgang Pfaff, Christopher Axline, Reinier W. Heeres, Nissim Ofek, Katrina Sliwa, Eric Holland, Chen Wang, Jacob Blumoff, Kevin Chou, Michael J. Hatridge, Luigi Frunzio, Michel H. Devoret, Liang Jiang, and Robert J. Schoelkopf. Quantum memory with millisecond coherence in circuit qed. *Physical Review B*, 94 (1):014506, July 2016. ISSN 2469-9969. doi:10.1103/PhysRevB.94.014506. (Cited on pages 69, 108, 143, and 193).

[190] Matthew James Reagor. *Superconducting Cavities for Circuit Quantum Electrodynamics*. PhD thesis, Yale University, 2016. (Cited on page 213).

[191] Patrick Rebentrost and Frank K Wilhelm. Optimal control of a leaking qubit. *Physical Review B*, 79(6):060507, 2009. (Cited on page 193).

[192] MD Reed, L DiCarlo, BR Johnson, L Sun, DI Schuster, L Frunzio, and RJ Schoelkopf. High-fidelity readout in circuit quantum electrodynamics using the jaynes-cummings nonlinearity. *Physical review letters*, 105(17):173601, 2010. (Cited on pages 222 and 223).

[193] G Rickayzen. Critical magnetic field of a thin film of superconductor. *Physical Review*, 138(1a):A73, 1965. (Cited on page 20).

[194] Alessandro Ridolfo, Martin Leib, Salvatore Savasta, and Michael J Hartmann. Photon blockade in the ultrastrong coupling regime. *Physical review letters*, 109(19):193602, 2012. (Cited on page 74).

[195] Jeremy M Sage, Vladimir Bolkhovsky, William D Oliver, Benjamin Turek, and Paul B Welander. Study of loss in superconducting coplanar waveguide resonators. *Journal of Applied Physics*, 109(6):063915, March 2011. doi:10.1063/1.3552890. (Cited on pages 106 and 213).

[196] N Samkharadze, A Bruno, P Scarlino, G Zheng, D P DiVincenzo, L DiCarlo, and L M K Vandersypen. High-kinetic-inductance superconducting nanowire resonators for circuit qed in a magnetic field. *Physical Review Applied*, 5(4):044004, 2016. (Cited on page 97).

[197] Jay N Sasserath and John Vivalda. Sloped niobium etching using cf4 and o2. *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, 8(6):3914–3919, 1990. (Cited on page 158).

[198] Brendan Saxberg, Andrei Vrajitoarea, Gabrielle Roberts, Margaret G Panetta, Jonathan Simon, and David I Schuster. Disorder-assisted assembly of strongly correlated fluids of light. *Nature*, 612(7940):435–441, 2022. (Cited on pages 172 and 173).

[199] Yannick Schön, Jan Nicolas Voss, Micha Wildermuth, Andre Schneider, Sebastian T. Skacel, Martin P. Weides, Jared H. Cole, Hannes Rotzinger, and Alexey V. Ustinov. Rabi oscillations in a superconducting nanowire circuit. *npj Quantum Materials*, 5(18): 1–5, March 2020. ISSN 2397-4648. doi:10.1038/s41535-020-0220-x. (Cited on pages 118 and 136).

[200] David Isaac Schuster. *Circuit quantum electrodynamics*. PhD thesis, Yale University, 2007. (Cited on pages 91 and 228).

[201] DI Schuster, Andreas Wallraff, Alexandre Blais, L Frunzio, R-S Huang, J Majer, SM Girvin, and RJ Schoelkopf. ac stark shift and dephasing of a superconducting qubit strongly coupled to a cavity field. *Physical Review Letters*, 94(12):123602, March 2005. ISSN 1079-7114. doi:10.1103/PhysRevLett.94.123602. (Cited on pages 188 and 229).

[202] Kartik Senapati, Mark G Blamire, and Zoe H Barber. Phase periodic conductance oscillations at subgap andreev resonances in nb–al–alo x–al–nb tunnel junctions. *Applied Physics Letters*, 97(16):162503, 2010. (Cited on pages 152 and 167).

[203] K Serniak, S Diamond, M Hays, V Fatemi, S Shankar, L Frunzio, RJ Schoelkopf, and MH Devoret. Direct dispersive monitoring of charge parity in offset-charge-sensitive transmons. *Physical Review Applied*, 12(1):014052, 2019. (Cited on page 173).

[204] Abigail Shearrow, Gerwin Koolstra, Samuel J Whiteley, Nathan Earnest, Peter S Barry, F Joseph Heremans, David D Awschalom, Erik Shirokoff, and David I Schuster. Atomic layer deposition of titanium nitride for quantum circuits. *Applied Physics Letters*, 113 (21):212601, 2018. (Cited on pages 97, 101, 109, 119, 130, 133, 134, and 287).

[205] Erik Shirokoff, Peter S Barry, Charles M Bradford, Goutam Chattopadhyay, Peter Day, Simon Doyle, Steve Hailey-Dunsheath, Matthew I Hollister, Attila Kovács, Christopher McKenney, et al. Mkid development for superspec: an on-chip, mm-wave, filter-bank spectrometer. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI*, volume 8452, page 84520r. International Society for Optics and Photonics, Spie, September 2012. doi:10.1117/12.927070. (Cited on pages 97, 196, and 216).

[206] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, July 2006. URL https://epubs.siam.org/doi/10.1137/S0097539795293172. (Cited on page 3).

[207] Wacław Sierpinski. Sur une nouvelle courbe continue qui remplit toute une aire plane. *Bull. Acad. Sci. Cracovie (Sci. math. et nat. Serie A)*, pages 462–478, 1912. (Cited on page 127).

[208] M A Skvortsov and M V Feigel'man. Superconductivity in disordered thin films: giant mesoscopic fluctuations. *Physical review letters*, 95(5):057002, 2005. (Cited on page 100).

[209] Aaron Somoroff, Quentin Ficheux, Raymond A. Mencia, Haonan Xiong, Roman Kuzmin, and Vladimir E. Manucharyan. Millisecond coherence in a superconducting qubit. *Phys. Rev. Lett.*, 130(26):267001, June 2023. ISSN 1079-7114. doi:10.1103/PhysRevLett.130.267001. (Cited on page 251).

[210] Mark J Sowa, Yonas Yemane, Jinsong Zhang, Johanna C Palmstrom, Ling Ju, Nicholas C Strandwitz, Fritz B Prinz, and J Provine. Plasma-enhanced atomic layer deposition of superconducting niobium nitride. *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, 35(1):01b143, 2017. (Cited on pages 97, 98, 100, 286, and 287).

[211] Gheorghe Stan, Stuart B Field, and John M Martinis. Critical field for complete vortex expulsion from narrow superconducting strips. *Physical review letters*, 92(9):097003, 2004. (Cited on page 113).

[212] Peter Stano and Daniel Loss. Review of performance metrics of spin qubits in gated semiconducting nanostructures. *Nat. Rev. Phys.*, 4:672–688, October 2022. ISSN 2522-5820. (Cited on page 2).

[213] Leandro Stefanazzi, Kenneth Treptow, Neal Wilcer, Chris Stoughton, Collin Bradford, Sho Uemura, Silvia Zorzetti, Salvatore Montella, Gustavo Cancelo, Sara Sussman, et al. The qick (quantum instrumentation control kit): Readout and control for qubits and detectors. *Review of Scientific Instruments*, 93(4):044709, 2022. (Cited on page 171).

[214] WC Stewart. Current-voltage characteristics of josephson junctions. *Applied physics letters*, 12(8):277–280, 1968. (Cited on page 151).

[215] Hubert Stokowski, Marek Pechal, Emma Snively, K. S. Kevin Multani, Paul B. Welander, Jeremy Witmer, Emilio A. Nanni, and Amir H. Safavi-Naeini. Towards Millimeter-Wave Based Quantum Networks. In *2019 44th International Conference on Infrared, Millimeter, and Terahertz Waves (IRMMW-THz)*, pages 1–2, September 2019. doi:10.1109/IRMMW-THz.2019.8874171. (Cited on pages 196, 201, and 216).

[216] H Sugiyama, A Fujimaki, and H Hayakawa. Characteristics of high critical current density josephson junctions with nb/alo/sub x//nb trilayers. *IEEE Transactions on Applied Superconductivity*, 5(2):2739–2742, 1995. (Cited on page 152).

[217] Aziza Suleymanzade, Alexander Anferov, Mark Stone, Ravi K. Naik, Andrew Oriani, Jonathan Simon, and David Schuster. A tunable high-q millimeter wave cavity for hybrid circuit and cavity qed experiments. *Appl. Phys. Lett.*, 116(10), March 2020. ISSN 0003-6951. doi:10.1063/1.5137900. (Cited on page 44).

[218] Y. Sunada, S. Kono, J. Ilves, S. Tamate, T. Sugiyama, Y. Tabuchi, and Y. Nakamura. Fast readout and reset of a superconducting qubit coupled to a resonator with an intrinsic purcell filter. *Phys. Rev. Appl.*, 17(4):044016, April 2022. ISSN 2331-7019. doi:10.1103/PhysRevApplied.17.044016. (Cited on page 2).

[219] L J Swenson, P K Day, B H Eom, H G Leduc, N Llombart, C M McKenney, O Noroozian, and J Zmuidzinas. Operation of a titanium nitride superconducting microresonator detector in the nonlinear regime. *Journal of Applied Physics*, 113(10): 104501, 2013. (Cited on pages 75, 110, 112, and 114).

[220] Evgeny F Talantsev and Jeffery L Tallon. Universal self-field critical current for thin-film superconductors. *Nature communications*, 6:7820, 2015. (Cited on page 113).

[221] S Tasker, RD Chambers, and JPS Badyal. Surface defluorination of ptfe by sodium atoms. *The Journal of Physical Chemistry*, 98(47):12442–12446, 1994. (Cited on page 160).

[222] Erik A Tholén, Adem Ergül, Evelyn M Doherty, Frank M Weber, Fabien Grégis, and David B Haviland. Nonlinearities and parametric amplification in superconducting coplanar waveguide resonators. *Applied physics letters*, 90(25):253509, 2007. (Cited on pages 115 and 116).

[223] JH Thomas III and LH Hammer. An x-ray photoelectron spectroscopy study of sulfur hexafluoride etchant residue on silicon and silicon dioxide. *Journal of Vacuum Science & Technology B: Microelectronics Processing and Phenomena*, 5(6):1617–1621, 1987. (Cited on page 158).

[224] Michael Tinkham. *Introduction to superconductivity*. Courier Corporation, 2004. (Cited on pages 9, 21, 23, 24, 25, 26, 27, 29, 32, 81, 82, 108, and 213).

[225] Sergey K Tolpygo and Denis Amparo. Fabrication-process-induced variations of nb/al/alox/nb josephson junctions in superconductor integrated circuits. *Superconductor Science and Technology*, 23(3):034024, 2010. (Cited on pages 151 and 152).

[226] Sergey K Tolpygo and Vasili K Semenov. Increasing integration scale of superconductor electronics beyond one million josephson junctions. In *Journal of Physics: Conference Series*, volume 1559, page 012002. IOP Publishing, 2020. (Cited on pages 2 and 142).

[227] Sergey K Tolpygo, Denis Amparo, Richard T Hunt, John A Vivalda, and Daniel T Yohannes. Diffusion stop-layers for superconducting integrated circuits and qubits with nb-based josephson junctions. *IEEE transactions on applied superconductivity*, 21(3):119–125, 2010. (Cited on page 145).

[228] Sergey K. Tolpygo, Vladimir Bolkhovsky, Terence J. Weir, Leonard M. Johnson, Mark A. Gouker, and William D. Oliver. Fabrication process and properties of fully-planarized deep-submicron nb/al– AlO_x/Nb josephson junctions for vlsi circuits. *IEEE Transactions on Applied Superconductivity*, 25(3):1–12, 2015. doi:10.1109/tasc.2014.2374836. (Cited on pages 144, 151, 164, and 166).

[229] Osami Tsukamoto. Ac losses in a type ii superconductor strip with inhomogeneous critical current distribution. *Superconductor Science and Technology*, 18(5):596, 2005. (Cited on page 113).

[230] John R Tucker and Marc J Feldman. Quantum detection at millimeter wavelengths. *Reviews of Modern Physics*, 57(4):1055, 1985. (Cited on page 97).

[231] JP Turneaure and Ira Weissman. Microwave surface resistance of superconducting niobium. *Journal of Applied Physics*, 39(9):4417–4427, 1968. (Cited on page 151).

[232] Kongpop U-yen, Ari D. Brown, Samuel H. Moseley, Omid Noroozian, and Edward J. Wollack. A Cryogenic Waveguide Mount for Microstrip Circuit and Material Characterization. *IEEE Transactions on Applied Superconductivity*, 27(4):1–4, June 2017. doi:10.1109/tasc.2016.2645740. (Cited on pages 196 and 216).

[233] Sergio O Valenzuela, William D Oliver, David M Berns, Karl K Berggren, Leonid S Levitov, and Terry P Orlando. Microwave-induced cooling of a superconducting qubit. *Science*, 314(5805):1589–1592, 2006. (Cited on pages 4, 183, and 251).

[234] L. Valletti, S. Fantauzzi, and F. Di Paolo. An Innovative Lens Type FinLine to Microstrip Transition. *Journal of Infrared, Millimeter, and Terahertz Waves*, 43(7):628–653, August 2022. doi:10.1007/s10762-022-00869-z. (Cited on page 198).

[235] S van Berkel. *Characterization of Printed Transmission Lines at High Frequencies*. PhD thesis, 2015. (Cited on pages 66 and 197).

[236] J Verjauw, A Potočnik, M Mongillo, R Acharya, F Mohiyaddin, G Simion, A Pacco, Ts Ivanov, D Wan, A Vanleenhove, et al. Investigation of microwave loss induced by oxide regrowth in high-q niobium resonators. *Physical Review Applied*, 16(1):014018, July 2021. doi:10.1103/PhysRevApplied.16.014018. (Cited on pages 143, 144, 145, 156, 161, 178, 179, 180, 182, 185, 214, 215, 216, and 217).

[237] I. V. Vernik, A. F. Kirichenko, O. A. Mukhanov, and T. A. Ohki. Energy-efficient and compact ersfq decoder for cryogenic ram. *IEEE Transactions on Applied Superconductivity*, 27(4):1–5, 2017. doi:10.1109/tasc.2016.2646926. (Cited on pages 2 and 142).

[238] R Vijay, MH Devoret, and I Siddiqi. Invited review article: The josephson bifurcation amplifier. *Review of Scientific Instruments*, 80(11):111101, 2009. (Cited on page 142).

[239] Daniel F Walls and Gerard J Milburn. *Quantum optics*. Springer Science & Business Media, Berlin, Germany, 2007. ISBN 978-3-540-28574-8. (Cited on pages 56, 57, 62, 70, 72, 73, 74, 90, 114, and 137).

[240] Changqing Wang, Ivan Gonin, Anna Grassellino, Sergey Kazakov, Alexander Romanenko, Vyacheslav P Yakovlev, and Silvia Zorzetti. High-efficiency microwave-optical quantum transduction based on a cavity electro-optic superconducting system with long coherence time. *npj Quantum Information*, 8(1):149, 2022. (Cited on pages 5 and 251).

[241] Chen Wang, Christopher Axline, Yvonne Y Gao, Teresa Brecht, Yiwen Chu, Luigi Frunzio, MH Devoret, and Robert J Schoelkopf. Surface participation and dielectric loss in superconducting qubits. *Applied Physics Letters*, 107(16):162601, 2015. (Cited on pages 65, 66, 156, 163, 173, 177, 178, and 192).

[242] H Wang, M Hofheinz, J Wenner, M Ansmann, R C Bialczak, M Lenander, Erik Lucero, M Neeley, A D O'Connell, D Sank, et al. Improving the coherence time of superconducting coplanar resonators. *Applied Physics Letters*, 95(23):233508, December 2009. doi:10.1063/1.3273372. (Cited on pages 67, 106, 175, 178, and 213).

[243] Haozhi Wang, S. Singh, C. R. H. McRae, J. C. Bardin, S.-X. Lin, N. Messaoudi, A. R. Castelli, Y. J. Rosen, E. T. Holland, D. P. Pappas, and J. Y. Mutus. Cryogenic single-port calibration for superconducting microwave resonator measurements. *Quantum Sci. Technol.*, 6(3):035015, June 2021. ISSN 2058-9565. doi:10.1088/2058-9565/ac070e. (Cited on page 204).

[244] Xiong Wang, Linghe Kong, Fanxin Kong, Fudong Qiu, Mingyu Xia, Shlomi Arnon, and Guihai Chen. Millimeter wave communication: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 20(3):1616–1653, 2018. doi:10.1109/comst.2018.2844322. (Cited on pages 2, 39, and 251).

[245] Zhixin Wang, Mingrui Xu, Xu Han, Wei Fu, Shruti Puri, S. M. Girvin, Hong X. Tang, S. Shankar, and M. H. Devoret. Quantum microwave radiometry with a superconducting qubit. *Physical Review Letters*, 126(18):180501, May 2021. ISSN 1079-7114. doi:10.1103/PhysRevLett.126.180501. (Cited on pages 4 and 183).

[246] Koki Watanabe, Keiji Yoshida, and Takeshi Aoki Kohjiro. Kinetic inductance of superconducting coplanar waveguides. *Jpn. J. Appl. Phys.*, 33(10r):5708, October 1994. ISSN 1347-4065. doi:10.1143/jjap.33.5708. (Cited on page 50).

[247] M Weides, RC Bialczak, M Lenander, E Lucero, Matteo Mariantoni, M Neeley, AD O'Connell, D Sank, H Wang, J Wenner, et al. Phase qubits fabricated with trilayer junctions. *Superconductor Science and Technology*, 24(5):055005, 2011. (Cited on pages 142, 143, and 173).

[248] Cao Wen-Hui, Yu Hai-Feng, Tian Ye, Yu Hong-Wei, Ren Yu-Feng, Chen Geng-Hua, and Zhao Shi-Ping. Nb/al-alox/nb junctions with controllable critical current density for qubit application. *Chinese Physics B*, 18(11):5044, 2009. (Cited on page 152).

[249] J. Wenner, M. Neeley, Radoslaw C. Bialczak, M. Lenander, Erik Lucero, A. D. O'Connell, D. Sank, H. Wang, M. Weides, A. N. Cleland, and John M. Martinis.

Wirebond crosstalk and cavity modes in large chip mounts for superconducting. *Supercond. Sci. Technol.*, 24(6):065001, March 2011. ISSN 0953-2048. doi:10.1088/0953-2048/24/6/065001. (Cited on page 188).

[250] Samuel James Whiteley. *Spin-phonon Interactions with Defects in Silicon Carbide*. PhD thesis, The University of Chicago, 2019. (Cited on page 57).

[251] SJ Williamson. Bulk upper critical field of clean type-ii superconductors: V and nb. *Physical review B*, 2(9):3545, 1970. (Cited on pages 153 and 155).

[252] Patrick Winkel, Kiril Borisov, Lukas Grünhaupt, Dennis Rieger, Martin Spiecker, Francesco Valenti, Alexey V. Ustinov, Wolfgang Wernsdorfer, and Ioan M. Pop. Implementation of a transmon qubit using superconducting granular aluminum. *Physical Review X*, 10(3):031032, August 2020. ISSN 2160-3308. doi:10.1103/PhysRevX.10.031032. (Cited on page 118).

[253] Chengkai Wu, Yong Zhang, Yukun Li, Huali Zhu, Fei Xiao, Bo Yan, and Ruimin Xu. Millimeter-Wave Waveguide-to-Microstrip Inline Transition Using a Wedge-Waveguide Iris. *IEEE Transactions on Microwave Theory and Techniques*, 70(2):1087–1096, February 2022. doi:10.1109/tmtt.2021.3123349. (Cited on page 198).

[254] Ze-Liang Xiang, Sahel Ashhab, J Q You, and Franco Nori. Hybrid quantum circuits: Superconducting circuits interacting with other quantum systems. *Reviews of Modern Physics*, 85(2):623, April 2013. doi:10.1103/RevModPhys.85.623. (Cited on pages 5 and 196).

[255] Mingrui Xu, Xu Han, Wei Fu, Chang-Ling Zou, and Hong X. Tang. Frequency-tunable high-q superconducting resonators via wireless control of nonlinear kinetic inductance. *Applied Physics Letters*, 114(19), May 2019. ISSN 0003-6951. doi:10.1063/1.5098466. (Cited on page 119).

[256] G. Yassin, P. K. Grimes, O. G. King, and C. E. North. Waveguide-to-planar circuit transition for millimetre-wave detectors. *Electronics Letters*, 44(14):866–867, July 2008. (Cited on pages 198, 199, and 201).

[257] Jen-Hao Yeh and Steven M. Anlage. In situ broadband cryogenic calibration for two-port superconducting microwave resonators. *Rev. Sci. Instrum.*, 84(3), March 2013. ISSN 0034-6748. doi:10.1063/1.4797461. (Cited on page 204).

[258] J.G.M. Yip, A.K. Jastrzebski, R.J. Collier, and Daiqing Li. The design of waveguide-to-finline taper transitions at millimetre wave frequencies. In *14th International Conference on Microwaves, Radar and Wireless Communications. MIKON - 2002. Conference Proceedings (IEEE Cat.No.02EX562)*, volume 1, pages 282–285 vol.1, May 2002. doi:10.1109/mikon.2002.1017851. (Cited on pages 198, 200, and 201).

[259] Keiji Yoshida, Mohammad Sajjad Hossain, Takanobu Kisu, Keiji Enpuku, and Kaoru Yamafuji. Modeling of kinetic-inductance coplanar striplin with nbn thin films.

Japanese journal of applied physics, 31(12r):3844, December 1992. ISSN 1347-4065. doi:10.1143/jjap.31.3844. (Cited on pages 50 and 104).

[260] Yang Yu, Siyuan Han, Xi Chu, Shih-I Chu, and Zhen Wang. Coherent temporal oscillations of macroscopic quantum states in a josephson junction. *Science*, 296(5569): 889–892, 2002. (Cited on pages 143, 173, and 180).

[261] Yang Yu, D Nakada, Janice C Lee, Bhuwan Singh, DS Crankshaw, TP Orlando, Karl K Berggren, and William D Oliver. Energy relaxation time between macroscopic quantum levels in a superconducting persistent-current qubit. *Physical review letters*, 92(11): 117904, 2004. (Cited on pages 142 and 173).

[262] Yang Yu, William D Oliver, D Nakada, Janice C Lee, Karl K Berggren, and TP Orlando. Energy relaxation times in a nb persistent current qubit. *IEEE transactions on applied superconductivity*, 15(2):845–848, 2005. (Cited on pages 142 and 173).

[263] B Yurke. Squeezed-state generation using a josephson parametric amplifier. *Josa B*, 4 (10):1551–1557, 1987. (Cited on page 116).

[264] Bernard Yurke and Eyal Buks. Performance of cavity-parametric amplifiers, employing kerr nonlinearites, in the presence of two-photon loss. *Journal of lightwave technology*, 24(12):5054–5066, 2006. (Cited on pages 74, 106, 109, 110, 111, 112, 114, and 135).

[265] C H Zhang, J B Wu, B B Jin, Z M Ji, L Kang, W W Xu, J Chen, M Tonouchi, and P H Wu. Low-loss terahertz metamaterial from superconducting niobium nitride films. *Optics express*, 20(1):42–47, 2012. (Cited on page 97).

[266] Helin Zhang, Srivatsan Chakram, Tanay Roy, Nathan Earnest, Yao Lu, Ziwen Huang, D. K. Weiss, Jens Koch, and David I. Schuster. Universal fast-flux control of a coherent, low-frequency qubit. *Phys. Rev. X*, 11(1):011010, January 2021. ISSN 2160-3308. doi:10.1103/PhysRevX.11.011010. (Cited on page 251).

[267] Jian-Song Zhang, Wei Zeng, and Ai-Xi Chen. Effects of cross-kerr coupling and parametric nonlinearity on normal mode splitting, cooling, and entanglement in optomechanical systems. *Quantum Information Processing*, 16(6):163, 2017. (Cited on page 116).

[268] RUICHEN Zhao, Sungoh Park, Tongyu Zhao, Mustafa Bal, CRH McRae, Junling Long, and DP Pappas. Merged-element transmon. *Physical Review Applied*, 14(6): 064006, 2020. (Cited on pages 84, 143, and 182).

[269] Jonas Zmuidzinas. Superconducting microresonators: Physics and applications. *Annu. Rev. Condens. Matter Phys.*, 3(1):169–214, 2012. (Cited on pages 16, 17, and 120).

APPENDIX A

DERIVATION OF RESONATOR RESPONSES

Direct coupling here refers to a direct connection from the LRC resonator to the transmission lines.

A.1 1-Port Shunt Resonator

Here we will calculate the complex S_{11} parameters for a simple parallel LRC resonator, following the derivation in [184], and ending up with an abstraction of the Rank-1 S matrix in terms of internal quality factor (Unloaded Q) Q_i, external quality factor (coupling Q) Q_e, and resonant frequency ω_0.

A.1.1 Derivation

For an arbitrary load impedance Z_L shunting an input line with impedance Z_0 [184] we can write:

$$S_{11} = \frac{Z_L - Z_0}{Z_L + Z_0} \tag{A.1}$$

In our case, the load impedance is that of a parallel LRC resonator (Figure 3.13). We can write the effective impedance as:

$$Z_L = Z_R||Z_L||Z_C = \frac{1}{i\omega C + \frac{1}{i\omega L} + \frac{1}{R}} \tag{A.2}$$

Expanding about the resonance frequency, we set $\omega = \omega_0 + \Delta\omega$, where the resonant frequency is $\omega_0^2 = 1/LC$. For small $\Delta\omega \ll \omega_0$ we can write:

$$\frac{1}{\omega} = \frac{1}{\omega_0 + \Delta\omega} = \frac{1/\omega_0}{1 + \frac{\Delta\omega}{\omega_0}} \simeq \frac{1}{\omega_0}\left(1 - \frac{\Delta\omega}{\omega_0}\right) \tag{A.3}$$

Using this, we transform the load impedance as follows:

$$Z_L = \frac{1}{\frac{1}{R} + i(\omega_0 + \Delta\omega)C + \frac{1}{i\omega_0 L}(1 - \frac{\Delta\omega}{\omega_0})} \tag{A.4}$$

$$= \frac{1}{\frac{1}{R} + i\omega_0 C + i\Delta\omega C - \frac{i}{\omega_0 L} + \frac{i\Delta\omega}{\omega_0^2 L}} \tag{A.5}$$

$$= \frac{1}{\frac{1}{R} + i\sqrt{\frac{C}{L}} + i\Delta\omega C - i\sqrt{\frac{C}{L}} + \frac{i\Delta\omega LC}{L}} \tag{A.6}$$

$$= \frac{1}{\frac{1}{R} + 2i\Delta\omega C} \tag{A.7}$$

$$= \frac{R}{1 + 2i\Delta\omega RC} \tag{A.8}$$

We now introduce the Quality factors. For a parallel LRC circuit [184]:

$$\text{coupling Q}: \qquad Q_e = \frac{Z_0}{\omega_0 L} = \omega_0 Z_0 C = \frac{\omega_0}{\kappa} \tag{A.9}$$

$$\text{internal Q}: \qquad Q_i = \frac{R}{\omega_0 L} = \omega_0 RC = \frac{\omega_0}{\gamma} \tag{A.10}$$

Using this definition, we can re-write:

$$Z_L = \frac{R}{1 + 2iQ_i\frac{\Delta\omega}{\omega_0}} = \frac{R}{1 + 2iQ_i\Delta} \tag{A.11}$$

Where we have used shorthand $\Delta = \Delta\omega/\omega_0$ for reduced frequency detuning. Substituting into the expression for S_{11} gives:

$$S_{11}(\omega) = \frac{R - Z_0(1 + 2iQ_i\Delta)}{R + Z_0(1 + 2iQ_i\Delta)} \tag{A.12}$$

$$= \frac{\omega_0 LQ_i - Q_e\omega_0 L(1 + 2iQ_i\Delta)}{\omega_0 LQ_i + Q_e\omega_0 L(1 + 2iQ_i\Delta)} \tag{A.13}$$

$$S_{11}(\omega) = \frac{2Q_i}{Q_e + Q_i + 2iQ_eQ_i\Delta} - 1 \tag{A.14}$$

Which is the familiar result. We can get an even neater expression if we transform Q's into κ's:

$$S_{11}(\omega) = \frac{2/\gamma}{\frac{\kappa+\gamma}{\kappa\gamma} + 2i\frac{\omega_0}{\kappa\gamma}\Delta} - 1 \tag{A.15}$$

$$= \frac{\kappa}{\kappa+\gamma}\frac{1}{\frac{1}{2} + i\delta} - 1 \tag{A.16}$$

Where we have introduced reduced detuning $\delta = \Delta\omega/(\kappa+\gamma)$

Summary and Variations

$$S_{11}(\omega) = \frac{2Q_i}{Q_e + Q_i + 2iQ_eQ_i\Delta} - 1 \tag{A.17}$$

$$= \frac{\kappa}{\kappa+\gamma}\frac{1}{\frac{1}{2} + i\delta} - 1 \tag{A.18}$$

$$\text{coupling Q} = \text{external Q}: \qquad Q_e = \frac{Z_0}{\omega_0 L} = \omega_0 Z_0 C = \frac{\omega_0}{\kappa} \tag{A.19}$$

$$\text{unloaded Q} = \text{internal Q}: \qquad Q_i = \frac{R}{\omega_0 L} = \omega_0 RC = \frac{\omega_0}{\gamma} \tag{A.20}$$

$$\text{loaded Q} = \text{total Q}: \qquad \frac{1}{Q_T} = \frac{1}{Q_e} + \frac{1}{Q_i} = \frac{\kappa+\gamma}{\omega_0} \tag{A.21}$$

Asymmetry

The coupling to the circuit is represented by a short section of transmission line, which scatters the signal before reaching the resonator, and introduces a phase $\beta l \approx \phi'$ to the signal that reaches the resonator. The modified reflection coefficient S_11 can be derived

using ABCD matrices [184] (see Appendix A for details), and approximated for small ϕ' as:

$$S_{11} = \frac{Z_R + \frac{-iZ_0Z_c\cos\phi' + Z_c^2\sin\phi'}{-iZ_c\cos\phi' - Z_0\sin\phi'}}{Z_R + \frac{-iZ_0Z_c\cos\phi' + Z_c^2\sin\phi'}{-iZ_c\cos\phi' + Z_0\sin\phi'}} \tag{A.22}$$

$$= \frac{Z_R - Z_0 + i\phi'\left(\frac{Z_0^2}{Z_c} - Z_c\right)}{Z_R + Z_0 + i\phi'\left(\frac{Z_0^2}{Z_c} + Z_c\right)} \tag{A.23}$$

Here we have done a first-order approximation of the numerator and denominator for small ϕ'. Making the replacement

$$Z_0 + i\phi' \rightarrow Z_0^* = |Z_0^*|e^{-i\phi} \tag{A.24}$$

We obtain an equation similar to Equation A.12:

$$S_11 = \frac{R - Z_0^*(1 + 2iQ_i\Delta)}{R + Z_0^*(1 + 2iQ_i\Delta)} \tag{A.25}$$

From this point we can proceed as before, except we define the complex coupling quality factor

$$Q_e^* = |Q_e^*|e^{i\phi} \tag{A.26}$$

A.1.2 2-Port Shunt Resonator

Using the results from the previous derivation, we can extend the results for a 2 port network.

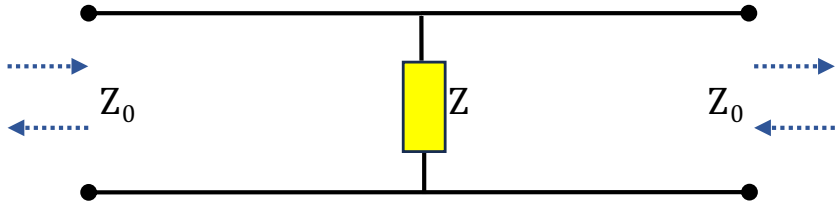

Figure A.1: Impedance network for a shunt impedance in between two transmission lines.

Derivation 1

For a shunting impedance Z_L (Fig. A.1) we have the general 2 port ABCD matrix [pozar]:

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 1/Z_L & 1 \end{pmatrix} \tag{A.27}$$

From which we can directly obtain the relevant S matrix parameters:

$$S_{11} = \frac{A + B/Z_0 - CZ_0 - D}{A + B/Z_0 + CZ_0 + D} \tag{A.28}$$

$$S_{21} = \frac{2}{A + B/Z_0 + CZ_0 + D} \tag{A.29}$$

(Note that while ABCD matrices can be exponentiated to get an ABCD matrix for a more complex system, this is not possible with S matrices!) Substituting values from the ABCD matrix into the expressions above gives:

$$S_{11} = \frac{-Z_0/Z_L}{2 + Z_0/Z_L} \tag{A.30}$$

$$S_{21} = \frac{2}{2 + Z_0/Z_L} \tag{A.31}$$

Substituting Z_L from Eqn. A.11, and setting $Z_0 = 2Q_e\omega_0 L$ (since the LRC sees two loads Z_0 in parallel now), and $R = Q_i\omega_0 L$ we get:

$$S_{11} = -\frac{Z_0(1 + 2iQ_i\Delta)}{2R + Z_0(1 + 2iQ_i\Delta)} \tag{A.32}$$

$$= -\frac{2Q_e\omega_0 L(1 + 2iQ_i\Delta)}{2Q_i\omega_0 L + 2Q_e\omega_0 L(1 + 2iQ_i\Delta)} \tag{A.33}$$

$$= -\frac{Q_e + 2iQ_eQ_i\Delta}{Q_i + Q_e + 2iQ_eQ_i\Delta} \tag{A.34}$$

$$= \frac{\kappa}{\kappa + \gamma}\frac{1}{1 + 2i\delta} - 1 \tag{A.35}$$

$$S_{21} = \frac{2R}{2R + Z_0(1 + 2iQ_i\Delta)} \tag{A.36}$$

$$= \frac{2Q_i\omega_0 L}{2Q_i\omega_0 L + 2Q_e\omega_0 L(1 + 2iQ_i\Delta)} \tag{A.37}$$

$$= \frac{Q_i}{Q_i + Q_e + 2iQ_eQ_i\Delta} \tag{A.38}$$

$$= \frac{\kappa}{\kappa + \gamma}\frac{1}{1 + 2i\delta} \tag{A.39}$$

And, as a sanity check, we can verify that $S_{21} + S_{11} = 1$.

Derivation 2

We can alternatively start with direct expressions for S_{21} and S_{11} given a shunting impedance Z_L:

$$S_{11} = \frac{Z_0||Z_L - Z_0}{Z_0||Z_L + Z_0} = \frac{-Z_0}{2Z_L + Z_0} \tag{A.40}$$

$$S_{21} = \frac{2Z_0||Z_L}{Z_0||Z_L + Z_0} = \frac{Z_L}{2Z_L + Z_0} \tag{A.41}$$

Which reduce to the expressions in Eqn. A.31, from which we can procede in the same way as the previous section.

Summary and Variations

$$S_{11}(\omega) = -\frac{Q_e + 2iQ_eQ_i\Delta}{Q_i + Q_e + 2iQ_eQ_i\Delta} = \frac{\kappa}{\kappa + \gamma}\frac{1}{1 + 2i\delta} - 1 \tag{A.42}$$

$$S_{21}(\omega) = \frac{Q_i}{Q_i + Q_e + 2iQ_eQ_i\Delta} = \frac{\kappa}{\kappa + \gamma}\frac{1}{1 + 2i\delta} \tag{A.43}$$

$$Q_e = \frac{Z_0}{2\omega_0 L} = \omega_0 C Z_0/2 \tag{A.44}$$

A.2 Side Coupled LRC Resonator

More often than not, direct coupling is undesirable due to the lack of control over the coupling Q. It is common to couple the resonator mode to the transmission line mode using a capacitor.

A.2.1 2-Port Transmission

The circuit network in question is identical to Fig. A.1, with the addition of a capacitor C_k in series with the LRC circuit. This gives us an additional term in the impedance from the previous section (Eqn. A.11):

$$Z_L = \frac{1}{i\omega C_k} + \frac{R}{1 + 2iQ_i\Delta} \tag{A.45}$$

where $\Delta = (\omega - \omega_0)/\omega_0$. The addition of the capacitor will redefine our definition of Q_e, but also split the resonant frequency into a resonance and anti-resonance. These frequencies are simply the solution of $\Im(Z_L) = 0$. Rewriting Z_L in a suggestive form:

$$Z_L = \frac{-i}{\omega C_k} + \frac{R(1 - 2iQ_i\Delta)}{1 + 4Q_i^2\Delta^2} \tag{A.46}$$

we can solve:

$$0 = \frac{-1}{\omega C_k} - \frac{2RQ_i\Delta}{1 + 4Q_i^2\Delta^2} \tag{A.47}$$

Solving and substituting our definition for internal loss $R = \omega_0 L Q_i$ (which is unchanged), we get:

$$\omega_0' = \frac{\omega_0\left(Q_i^2(-4 + C_k L\omega_0^2) \pm Q_i\sqrt{-4 + LC_k\omega_0^2(2 + LC_kQ_i^2\omega_0^2)}\right)}{2Q_i^2(-2 + LC_k\omega_0^2)} \tag{A.48}$$

Taking $Q_i \gg 1$, we get:

$$\omega_0' = \frac{\omega_0(-4 + LC_k\omega_0^2 \pm LC_k\omega_0^2)}{-4 + 2LC_k\omega_0^2} \tag{A.49}$$

$$= \begin{cases} \frac{2}{2-LC_k\omega_0^2} \\ \omega_0 \end{cases} \tag{A.50}$$

The frequency of interest is the one shifted lower, which is a resonance, while the unchanged frequency ω_0 is an anti-resonance.

Extreme low coupling limit

Generally we expect to have a very weak coupling (and thus high coupling Q) which should recover a symmetric resonator response. If we make the assumtion that the coupling capacitor is small, and thus the shift is small (ie. $LC_k\omega_0^2 \ll 1$), we can rewrite the frequency shift as

$$\Delta\omega_0 = \omega_0' - \omega_0 = -\frac{LC_k\omega_0^3}{2 - LC_k\omega_0^2} \simeq -\frac{LC_k\omega_0^3}{2} \tag{A.51}$$

We can now expand the load impedance Z_L about frequencies close to the resonance ω_0'. (note that since C_k is small, $\Delta\omega_0$ is also small, but expanding about ω_0 gives a different result). Since at ω_0', $\Im(Z_L) = 0$, then $Z_L(\omega_0') = \Re(Z_L(\omega_0'))$. Thus,

$$Z_L\Big|_{\omega_0'} \simeq \Re(Z_L)\Big|_{\omega_0'} + (\omega - \omega_0')\frac{\partial Z_L}{\partial\omega}\Big|_{\omega_0'} \tag{A.52}$$

$$\Re(Z_L)\Big|_{\omega_0'} = \frac{LQ_i\omega_0}{1 + (2Q_i\frac{\omega-\omega_0}{\omega_0})} \simeq \frac{LQ_i\omega_0}{1 + (2Q_i\frac{LC_k\omega_0^2}{2})^2} \tag{A.53}$$

APPENDIX B

FABRICATION RECIPES

In our field, nobody is ever quite sure exactly which parts of fabrication matter so these types of sections typically get relegated supplemental materials. Since we go into quite a bit of detail about junction fabrication, I'll instead try to organize the leftover processes into recipe steps.

B.1 Substrate Preparation

A low loss dielectric is imperative for developing low-loss superconducting devices. The optimal choices so far (primarily for microwave devices, however we assume this translates to millimeter-wave frequencies as well) are either single-crystal silicon [153] or single-crystal sapphire (Al_2O_3)[176, 188**?**]. Although other crystals (MgO, SiO2) are also sometimes used, silicon and sapphire yield superior properties [1]. In both cases, since the electric fields are typically concentrated within a few microns of the surface the surface dielectric losses play a significant role [188]. To remedy this, we do our best to clean and prepare as pristine of a surface as possible before starting further fabrication steps.

All devices are fabricated on C-plane (0001) Sapphire wafers with a diameter of 50.8 mm. The thickness of the wafer depends on the device frequency: for microwave devices, we use 430 μm-thick sapphire, for K-band 330 μm-thick, and for W-band 100 μm-thick (these are significantly more fragile!). Wafers are cleaned with organic solvents (toluene, acetone, Methanol, isopropanol, and DI water) in an ultrasonic bath to remove organic contamination. As received, the wafer surface has typically undergone cutting and polishing, which leaves behind a surface which while smooth is fairly damaged. To fix this, the wafers are annealed

1. We empirically determined this since our early W band resonators fabricated on quartz had internal quality factors consistently below 10^4

at 1250°C in air[2] for 1.5 Hours: this helps heal the crystal lattice near the surface and as an added benefit burns off contaminants. Prior to film deposition, wafers undergo a second clean with organic solvents (toluene, acetone, methanol, isopropanol, and DI water) in an unltrasonic bath, which helps remove dust and additional organic contaminants that acuumulated during and after the anneal. Following this, the wafers undergo a 2 minute clean in 50°C Nano-Strip™ etch (sulfuric acid and hydrogen peroxide) which helps further remove some damaged and contaminated surface layers of the sapphire. Finished wafers are rinsed with high purity DI water, leaving a clean surface.

B.2 Atomic Layer Deposition

Atomic layer deposition is a conformal process, so I do not recommend it for use on chips, since the deposition method also coats an exponentially thin layer of the partially exposed bottom edge of the substrate (something something gas and plasma diffusion has a characteristic length). For a chip this means the backside of the device now has a thin (and extremely lossy) film of metal! No good. But on a wafer we typically throw away the edge chips anyways! So performing the deposiiton on wafer-scale nearly always gives good results.

B.2.1 Niobium Nitride

Wafers are loaded into an Ultratech Fiji G2 plasma enhanced atomic layer deposition tool for metallization, where they first undergo a 1 hour bake at 300°C under vacuum continuously purged with 5 sccm of argon gas. Chamber walls should match the substrate temperature, at least up to 300 degrees C. The deposition parameters and machine configuration are adapted from Ref. [210]. (*t*-butylimido)tris(diethylamido)-niobium(V) (TBTDEN) is used as the niobium precursor, which is kept at 100°C and delivered by a precursor Boost™ system,

2. We tried annealing wafers in nitrogen and this ends up damaging the surface further as oxygen leaves the crystal

which introduces argon gas into the precursor cylinder to promote material transfer of the low vapor pressure precursor to the wafer [210]. The deposition cycle consists of three 0.5 second pulses of boosted TBTDEN followed by 40 seconds of 300 W plasma consisting of 80 sccm hydrogen and 5 sccm nitrogen. A flow of 5 sccm of nitrogen and 10 sccm of argon is maintained throughout the deposition process. After deposition the wafer is allowed to passively cool to 250°C under vacuum (in the chamber).

B.2.2 Titanium Nitride

See Ref. [204] for more details and microwave TiN characterization. Wafers are also loaded into an Ultratech Fiji G2 plasma enhanced atomic layer deposition system, where they are first baked at 270°C under vacuum continuously purged with 5 sccm of argon gas. Tetrakis(diethylamido)titanium(IV) (TDMAT) is used as the titanium precursor, and is delivered in 0.5 second pulses along with argon, followed by a 300 W nitrogen-argon plasma. The wafer is maintained at 270°C during and after the deposition until unloading.

B.3 Lithography

Following deposition, wafers (or chips) undergo a brief dehydration bake at > 90°C in atmosphere for a minute or so before spinning resist. For optical lithography, to avoid defocusing from wafer deformation, wafers are mounted to a Silicon handle wafer with AZ MiR 703 (or thicker AZ MiR 1518) photoresist cured at 115°C. Wafers are then coated with ~ 1 μm of positive I-line photoresist AZ MiR 703, and raster exposed with a 375 nm laser Heidleberg MLA150 Direct Writer. Defocus and dose are optimized for best resolution, typically relying on lower dose values to ensure small features are not blurred away. If necessary, the wafer is realigned until rotation is less than 5 mRad to ensure the best resolution grid.

For electron Beam lithography, wafers are first coated with ~ 200 nm of negative resist ARN 7520, followed by 40 nm of conductive resist 'Elektra' AR PC 5090. This stack leads to

significant crosslinking and thick residues after plasma etching however. APR 6200(a ZEP equivalent) is better suited for etching, and useful for resolving sub 100 nm scale features (such as the nanowires used in the fractal work). This resist is hydrophobic however, so we instead use a 5–10 nm conductive layer of thermally evaporated gold.

The spun wafers are exposed with a Raith EBPG5000 Plus E-Beam Writer. The conductive coating is removed by either a 60 second DI water quench (for elektra) or with an acqueous iodine etch in the case of gold. Both optical and negative E-Beam resists are baked at 110°C to further harden the resist, and then developed for 60 seconds in AZ MIF 300, followed by a 60 second quench in DI water. ARP 6200 is developed with amyl acetate, followed by a quench in IPA.

B.4 Reactive Ion Etching

Metal films are etched in a Plasma-Therm inductively coupled plasma reactive ion etcher. Etch chemistry, substrate etch depth and etch time have been shown to affect planar resonator quality factors [132], in particular due to the formation of cross-linked polymers at the metal-resist interface after the bulk metal is etched away (see Chapter 7 for details). For this reason we scale sample etch times to metal thickness, with a fixed over-etch time of between 10 and 30 seconds to ensure complete metal removal. As a general process for etching niobium or tantalum-containing films, we use a low-pressure fluorine based ICP etch chemistry with a plasma consisting of 15 sccm SF_6, 40 sccm CHF_3, and 10 sccm Ar. ICP and bias powers are kept at 100 W, and the substrate is maintained at a temperature of 20°C. However as discussed in detail in Chapter 7 we find that different etch chemistries yield better results.

B.5 Resist Removal and Dicing

Following etching, the resist is stripped in 80°C Remover PG (N-Methyl-2-Pyrrolidone), which contains a small percentage of surfactant: this assists in mechanically removing undissolved resist chunks. This step also releases the wafer from the Silicon carrier wafer (though freuqently requires additional coaxing with extensive sonication).

The wafers are then cleaned with organic solvents (acetone, isopropanol, and sometimes DI water), coated with a ~ 2 μm protective layer of photoresist, and diced into chips with a diamond blade. This process can be fairly destructive for thin wafers and small chips, so we use slow blade travel speeds of 0.5 mm/sThese are stripped of protective resist with 80°C Remover PG, cleaned with organic solvents (acetone, isopropanol, and DI water), dried either on an unpolished Sapphire carrier wafer in atmosphere at 80°C, or directly with nitrogen (though this can easily blow away the chips!). Finally the finished chips are mounted with indium (in various fashions) described in the text.

APPENDIX C

DISENTANGLING HETERODYNE QUBIT IMAGES

The heterodyne method of generating the millimeter-wave qubit pulse described in Chapter 10 generates a number of signals, making it difficult to determine which pulse image is exciting the real qubit transition. We show a particularly complex example in Figure C.1 to C.3. By sweeping qubit LO frequency and qubit IF frequency and organizing the results as a function of lower and upper sideband frequencies as well as the LO frequency allows us to determine which feature is caused by which signal. Features that remain constant from a particular sweep suggest the feature is excited by that paticular image.

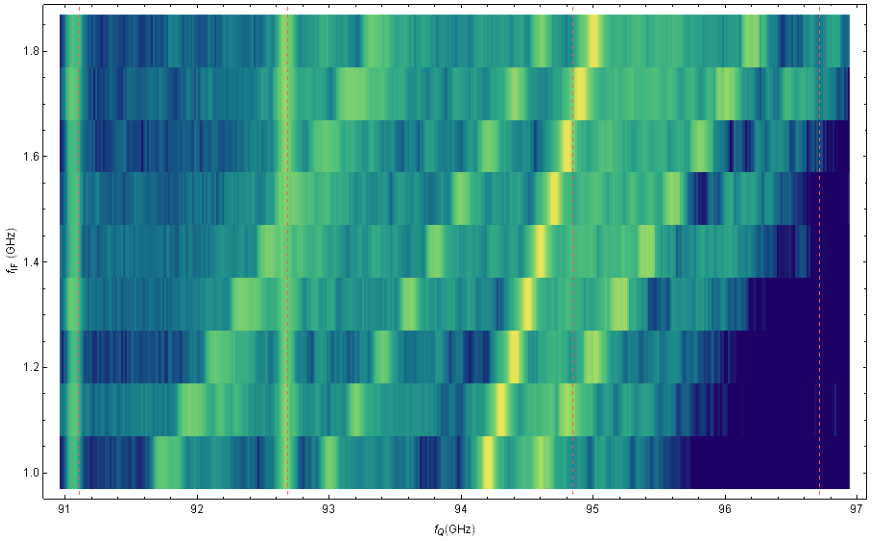

Figure C.1: Continuous-wave spectroscopy using the millimeter-wave heterodyne qubit pulse measurement setup, plotted with respect to the qubit upper sideband frequency for increasing qubit IF pulse frequencies.

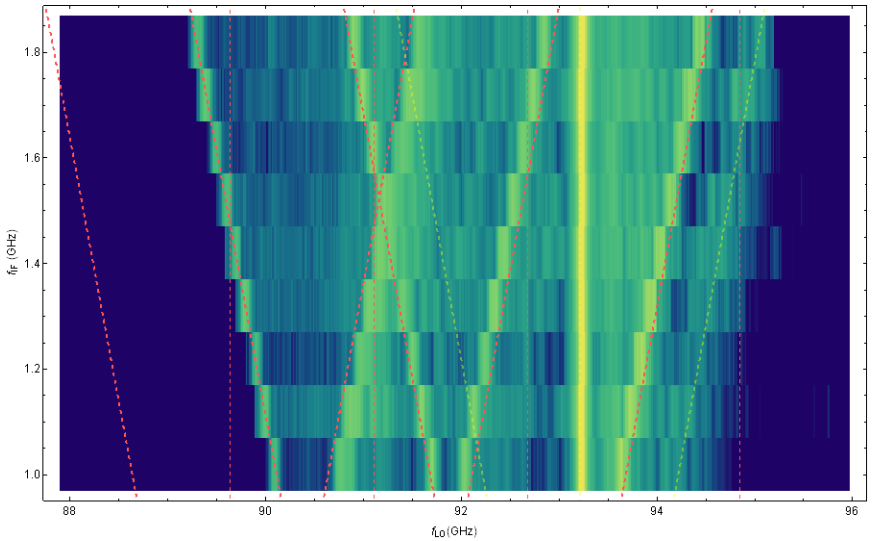

Figure C.2: Continuous-wave spectroscopy using the millimeter-wave heterodyne qubit pulse measurement setup, plotted with respect to the qubit LO frequency.

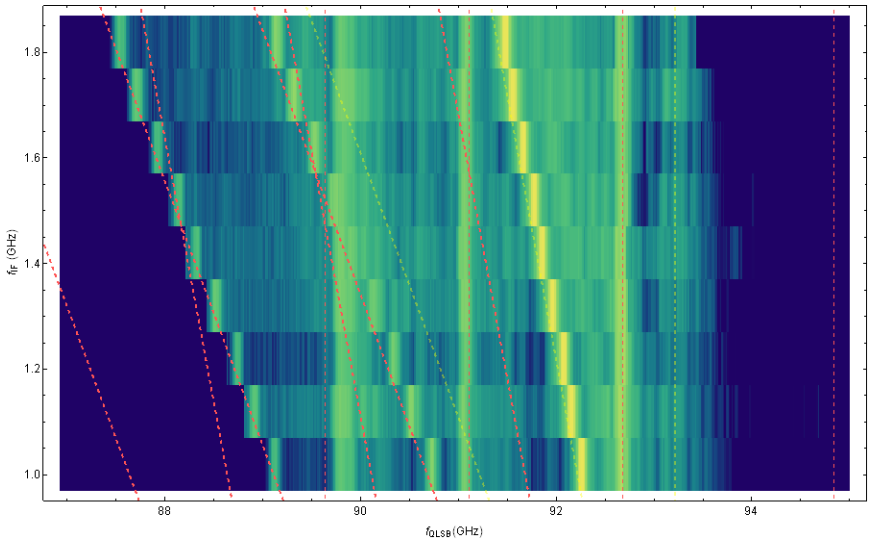

Figure C.3: Continuous-wave spectroscopy using the millimeter-wave heterodyne qubit pulse measurement setup, plotted with respect to the qubit lower sideband frequency for increasing qubit IF pulse frequencies.