THE UNIVERSITY OF CHICAGO

MULTIMODAL AND ULTRA HIGH-Q SUPERCONDUCTING NIOBIUM CAVITIES FOR CIRCUIT QUANTUM ELECTRODYNAMICS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE PRITZKER SCHOOL OF MOLECULAR ENGINEERING
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
ANDREW ENRIQUE ORIANI

CHICAGO, ILLINOIS
DECEMBER 2022

To my mother and father, my brother Alex, and to all of the amazing educators who have helped me along the way.

If what shone afar so grand

Turn to nothing in thy hand

On again! The virtue lies

In the struggle, not the prize.

R.M. Milnes

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

Something that no one tells you is how much of a collaborative effort a PhD can be. With that said, my greatest anxiety is *not* properly acknowledging everyone who helped me over my years. To begin I would like to acknowledge the multimode and 3D cavity teams for always being helpful. These include Ankur Agrawal, Kevin He, Akash Dixit, and Riju Banerjee. I would like to also acknowledge Ravi Naik, who dragged me into the fold with the 3D flute and multimode cavities. On the theory side the input from Liang Jiang and his group, while not directly applicable to this thesis, was a driving force in many of the experimental questions asked. I would also like to recognize Clai Owens, who helped early on with many of the machining problems, and was a close collaborator (albeit with *very* different approach) on many of the machining and manufacturing issues involving these cavities. Alex "Sasha" Anferov, who always had great insights and careful analysis of the problem, and acted as a sanity check for much of this work, and was at times a close cryogenics collaborator as well. Gerwin Koolstra taught me many basic ideas in cryogenic RF measurement, and was always so generous with his time. My office mates Nate Earnest, Brendan Saxberg, and Helin Zhang were always a productive source of distraction and a good laugh. My (second) deskmate Andrei Vrajitoarea was always a great source of wisdom and insight. Fang Zhao, who brought much needed material science chops to our lab was *indispensable* to this thesis, and a pleasure to work with. Gabrielle Roberts and Meg Panetta, who besides being great scientists also did a lot to add to the lab culture and make the lab a more pleasant place to be. My former understudy, who has since become an accomplished scientist himself, Larry Chen, who helped on a lot of things early on. I would also like to recognize Justin Jureller and Luigi Mazzenga who were instrumental in keeping MRSEC and the machine shop in good working order, and always being so helpful. Also Brenda Thomas, who kept the lab running, and was always a pleasure to be around. I would also like to acknowledge Rhys Povey, my roommate, who was always a source of good conversations and a good time. My

partner, Rachel, who was always patient and encouraging during this process. Second to last, I would like to acknowledge Srivatsan "Vatsan" Chakram, my first deskmate, and in many ways a co-advisor, who was always so supportive, and was an amazing teacher, both in science and in life. I will always cherish our late night Sunday Lao Sze Chaun dinners. Finally, I would like to thank my advisor, Dave, who always challenged me in intellectually stimulating ways, and taught me a lot of important lessons in how to conduct science. His endless encouragement, kind words, precociousness, and unique ability to motivate people to reach there full potential, are all things I will carry with me. I have no doubt that our paths will cross again.

ABSTRACT

The high coherence afforded by 3D superconducting cavities, combined with strong engineered interactions via circuit quantum electrodynamics (cQED), has made 3D cavities a leading platform for studying quantum phenomena. In this thesis, we will develop a 3D architecture that leverages these techniques to allow for the addressing of quantum states encoded onto individual modes of a multimodal cavity. We will show that by engineering the mode dispersion and electric field, we can achieve qubit-cavity cooperativities in excess of 1×10^9, while maintaining coherence times approaching 2ms over 9 modes. In order to scale the number of accessible modes, we will introduce niobium cavities as a platform for reaching higher coherence times. First, we will discuss developing a repeatable technique for producing single-mode 3D niobium cavities with loaded single-photon coherence times in excess of 15ms and internal quality factors of greater than 1.5×10^9. Next we will examine the effects that surface processing and chemistry have on the cavity performance. Finally, we will outline the intricacies of implementing cQED in such high-coherence devices, before demonstrating a single-mode quantum memory with energy relaxation times approaching 10ms. By extending cQED to 3D multimodal platforms, and developing techniques that produce coherence times in excess of 10s of milliseconds, we hope to enable new inquiry into the application of 3D cQED for exploring quantum information and quantum optics phenomena.

CHAPTER 1

INTRODUCTION

1.1 Quantum memories with 3D cavities

The need for more hardware efficient and scalable quantum systems is necessary to reach the long-term milestones and promises that quantum technologies hold. The placement of superconducting quantum systems, with their large range of tunability, interaction strength, and high intrinsic anharmonicities, make them a good candidate for creating large scale quantum computers. In this section we will layout the historical and theoretical groundwork for quantum memories, and explain the role that 3D cavities may play in building larger quantum systems , setting the stage for the proceeding chapters of this thesis.

1.1.1 Background to quantum memories

Any system that a quantum state can be encoded onto can act as a quantum memory. By analogy, any classical transistor, magnetic domain, or even a complex polymer, like DNA, can store information, however the choice of storage media is determined by the needs of the specific application. Beyond just memory, modern integrated processors employ a huge variety of transistor design and architecture to better optimize the performance of the system as a whole. By comparison, quantum computation, which is still very much in its infancy, has no delineation; Current state-of-the art quantum processors, like Google's Sycamore QPU, make no distinction between a qubit used for logical operations, and ones sitting idle, "storing" the state between operational cycles [12]. Right now this is not much of a concern, but as the system size grows, and the number of operations increases, local quantum registers, cache, and eventually "non-volatile" storage will be required, with a hierarchy in operational speed and coherence times existing between these components.

The building blocks of these superconducting processes, the transmon qubit, which will

be discussed in more detail in the next section, coupled to resonators via circuit quantum electrodynamics (cQED), currently represents the basis of the logical hardware of these systems [104, 183, 164]. Historically the transmon represented a huge leap forward in superconducting quantum systems, with inherent insensitivity to charge noise that plagued the performance of earlier charge qubit designs. By also coupling these qubits to a resonant cavity there was also an enhancement in the coupling, leading to a higher fidelity and operational speed. Combined with dispersive readout techniques [151, 184] and continuous improvements in fabrication and design, modern state-of-the-art transmon qubits can now reach coherence times in excess of $300\mu s$ [140, 188]. Other than the transmon there has been the creation of qubits that are intrinsically protected from charge noise, known as the fluxonium qubit, which has demonstrated coherence times in excess of 1ms [65, 169].

An alternative route to creating long-lived quantum memories, however, may be to skip the superconducting circuit altogether, and instead off-load the memory to systems that exhibit intrinsic long-term coherences. NV centers and neutral atom systems have both demonstrated coherence times in excess of 5 seconds [20, 4], while trapped ions have shown lifetimes on the order of minutes [191]. These electronic and nuclear transitions, however, occur at much high frequencies than the aforementioned superconducting systems, making interoperability of the two difficult without the advent of quantum limited transduction. In addition to this, the coupling rates of these systems necessitate much longer gate operations and in some cases readout times, making fast encoding and retrieval challenging. Nonetheless, these systems hold huge promise for both quantum memory and computation, and will nevertheless be integral parts of a future quantum ecosystem.

Ideally we would like to find a system that has interoperability with existing superconducting and cQED systems, with longer coherence times and similar coupling rates. One solution is to utilize the resonant cavity instead of the qubit for the storage of the quantum state. We can encode our logical state onto the excitations of our resonant cavity, which is

described exactly by an harmonic oscillator. Because of the intrinsic linearity of a harmonic system, we need a way to effectively initialize and control the number-state of our oscillator. cQED provides the answer for how, through the introduction of a conditional non-linearity, we can enact exact control over the logical encoding and readout of our resonator state [149]. By using resonators that have very dilute electromagnetic fields, i.e. the field energy *mostly* lives in the lossless vacuum of the resonant cavity, we can have cavity photon lifetimes that far exceed most, if not all, superconducting qubits.

The first demonstration of coupling a transmon qubit to a 3D superconducting cavity [135] laid the groundwork for future advancements in quantum memories, including the creation of a quantum memory with $> 1ms$ coherence times [148], universal control of the oscillator state [88, 87], and later error correction [152, 35]. These techniques have since been expanded to two-cavity systems, with the demonstration of 2-mode exchange and entanglement [75], where they have even been used to demonstrate molecular dynamic simulations [189].

These techniques showed that it was possible to efficiently store and control quantum states encoded onto a cavity's resonant mode, but to make a memory, made up of multiple quantum state, would require the advent of a multimodal resonant system, where each mode could be addressed by one or more non-linear elements. The first demonstration of a multimodal memory, which utilized a single transmon qubit coupled to a chain of capacitively coupled coplanar waveguide resonators, showed that it was possible to address an arbitrary number of modes, and conduct universal gate operations on and between any number of those modes [124]. This result also showed some of the advantages of such a multimodal system for use as both a memory and as processor, namely that only a single qubit, and its requisite control hardware, were required to interact with the entire register of harmonic oscillators, thereby creating a scalable and hardware efficient system.

The 2D multimodal quantum memory was an important first step in the creation of

scalable quantum memory, it also showed the importance that mode coherence times had on the ability to create large arbitrary entangled states between modes. In the case of the 2D design, the main contributor to errors was the limited mode coherences of $1 - 6\mu s$, which was lower than the qubit relaxation time of $\sim 10\mu s$, leading to a system where accumulated errors during the idling of the mode were the dominant source of infidelity. With this in mind the next natural step was to adapt the ideas of the 2D multimode design to 3D cavities, where one could utilize the aforementioned high mode coherence times and existing control techniques.

In the following few sections we will discuss some of the technical background and operating principles of quantum circuits and circuit quantum electrodynamics as the basic building blocks of a quantum memory, before setting off to solve the problem at the heart of this thesis: how to build the best quantum memory to date.

1.1.2 Cavity quantum electrodynamics

In the previous section we spent some time going over the short (but rich) history of superconducting quantum memories, but we have not explained the physics that underpins their operation. Circuit quantum electrodynamics (cQED) is simply an extension of cavity quantum electrodynamics, which used a cavity to enhance the light-matter interaction between atoms and photons. This enhancement comes about since any spontaneous emission from the atom that has a frequency outside of a cavities resonant frequency will be suppressed, while resonant interactions, in a limit where the cavity and atom decay rate is small when compared to the rate that the atom can emit and readsorb a photon with the cavity, will lead to a coherent interaction between the two. This coherent exchange from a cavity to an atom excitation, known as a *Rabi* oscillation, is just one result of this interaction and of the rich physics possible in these systems. This section will discuss the basic principles of this atom-cavity interaction, but only to the point that intuition necessary to understand later

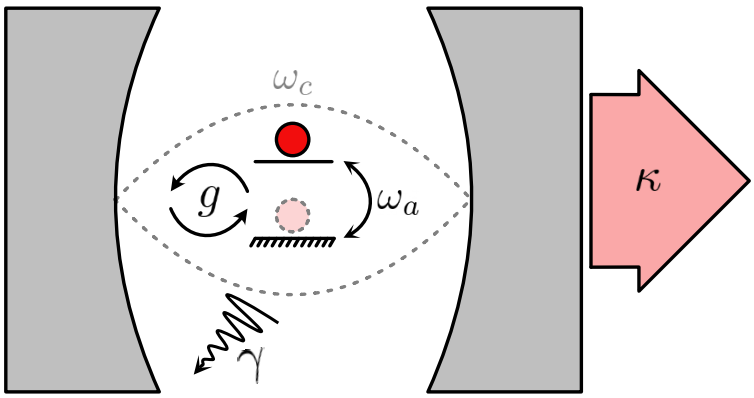

Figure 1.1: **Cavity QED interaction**: A two-level system is coupled to the electromagnetic field of a cavity via dipole interaction. A coherent interaction rate g, describes the rate of exchange of energy between a cavity photon and an atomic excitation. The total loss rate of the cavity, κ, represents incoherent loss in the form of cavity dissipation and leakage. In addition to this, atomic decay into parasitic modes of the cavity is described by γ. The scale of γ and κ versus the coupling rate determine, in addition to the cavity-atom frequency detuning, which coupling regime the system operates in.

sections can be gleaned. For more in-depth discussions on cavity QED and quantum optics see Ref. [185, 72, 164, 27, 86].

The aforementioned phenomena can be described by a two-level system that interacts with a cavity via a dipole coupling to the electromagnetic field. We can compactly write this interaction in the form of the Jaynes-Cummings Hamiltonian, written below using RWA and in the absence of external driving or non-unitary dissipation terms:

$$H_{\mathrm{JC}} = \hbar\omega_r(a^\dagger a + 1/2) + \hbar\frac{\omega_a}{2}\sigma_z + \hbar g(a^\dagger\sigma^- + a\sigma^+) \tag{1.1}$$

The first term in Eq. 1.1 describes the cavity as a quantum harmonic oscillator (QHO) with a resonant frequency ω_r, the second term describes our atom as a spin-1/2 system with an atomic transition energy of $\hbar\omega_a$, and the last term represents the dipole coupling between the two, with g being the rate of exchange between the atom and cavity, where $a^\dagger\sigma^-$ represents an emission of a photon from the atom to the cavity, and $a\sigma^+$ describes the absorption of the photon by the atom.

The above describes the coherent interaction of the atom and cavity. In addition to this there are competing incoherent processes, such as decay, due to cavity leakage or dissipation by the cavity that lead to a decay rate κ. Besides cavity decay processes, the atom can also decay at a rate γ, either radiatively into other parasitic modes of the cavity, or non-radiatively. In actual cavity QED systems, where the atomic system is often moved through the cavity, there is an additional transit timescale that must be considered, however our systems fix the atom to a point within the cavity.

As we will see, in the case of cQED systems, the interaction rate is typically much greater than the decay rate. The implication of this is important, in that the more time a cavity photon can interact with the atomic state before decay, or vice-versa, the more sensitive the atom (or cavity) is to the state of the other. Two parameters that encapsulate this are the critical atom and photon numbers:

$$N_0 = \frac{2\gamma\kappa}{g^2} \tag{1.2a}$$

$$m_0 = \frac{\gamma^2}{2g^2} \tag{1.2b}$$

Where N_0 is the critical atom number, or how many atoms are required to split the cavity line width, or in other words, how sensitive the cavity is to an atom excitation. m_0 is the critical photon number, or how many photons are needed to measure the atomic state. It should be noted that, in circuit QED, the critical atom number, N_0, is sometimes expressed as the cooperativity factor $C = N_0^{-1}$, and is often used as a metric to describe how many times the cavity can probe the qubit state before decay. This reflects the need of quantum computation to do many operations on a state during the course of a logical operation.

The scale of atom-cavity detuning, $\Delta = \omega_a - \omega_r$ versus the coupling rate determines the

way the physics of the system behaves. Previously we discussed the case where the qubit is resonant with the cavity, or $\Delta \to 0$. In this limit, the resonator photon number states, $|n\rangle$ and the atomic ground and excited states, $|g\rangle$ and $|e\rangle$, are no longer eigenstates of the hamiltonian, but instead the cavity and atom eigenstates include a contribution from the coupling term, leading to a hybridization of energy levels, with the new energies split by $2g\sqrt{n}$. This manifests in a splitting of the cavity resonance, with the linewidth of the two peaks set by the cavity and atom decay rates. As the linewidth approaches that of the splitting, or when the decay rate approaches g, the states become indistinguishable.

The above strong-coupling limit can be used to create and study interesting quantum optics phenomena, it means that the individual characteristics of the cavity and atom are indistiguishable. By detuning the cavity and atom frequencies, to a point where $\Delta \gg g$, the cavity photon only indirectly interacts with the atom via a dispersive coupling. The result is that the cavity frequency experiences a small state-dependent shift based on the atom state, and vice-versa. We can apply apply the unitary transformation $U = \exp\Big[\frac{g}{\Delta}(a^\dagger\sigma^- + a\sigma^+)\Big]$ to Eq. 1.1, and expanding in orders of g/Δ before truncating to second-order in g to get the *dispersive hamiltonian*:

$$UH_{\rm JC}U^\dagger \approx \hbar\left(\omega_r + \frac{g^2}{\Delta}\sigma_z\right)(a^\dagger a + 1/2) + \hbar\omega_a\sigma_z/2 \tag{1.3}$$

The above dispersive hamitonian has some nice features. For one, because no photon is directly adsorbed, both the atom state and photon number are conserved. Instead, we see that the cavity *shifts* by g^2/Δ dependent on the state of the atom, as illustrated above in Eq. 1.3. This shift, which we will write as χ, is the *dispersive shift*, and the primary way that readout is conducted. Because of the conservation of the atomic state, quantum non-demolition measurements (QND) can be performed.

Another nice feature of the dispersive limit is the symmetry of the interaction between the atom and the cavity. We can regroup the terms in Eq. 1.3, collecting state dependent

terms to arrive at the form:

$$H_{JC} \approx \hbar\omega_r(a^\dagger a + 1/2) + \frac{\hbar}{2}\left(\omega_a + \frac{2g^2}{\Delta}a^\dagger a + \frac{g^2}{\Delta}\right)\sigma_z \tag{1.4}$$

In the above form, the atomic transition frequency ω_a experiences a shift that is made up of two terms. The first is the standard dispersive shift $\chi \approx g^2/\Delta$, the second however is $2\chi a^\dagger a = 2n\chi$, where n is the the photon number state of the cavity. This second term is the *number-dependent* "stark" shift. These corresponding atom-cavity shifts, which are a result of the Heisenberg-uncertainty principle, and subsequently a result of measurement backaction, give us a powerful way of measuring both cavity *and* atom states while preserving the quantum state of either. It should be noted that Eqs. 1.3 and 1.4 both only consider the lowest laying energy states of the atom, g and e, which is a good approximation. In circuit based systems discussed in the next section, where the anharmonicity of our qubit is on the order of g, the above approximation for χ breaks down, and contributions from higher qubit energy levels must be considered.

The limits on how long one can probe the system is largely set by the rate of loss in the form of incoherent processes. Qubits that are tuned to be far off-resonance from the cavity have inherent protection, since photons emitted outside of the bandwidth of the resonator are largely suppressed. To readout a qubit state, however, it is natural to overcouple the qubit to a resonator which has a decay rate larger than the qubit's decay time ($\kappa \gg \gamma$). We can express this rate using Fermi's golden rule [164]:

$$\gamma_\kappa = \kappa |\langle -, n-1| a |+, n\rangle|^2 \tag{1.5}$$

Where the eigenstates $|-, n\rangle$ and $|+, n\rangle$ in terms of atom and cavity excitations are:

$$|-, n\rangle = |g, n\rangle - \frac{g\sqrt{n}}{\Delta}|e, n-1\rangle \tag{1.6a}$$

$$|+,n\rangle = \frac{g\sqrt{n}}{\Delta}|g,n\rangle + |e,n-1\rangle \tag{1.6b}$$

Substituting Eqs. 1.6 (a) and (b) into Eq. 1.5 we get an effective decay rate of the qubit *through* the lossy resonator to be:

$$\gamma_\kappa \approx \left(\frac{g}{\Delta}\right)^2 \kappa \tag{1.7}$$

This rate, which represents the enhanced emission of the qubit through the resonator is known as the *Purcell* loss. In the design of cavity-qubit systems, this value represents a limit to the effective performance of our qubit. As we can see from the symmetry of the dispersive eigenstates in Eqs. 1.6 (a) and (b), the cavity can similarly experience an enhanced emission by being coupled to a lossy qubit. This rate, sometimes denoted as a *reverse purcell* effect, is similarly given as:

$$\kappa_\gamma \approx \left(\frac{g}{\Delta}\right)^2 \gamma \tag{1.8}$$

As written, Eqs. 1.7 and 6.2.2 are photon-number independent. In reality this is only a first-order approximation. The above values however set upper bounds in the qubit or cavity coherence times for a given set of design parameters. In the case of the effective lifetime of a qubit coupled to a lossy readout, the effective lifetime can be given as:

$$T_{1,\text{eff}} = \left(\frac{1}{T_{1,q}} + \frac{1}{T_{1,\kappa}}\right)^{-1} \tag{1.9}$$

Where $T_{1,q} = \gamma^{-1}$ is the intrinsic coherence time of the qubit, and $T_{1,\kappa} = \gamma_\kappa^{-1}$ is the purcell limited lifetime. For a real-world system, for $\gamma/\gamma_\kappa = 10$, the effective coherence time of the qubit will be 10% *lower* than the intrinsic qubit lifetime, meaning that the ratio of γ and γ_κ should be carefully considered when designing a cavity-qubit system where qubit

coherence is important. Inversely, we will see how this manifests in ultra-high coherence cavity systems, where the qubit lifetime, and dispersive coupling rate, set the main limit in the cavity's coherence time.

The above dynamics have all been defined in the rotating-frame, which have effectively removed any time-dependence from our derivations up to this point. While this treatment is pedagogical in nature, in practice time-dependent drives are necessary to control and measure the state of the system. In later chapters we will show what the dynamics of the above system looks like under the influence of simple time-dependent drives, however the time-dependence of these systems, and the use of time-varying drives to generate non-trivial quantum states, is largely outside of the scope of this thesis. For that discussion about the platforms discussed in this dissertation, Kevin He's thesis should be compulsory reading.

1.1.3 Superconducting quantum circuits

This section will provide some working knowledge in superconducting quantum circuits as they will pertain to the 3D quantum memories discussed in later sections. This section should not be seen as an exhaustive explanation of these topics, but instead provide some intuition and vocabulary to understand the methodologies and logic of later sections. For more thorough descriptions of the following please reference the foundational texts given in Refs. [25, 104, 164, 29, 53, 78].

In the previous section we outlined the basic tenants of cavity quantum electrodynamics, assuming a generic two-level system with an atomic transition frequency ω_a. In our case the "atom", or as we will call it from now on, the *qubit*, is actually a discrete circuit. We can begin with the simplest harmonic circuit we can think of: an LCR circuit. In Ch. 2 we will go in far more depth about the properties of LCR circuits as the model for our high-coherence cavity systems, however, in this discussion, we will simplify matters and only look at the qubit as an ideal LC circuit. This basic abstraction is shown in Fig. 1.2. We can do this

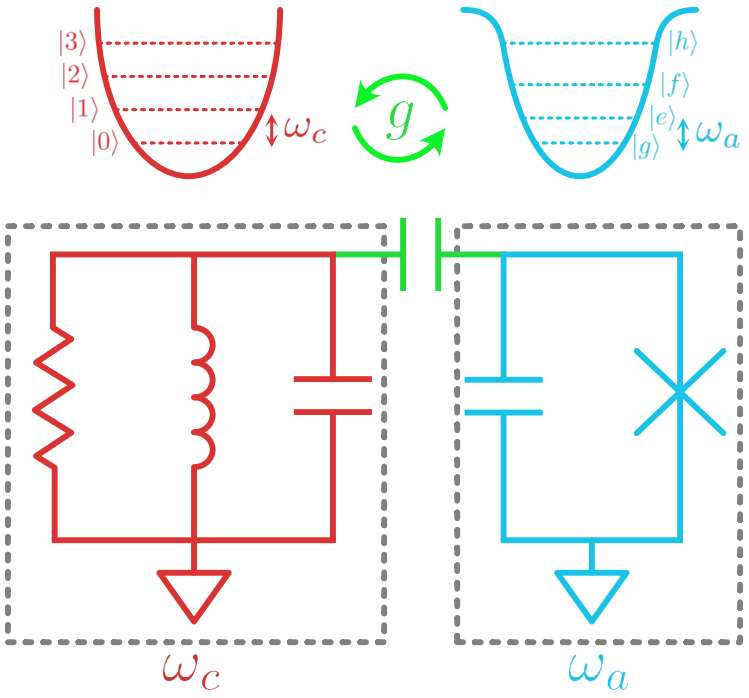

Figure 1.2: **Ciruit elements of cQED system** Above is the canonical LCR oscillator that will be the basis of both the cavity (see Ch. 2), or, as we will see in this section, the qubit, when we introduce a non-linear element to the above system.

because, in practice, our circuit is a superconductor, thereby making the resistive component very small (but as we will see, not zero). In a later chapter we will discuss the properties of superconductors at microwave frequencies, but for now we will assume that $\omega_a \ll \Delta_g$ and that $T \ll T_c$, where Δ_g and T_c are the superconducting gap energy and superconducting transition temperature respectively. In this regime the current flowing through the inductor experiences no loss.

The total energy of the system is stored in both the capacitor, as built up charge, and as current passing through an inductor. We can write the Lagrangian that describes this circuit as:

$$\mathcal{L} = \frac{1}{2}LI^2 - \frac{1}{2}\frac{Q^2}{C} \tag{1.10}$$

We can note that the current is $I = dQ/dt = \dot{Q}$. From this we can solve for the Euler-

Lagrange equations of motion, which give us the expected harmonic frequency of the circuit, $\omega_h = (LC)^{-1/2}$, and the momentum conjugate $\Phi = LI$. We can now write the Hamiltonian in terms of conjugate variables Φ and Q:

$$H = \frac{\Phi^2}{2L} + \frac{Q^2}{2C} \tag{1.11}$$

In the above the the value ϕ is described as the *flux* through the inductor, while Q is the charge stored on the capacitor. Here Q is treated as a single degree of freedom, which has some mass L and experiences a spring constant C. We can make our system quantum by giving our conjugate variables nice hats and imposing the canonical commutation relation:

$$[\hat{\Phi}, \hat{Q}] = -i\hbar \tag{1.12}$$

From this we can write the quantum operators $\hat{\phi}$ and $\hat{Q}$ in term of annihilation and creation operators a and $a^\dagger$:

$$\hat{\Phi} = \sqrt{\frac{\hbar Z_0}{2}}(a + a^\dagger) \tag{1.13a}$$

$$\hat{Q} = i\frac{\hbar}{2Z_0}(a^\dagger - a) \tag{1.13b}$$

Noting that $[a, a^\dagger] = 1$. He have defined a new variable, $Z_0 = \sqrt{L/C}$, which is the characteristic impedance of our LC oscillator. We can rewrite Eq. 1.11 in terms of a and $a^\dagger$ to get the canonical form of the quantum harmonic oscillator Hamiltonian:

$$H = \hbar\omega_h\left(a^\dagger a + \frac{1}{2}\right) \tag{1.14}$$

Where we note that $a^\dagger a = \hat{n}$, which is the number operator of our oscillator. At this point we have not done anything that interesting, but just showed that an LC oscillator can

be quantized and subsequently functions as a quantum harmonic oscillator. At the end of the day, this circuit would make a pretty lousy qubit, since there is no way to confine the otherwise infinite Hilbert space of the above system to just two levels to act as our logical basis. To do that, we need to make the frequency of our LC circuit *dependent* on the number of excitation, i.e. nonlinear.

One way to do that is start with a circuit element that does not exhibit linear $I - V$ characteristics. One natural choice is the Josephson junction, which is simply a tunnel barrier separating two superconducting islands. The cooper pairs[1] on each island can be described by a superconducting order parameter which leads to a relative time dependent phase, $\varphi(t)$, between the wave-functions of the cooper pairs on each island. The result of this boundary condition problem is that the voltage and current of these electrons become dependent of this phase difference, leading to the Josephson relations:

$$V(t) = \frac{\Phi_0}{2\pi}\frac{\partial\varphi(t)}{\partial dt} \tag{1.15a}$$

$$I(t) = I_c \sin\varphi(t) \tag{1.15b}$$

Where $\Phi_0 = h/(2e)$ is equivalent to the flux-quantum, and I_c, or critical current, is the current in which the tunneling cooper pairs break into conduction electrons. We can plug Eq. 1.15 (b) into (a) and get:

$$V(t) = \frac{\Phi_0}{2\pi I_c, \cos\varphi(t)}\frac{\partial I(t)}{\partial t} = L_j(\varphi(t))\dot{I}(t) \tag{1.16}$$

We see that the above time dependent relationship between I and V is just Faraday's law. Here we have lumped the prefactor of our current time-derivative into a *Josephson*

1. To see discussions on this phenomena see Sec. 2.2.4 and Ref. [175]

inductance, L_j. As we can see, the inductance is dependent on the Josephson phase. To see how this may effect our LC oscillator, we can determine the stored energy of our Josephson inductor by integrating the product of the voltage and current over time:

$$U_j(t) = \int_0^t I(t)V(t)dt = -E_j \cos\varphi(t) \tag{1.17}$$

Where $E_j = \Phi_0 I_c/2\pi$ is the Josephson energy.

With Eq. 1.17 we can substitute it back into Eq. 1.11. First we will group several parameters, first by defining $Q = ne$ where n is the charge number. We can then define a charging energy $E_c = e^2/2C$.

$$H = E_c n^2 - E_j \cos\varphi(t) \tag{1.18}$$

So, what is the takeaway from all of this, and what are the implications? For small phase fluctuations across the junction, Eq. 1.18 reduces back into the LC oscillator Hamiltonian in Eq. 1.11 since $\cos\varphi(t) \sim 1$. But at large $\varphi(t)$, the non-linearity of the cosine term gives the harmonic oscillator a 4^{th}-order correction. This anharmonicity breaks the degeneracy *between* transition energies of our harmonic oscillator, allowing us to selectively confine our Hilbert space to the lowest two energy levels.

We can write Eq. 1.18 in the charge basis and phases basis as follows:

$$H = 4E_c(\hat{n} - n_g/2)^2 + \frac{E_j}{2}\sum_n (|n\rangle\langle n+1| + |n+1\rangle\langle n|) \tag{1.19a}$$

$$H = 4E_c\left(i\frac{\partial}{\partial\varphi} - n_g/2\right)^2 - E_j\cos\hat{\varphi}(t) \tag{1.19b}$$

In the charge basis we have written the number of charge in terms of a quantized integer number that is offset by a gate charge, which is the number of excess cooper pairs on

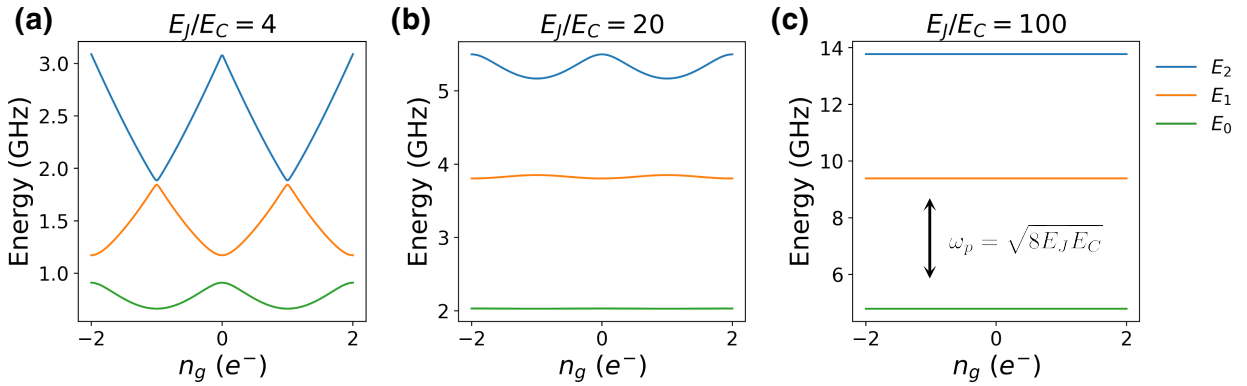

Figure 1.3: **Energy level splitting of a CPB for various** E_J/E_C:From the numerical solution to Eq 1.19 (b) we can see the change in dispersion, or dE/dn_g, for various ratios of E_j/E_c

the superconducting island. The inductive potential introduced by the Josephson junction introduces a tight-binding like operator that describes the coherent tunneling of discrete cooper pairs across the junction, with a tunneling energy equivalent to E_j. The phase-basis hamiltonian is derived by a transformation of $|n\rangle \rightarrow |\varphi\rangle$ and gives us the expected cosine correction.

The charge basis gives us intuition on the tunneling of cooper pairs, but requires us to diagonalize a sub-space of nearest neighbor charge-states to solve for the energy eigenvalues. The phase-basis hamiltonian on the other hand, which is a type of Mathieu equation, allows us to analytically solve the time-dependent Schrodinger equation, although higher order finite-differentiation can give enough numerical precision for most calculations.

From Fig. 1.3 we show the energy eigenvalues of the CPB Hamiltonian as a function of excess charge n_g. While our harmonic oscillator has inherited the nonlinearity of the Josephson junction, the anharmonic potential is now also dependent on these excess charge. The derivative $dE_{n,n+1}/dn_g$, or charge dispersion, gives us a sensitivity of our energies to excess charge. Small fluctuations in n_g effectively act as a time dependent perturbation leading to pure-dephasing of the qubit's state, and decoherence. For qubits that have charging energy similar to the tunneling energy $E_c \sim E_j$, we see that there exists values of n_g where the

qubit is first order insensitive to charge-noise. One can bias the CPB at this "sweet-spot" in excess charge, however doing so is challenging, and even at this point one is still sensitive to second order deviations in excess charge.

The solution to this problem is to decrease the charging energy of the CPB relative to the tunneling energy. From Fig. 1.3 we see that as the ratio of E_j/E_c increases, the effect of this dispersion decreases, but so does the anhamonicity that we fought so hard to introduce. The key however is to realize that the dispersion decreases *exponentially* with E_j/E_c, while the anharmonicity only decreases *algebraically*, meaning that a modest increase of E_j/E_c can greatly improve the operational lifetime of the qubit. This limit, where $E_j \gg E_c$ is what is known as the *Transmon Regime*, and qubits in this regime are subsequently called transmon qubits.

One nice property of the transmon qubit is that its anhamonicity is independent of the tunneling energy. This may seem counterintuitive, since the pre-factor of the anharmonic contribution in Eq. 1.18 is E_j. In the transmon limit, or where we can ignore the effect of n_g, we can show from Eq. 1.19 (b), once the $\cos\varphi$ term is expanded about small values φ, that the energy of the j^{th} level can be written as:

$$E_{trans,j} \approx \hbar\omega_a\left(j+\frac{1}{2}\right) - \frac{E_c}{12}(6j^2+6j+3) \tag{1.20}$$

Eq. 1.20 tells us that the transmon's anharmonicity is, to first order, only dependent on E_c, which we recall is determined by the capacitor of our now anharmonic LC circuit. Because of this, we can set E_c *first* to get the desired anharmonicity, before choosing an appropriate E_j that minimizes our dispersion, greatly simplifying the design. Here, ω_a, which is by analogy equivalent to the atomic transition energy in the previous section is therefore given by $\omega_a = \sqrt{8E_jE_c} - E_c$.

In a 3D cavity mode, to first order, the coupling between the qubit and cavity can be

thought of as a dipole [2], but in terms of a circuit analogy, the coupling rate g between the above transmon and a resonant cavity, which itself can be treated as a linear LC oscillator. The effective dipole coupling of a transmon to an LC oscillator to various qubit eigenstates can be given by the following [104]:

$$g_{ij} = 2\beta e V^0_{rms} \left(\frac{E_j}{8E_c} \right)^{1/4} \langle i | \left(b - b^\dagger \right) | j \rangle \tag{1.21}$$

Where $\beta = C_g/C_\sigma$, or the ratio between the capacitance to ground and the total capacitance of the coupled circuit, while V^0_{rms} is the zero-point voltage fluctuation on the resonator's capacitor, $V^0_{rms} = \sqrt{\hbar\omega_r/2C_r}$, and the product of $eV^0\beta$ is the bare dipole coupling constant g. The term $\langle i | \left(b - b^\dagger \right) | j \rangle$ gives the transmon charge matrix elements for various transitions, with the qubit eigenstates denoted by i and j. One point to note is that the effective capacitance of the CPB is modified by the additional capacitance to the resonator, and will be discussed in more detail in a later section.

As mentioned in the previous section, for typical cQED systems, the above value for $g \sim E_c$. In addition to this, the charging energy and subsequently the anharmonicity is much less than the detuning Δ in the dispersive regime, meaning that higher order states are all similarly detuned. The net effect of these two properties of cQED systems is that the effective dispersive shift has contributions from the $g - e$ transition as well as small additional shifts from higher energy states of the transmon. This leads to a dressed dispersive hamiltonian for the system:

$$H' = \frac{\hbar\omega'_a}{2} + \hbar(\omega'_r + \hbar\chi'\sigma_z)a^\dagger a \tag{1.22}$$

Here the effective dressed frequencies of the cavity and transmon are $\omega'_a = \omega_{01} + \chi_{01}$ and $\omega'_a = \omega_r - \chi_{01} - \chi_{12}$ respectively. The effective dispersive shift including the contribution

2. In reality the coupling is not so simple, as shown by Ref. [121]

of higher transitions then becomes $\chi' = (\chi_{12} - \chi_{01})/2$. We can compute χ for various transitions with:

$$\chi_{ij} = \frac{g_{ij}^2}{\omega_{ij} - \omega_r} \tag{1.23}$$

Plugging in Eq. 1.21 into Eq. 1.23 and computing the contribution of the first and second order transitions leads us to a χ' [33]:

$$\chi' \approx \frac{g^2}{2\Delta} \frac{\alpha}{(\Delta + \alpha)} \tag{1.24}$$

This is the effective dispersive shift from our qubit given the effect of the qubit's second-order transition. Higher order transitions can be considered, however their effect is small in comparison. For the purpose of this thesis, χ' as written above will be used when discussing the dispersive shift.

In a later section we will revisit the transmon when we discuss simulating the incorporation of a qubit into our cavity system, and discuss the way in which we model the coupling of this artificial atom to our microwave cavity via more advanced techniques.

1.2 Overview of thesis

The previous sections have laid the background and theoretical groundwork for the kinds of systems which will be discussed moving forward. In the remainder of this thesis we will see how we can use these fundamental building blocks, in particular the linear element of our cQED system, the cavity, to create an improved multimodal quantum architecture with state-of-the art mode coherences and qubit-cavity cooperativity. Next we will use this result as motivation for the creation of 3D cavities with improved coherences. We will do this by moving from aluminum to niobium as our superconducting material, discussing its unique properties and developments of a simplified etching technique capable of producing cavity

coherences $>$ 15 times that previously achieved in the same platform, with material loss comparable to state-of-the art accelerator cavities. We will then study the performance of these cavities as a function of manufacturing parameters. Finally we will correlate these results with material characterization of the surface and bulk properties. Next we will discuss the considerations for integrating a qubit into such high-coherence systems, before presenting a design that maintains cavity performance. We will then show that, through weak dispersive coupling, it is possible to integrate a lossy qubit into these niobium cavities without a significant degredation of the cavity performance. Finally we will discuss the challenges and techniques necessary for interacting with such a weakly coupled system.

CHAPTER 2

3D CAVITIES

2.1 Overview and motivation

In the previous chapter we discussed the mechanisms by which non-linear circuits can be coupled to microwave resonators and operated in a quantum limited manor. An important thread in cQED research has been the improvement of cavity and qubit coherence times. For quantum information processing longer coherences are crucial for increasing the circuit depth while decreasing the infidelity of quantum operations to reach true fault-tolerant quantum computing [102, 101, 55, 61, 171, 30]. In addition to computation, coherence times will play an important role in the scaling of quantum memories [124]. With this in mind, 3D superconducting cavities have emerged as a leading platform for everything from the creation of memories with coherence times of $T_1 > 1$ms [148, 39], to the creation of bosonic error-correcting codes and fault-tolerant systems [130, 92, 34, 152, 107]. These systems have been extended to multi-cavity systems demonstrating 2-mode gate operations [187, 158, 75].

Along with these results, cavities with coherence times in excess of $T_1 > 10$ms have been demonstrated in aluminum cylindrical cavities [147] and niobium coaxial cavities [89], with coherences $>$ 2s demonstrated in superconducting niobium SRF cavities at single-photon powers [157, 156]. To translate these high coherences into useful quantum systems however a number scientific and engineering problems must be solved.

While higher coherence modes can benefit any of the aforementioned systems, we will focus on building such cavities in the context of quantum memories, with more in-depth discussion about the way coherence allows for scaling quantum memory being done in Ch.4. This chapter however will lay the theoretical and technical groundwork for *what* goes into making a good cavity to begin with, motivating the novel *Flute Cavity* discussed in Ch.3.

2.2 Loss mechanisms in 3D cavities

2.2.1 Overview of loss

A perfect superconducting cavity would act like an ideal LC circuit, exhibiting zero dissipation and subsequently an infinite mode lifetime. In reality, like any real-world resonant circuit, superconducting RF cavities exhibit *some* dissipation. In later sections we will move to a quantum optics formalism, but for the time being lets build a more intuitive model for our microwave cavities using classical circuits. To first order we can model our resonator as a parallel RLC circuit with an additional series resistor as shown in fig. 2.1. The impedance, Z_{in} of this circuit is:

$$Z_{in}(\omega) = \left(\frac{1}{R_{tot}} + \frac{1}{j\omega L} + j\omega C \right)^{-1} \tag{2.1}$$

By stating that $\omega = \omega_0 + \Delta\omega$ and series expanding to lowest order in $\Delta\omega$ we can write the $Z_{in}(\omega)$ near resonance:

$$Z_{in}(\omega) \approx \frac{R_{tot}}{1 + 2j\Delta\omega R_{tot} C} \tag{2.2}$$

Here we see that, on resonance, or when $\omega = \omega_0$, $Z_{in}(\omega_0) \to R_{tot}$, where $R_{tot} = R_P + R_S^{-1}$. Our toy circuit, much like our real-world resonators, is a system with dissipation, and to quantify that dissipation we can look at the amount of energy *lost* per cycle, or in otherwords, 180 degree evolution of the phase. To do this we want a quantity that is simply the ratio of *stored energy* in both the magnetic and electric field components (or energy stored in either the inductive or capacitance element) vs the lost energy per unit time.

$$Q_0 = \omega \frac{W_m + W_e}{P_{loss}} \tag{2.3}$$

Where Q_0 is the *intrinsic quality factor*, sometimes denoted as Q_{int}. W_m and W_e are

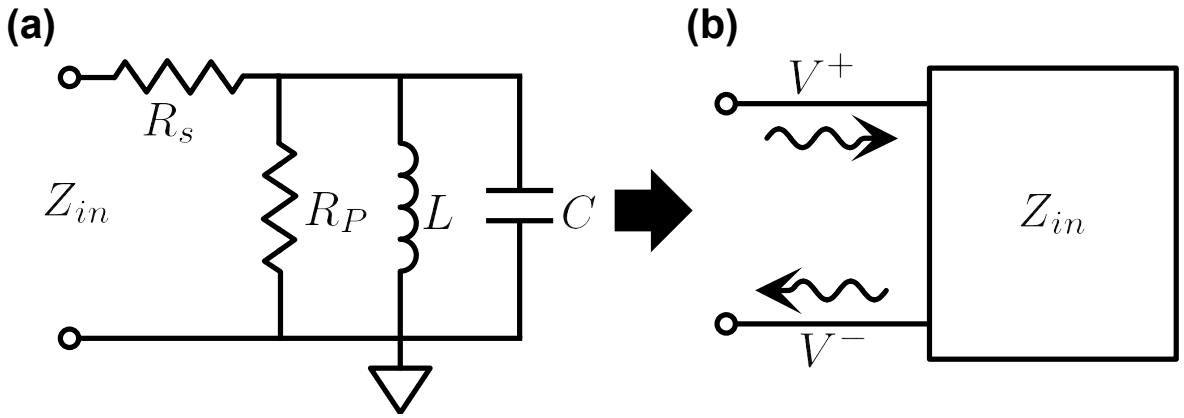

Figure 2.1: **Black-box equivalency of an RLC circuit** (a) shows the equivalent resonant circuit for a resonator which has a known impedance response $Z_{in}(\omega)$. By knowing the relative amplitudes and phase of V^+ and V^- it is possible to build equivalent circuit parameters to describe this resonator.

simply the energy stored in the magnetic and electric field of our system, with P_{loss} being the power loss through our resistor network R_P. For this circuit, the power loss is given by $P_{loss} = |V|^2/2R_{tot}$. On resonance the time-averaged stored magnetic and electric energies are equal, giving us that $Q_0 \propto R_P$ and $Q_0 \propto R_S^{-1}$. In essence Q_0 is the efficiency of our resonator, and depending on the material properties and use case of our resonant cavity, either R_P or R_S may be more important.

The resistance, R_P is sometimes called the *shunt impedance*, and for non-superconducting cavities, which typically have $R_P^{-1} \gg R_S$, has been the main point of optimization [93]. When R_P is small, the power dissipated across this resistance is high, which in cavities, manifests as surface heating and higher wall-power. In superconducting cavities $R_P > 1\times10^9$ [142], making its contribution to Q_0 for *most* situations negligible, but what about R_S?

Next lets finish writing Q_0 as a function of R_S. First we can compute the stored energy in the inductor by noting:

$$\int_0^{1/\omega} |P_m| dt = W_m = \frac{1}{4}\frac{|V|^2}{\omega^2 L} \tag{2.4}$$

Once again noting that at $\omega = \omega_0$ the stored electric and magnetic energies are equal,

we can then write eq. 2.3 as:

$$Q_0 = \frac{1}{\omega_0 L R_S} = \frac{\omega_0 C}{R_S} \tag{2.5}$$

Where we have also written Q_0 in terms of the energy stored in the capacitor.

Up to this point our treatment of the problem has assumed ideal passive components, meaning that L and C have no *intrinsic* loss. In reality these elements do have loss which are typically abstracted into R_S. Because of this, the value R_S is sometimes known as the *residual resistance* of the cavity and moving forward will be the primary material parameter to optimize. In the following sections we will expand on the above model and discuss the contributions of BCS conductivity, dielectric loss, and geometry on Q_0, before discussing the principals behind the $flute$ cavity design.

2.2.2 Coupling and Q_{ext}

Before discussing sources of loss in our otherwise ideal system, we have to address the effect of external coupling on the above circuit. In order to study the transients of our cavity we need to couple power into it. Doing so introduces an external load with a characteristic impedance Z_0. From the perspective of this coupled load, our resonator will have a frequency dependant impedance $Z_{in}(\omega)$ as defined in eq. 2.2. If we design our cavity such that the condition $Z_0 = Z_{in}(\omega_0)$, we would expect 100% of the power to be coupled into our circuit at resonance. Far away from resonance we expect the impedance mis-match to be large enough that 100% of the input power is reflected back into the port of our coupled device.

We can more rigorously describe this by looking at fig. 2.1 (b), which treats our resonant circuit as a black-box with an impedance of Z_{in}. At the port we can measure the incident and reflected voltage amplitudes V^+ and V^- respectively. If we attach a load with impedance Z_0, the current and voltage at the impedance is simply:

$$V_0 = V^+ + V^- \tag{2.6a}$$

$$I_0 = (V^+ - V^-)/Z_0 \tag{2.6b}$$

Setting the boundary condition that the current and voltage at Z_{in}:

$$Z_{in} = \frac{V_0}{I_0} = Z_0 \frac{V^+ + V^-}{V^+ - V^-} \tag{2.7}$$

What we really want to know, based on the frequency, is how much of the signal is reflected or absorbed into our circuit. This value is simply the ratio between reflected and incident voltage amplitudes, $S = V^-/V^+$. Substituting this into eq. 2.7 and solving for S in terms of the characteristic and input impedances we arrive at:

$$S_{11} = \frac{Z_{in} - Z_0}{Z_{in} + Z_0} \tag{2.8}$$

Where we use the standard nomenclature of denoting the scattering amplitude as S_{11} to represent a single-port reflective measurement. We can clearly see from eq. 2.8 that if the condition $Z_0 = Z_{in}(\omega_0)$ is met, then the scattering S_{11} will go to zero, whereas far from resonance, $Z_{in} > Z_0$, and the signal is scattered back into port 1. Fig. 2.2 illustrates this for a load Z_0 placed in series with our model circuit. To illustrate the dynamics we deliberately choose $Z_0 = \sqrt{L/C}$, with $\omega_0 = 1/\sqrt{LC}$. This means that, for $S_{11}(\omega_0) = 0$ to be true, on resonance $Z_{in} = R = Z_0$. Plotting $|S_{11}|^2$ for various ratios of Z_0/R in fig. 2.2 (a) illustrates this condition, but what does this mean in terms of the external quality factor?

In eq. 2.5 we showed the internal Q for our parallel circuit with a series resistance R_S. We can transform the parallel LC circuit values into an equivalent series circuit giving us $Q_0 = \omega_0 L_S/R_S$. The external quality factor for a series circuit can be expressed, in terms of the characteristic impedance, as $Q_{ext} = \omega_0 L_S/Z_0$. We can see that we can rewrite our ratio

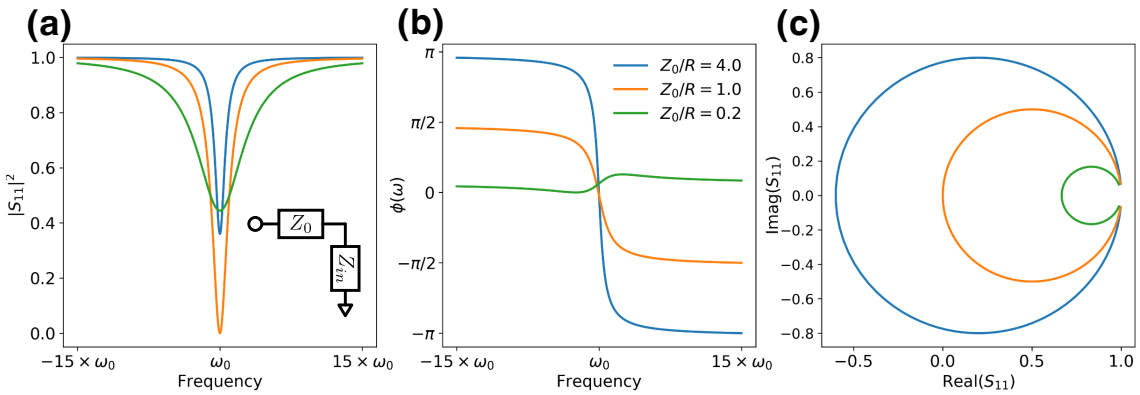

Figure 2.2: **Frequency response of phase and amplitude of a single-sided resonator**: (a) shows the scattering amplitude $|S_{11}|^2$ for various coupling regimes of $Q_{ext} > Q_{int}$, $Q_{ext} = Q_{int}$, or $Q_{ext} < Q_{int}$. (b) shows the phase response for these same regimes. (c) shows the resonance mapped onto the $I - Q$ plane, thereby encapsulating both magnitude and phase response.

of characteristic impedance to resistance in terms of Q_0 and Q_{ext} as $Q_0/Q_{ext} = Z_0/R_S$. This tells us that, when $Q_0 = Q_{ext}$, we have the *maximum* amount of power being coupled into our resonator. When this condition is met the resonator is *critically coupled*, where as if $Q_0 < Q_{ext}$ or $Q_0 > Q_{ext}$ the resonator is *undercoupled* or *overcoupled* respectively. This means that, just by looking at the magnitude (fig. 2.2 (a)) and phase response of our system (fig. 2.2 (b)), we can deduce not just the quality factor of the resonator, but also Q_0 and Q_{ext}, and subsequently the effective residual resistance R_S. This "black-box" treatment of the resonator will become indispensable when measuring 3D cavities, as the equivalent circuit parameters, L and C, are not well defined, as the electric and magnetic potential of our field is distributed over the entire mode volume of our system.

The final point in our discussion of Q_{ext} is the effect on the system's *total* quality factor, Q_{tot}. From the perspective of the resonator, the addition of Z_0, is the equivalent of adding an additional series resistance R_L. If we then ask what the *new* Q of our system is, we know it should be $Q = \omega_0 L_S/(R_L + R_S)$. Solving for this Q in terms of Q_{ext} and Q_0 finally gives the relationship:

$$Q_{tot} = \left(\frac{1}{Q_0} + \frac{1}{Q_{ext}}\right) \tag{2.9}$$

This new Q_{tot} represents the effective quality factor of the system. So far we have discussed this coupling as if it were a black-box with some known characteristic impedance that can be treated as a series resistance, but in later sections we will discuss how this coupling is controlled in 3D cavity systems, both for directly interrogating the cavity and coupling non-linear circuits to the cavity mode.

2.2.3 Input-output formalism for cavities

In the previous section we had derived the effective response of the cavity in the frequency domain based on the input and characteristic impedance of our system. These values are based on the particular circuit parameters of the LC oscillator, and the input load in question. In 3D cavity systems, however, the equivalent circuit parameters are not well defined. Instead, the inductive and capacitive elements (and their stored energy) are distributed. Because of this, it is beneficial to write the scattering, as described in Eq. 2.8, in terms of the internal quality factor, Q_0, and the coupling quality factor, Q_{ext}. By plugging in Eqs. 2.2 and 2.5 into Eq. 2.8:

$$S_{11}(\omega) = \frac{\omega_0(Q_{ext} - Q_0 e^{j\phi}) - 2j(\omega - \omega_0)Q_0 Q_{ext}}{\omega_0(Q_{ext} + Q_0 e^{j\phi}) + 2j(\omega - \omega_0)Q_0 Q_{ext}} \tag{2.10}$$

Here we have introduced a complex phase $e^{j\phi}$ that results in asymmetry in our lineshape due to impedance mismatch between our coupler and the output line. A common way to write Eq. 2.10 is in terms of the intrinsic cavity linewidth, κ_i, and the coupling rate, κ_c. The relation between Q_0 and κ_i is:

$$Q_0 = \frac{\omega_0}{\kappa_i} \tag{2.11}$$

Plugging into Eq. 2.10 gives:

$$S_{11}(\omega) = \frac{\kappa_i - \kappa_c e^{j\phi} - 2j(\omega - \omega_0)}{\kappa_i + \kappa_c e^{j\phi} + 2j(\omega - \omega_0)} \tag{2.12}$$

The above expression is more commonly used in cavity quantum electrodynamics and quantum optics [185, 86], but here we showed that it is equivalent to describing the scattering of an AC signal off of our idealized LC-circuit.

The above represents the result for a single sided cavity, however you can also derive an expression that describes the *transmission* of a signal through a cavity on resonance, as opposed to the reflected signal. For a parallel RLC oscillator, we can use the transformation provided by Table 4.2 in Ref. [144], and more carefully derived in Ref. [116] to calculate the scattering matrix element S_{21} to get the expression:

$$S_{21}(\omega) = \frac{1}{1 + 2jQ_L(\omega - \omega_0)/\omega_0} \tag{2.13}$$

A couple of things to note with the above equation, is that $|S_{21}|^2$ takes the form of a lorentzian, with $Q_L = (1/Q_0 + 1/Q_{ext})^{-1}$ being related to the lorentzian line width by noting that $\omega_0/Q_L = \Gamma$. The second point to note is that the above has been normalized to a peak height of 1. In reality the height of the peak is set by the input-output scattering rates. A more natural treatment for cavity systems, as before, is to use a quantum optics formalism. For a two-sided cavity, such as the one shown in Fig. 2.3 the measured value is determined by the ratio of input and output scattering amplitudes b_{in}/a_{out}.

Following the method outlined in Ref. [185], we can solve the equation of motion for the system to get:

$$S_{21}(\omega) = \frac{b_{in}}{a_{out}} = \frac{\sqrt{\gamma_1\gamma_2}}{\frac{\gamma_1 + \gamma_2}{2} + j(\omega - \omega_0)} \tag{2.14}$$

Above we have written the rates as γ, which is the more common notation in quantum

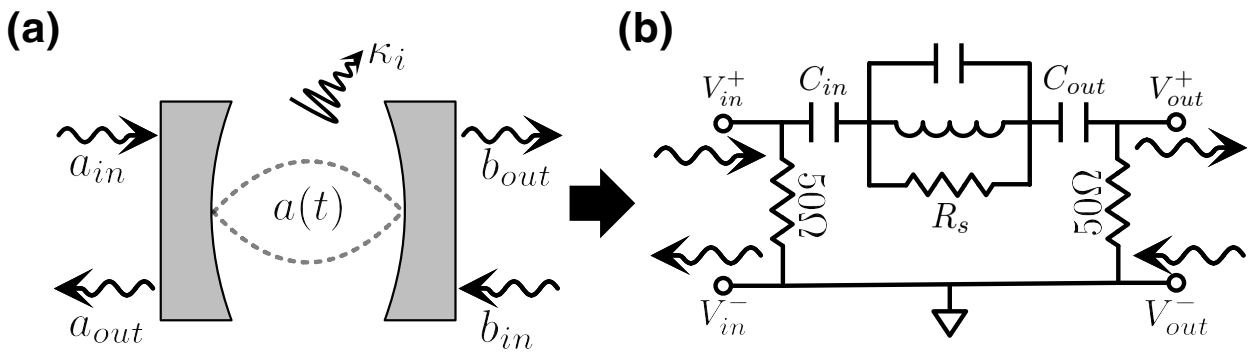

Figure 2.3: **The scattering off of a two-sided cavity** (a) The effective measured S_{21} for our system is the ratio of scattering amplitudes b_{in}/a_{out}. By solving for the equation of motion and solving for this ratio we can find the evolution of the cavity state $a(t)$. By taking the fFourier transform of this time-dependent state evolution we can recover the frequency response of our resonator as a function of input and output coupling. Here, κ_i represents internal loss of our cavity, which factors into the evolution of $a(t)$. (b) shows an effective circuit model for our two-sided cavity system. The RLC circuit is our resonant cavity, with R_s giving rise to the intrinsic linewidth, as described in the previous section. Our cavity is coupled via 50Ω ports which are capacitively coupled to the system. The effective input impedance of this coupling sets the Q_{ext} of the port.

optics, but it should be noted that $\gamma = 2\kappa$ as we had previously defined it. The above also assumes that the cavity has no intrinsic linewidth, but rather the quality factor is set only by output couplings. We see that for $\gamma_1 = \gamma_2 = \gamma$ and setting $\gamma = 2\omega_0/Q_L$, we get back Eq. 2.13. We can rewrite Eq. 2.14 in terms of κ, adding in the internal loss to arrive at the form commonly used:

$$S_{21}(\omega) = \frac{2\sqrt{\kappa_1\kappa_2}}{(\kappa_1 + \kappa_2 + \kappa_i) + 2j(\omega - \omega_0)} \tag{2.15}$$

Where we use the subscript 1,2 and i to denote the decay rates of coupler 1,2 and the internal decay rate respectively.

2.2.4 BCS Theory and non-ideal superconductors

Before we can begin a more detailed discussion of the cavity geometry and its effect on performance, we have to first discuss the nature of superconductivity in RF systems. This

section will not be an exhaustive discussion of the microscopic or phenomenological treatment of superconductivity, but will build a basic intuition for the system at hand. For a more detailed description, Refs. [175, 134, 133] are indispensible.

In a metal, the electrons can be treated as distributed "sea" of charge which follow Fermi statistics, with a distribution of energies up to the Fermi energy E_f. Cooper showed that, given even a very weak attraction between electrons, the formation of pairs of electrons into stable bound states was possible at low enough energy scales [51]. An explanation for this attraction came from the interaction of the electrons with the ionic lattice of the metal, whereby vibrations in the positively charged nuclei at the lattice site, lead to small scale polarization of the ionic site, with the excess positive charge attracting an additional electron. If this attraction is greater than the coulomb screening than a net attractive force between the two electrons is achieved. This interaction, which can be written in second-quantized form as a phonon-electron scattering Hamiltonian (or *pairing Hamiltonian*), can be solved using variational theory to give insight to the scale of energies, and subsequently temperatures, necessary for superconductivity to be achieved [18]. This treatment also gives rise to a few additional observations, namely that the wave functions of the combined pairs of electrons all overlap, as opposed to the normal state electrons that follow Fermi statistics, resulting in the paired electrons acting as a condensate. The gap in energy between this condensate of electron pairs and normal state electrons is large compared to the energy scale of intrinsic phenomena in the lattice, meaning that the condensate interacts little with its electronic environment. This result gives us the zero DC resistivity which characterize superconductivity as a whole, but more importantly, it gives us a way of connecting superconductivity to measurable quantities of our superconducting microwave cavities.

One outcome of this microscopic theory of superconductors, described by Bardeen, Cooper, and Schrieffer (BCS theory) in 1957, is the estimation of the energy gap, $\Delta(T)$, between bound electron pairs and unpaired quasi-particles. This temperature value can be extracted

numerically from the transcendental expression for the system at finite temperature:

$$\frac{1}{N(0)V} = \int_0^{\hbar\omega} \frac{\tanh \frac{1}{2}\beta(\xi^2 + \Delta^2)^{1/2}}{(\xi^2 + \Delta^2)^{1/2}} d\xi \tag{2.16}$$

Where ξ (not to be confused with electron coherence length discussed later) is a measure of electron energy relative to the Fermi energy, $N(0)$ is the number of cooper pairs at $T = 0K$, or rather the density of states at the electron Fermi surface. In addition, V is the *attractive* potential between electrons, and $\beta = 1/k_b T$. It can be shown that for weakly interacting, or "conventional" superconductors the above can be generalized, and the relationship between Δ and T can be described as a function of T/T_c, with T_c being when $\Delta(T) \to 0$, for $T \leq T_c$:

$$\frac{\Delta(T)}{\Delta(0)} \approx 1.74\left(1 - \frac{T}{T_c}\right)^{1/2} \tag{2.17}$$

With the relationship of T_c and the associated frequency $\hbar\omega_c = \Delta(T_c)$ being:

$$k_b T_c = 1.13\hbar\omega_c \exp(-1/N(0)V) \tag{2.18}$$

The above, plotted in Fig. 2.4(a), gives us a high level overview of several important details that will become apparent later in the discussion of RF superconductivity. The first is that, at temperatures above $\sim 0.4T_c$, the gap $\Delta(T) \approx \Delta(0)$, with a rapid decrease in the gap energy approaching T_c for $T > 0.8T_c$. A result of this is that, as the energy of the system increases, the gap energy goes to zero, the density of cooper pairs to normal state electrons must decrease. This ratio can be expressed as:

$$\frac{N_s(E)}{N(0)} = \frac{d\xi}{dE} = \begin{cases} \dfrac{E}{(E^2 - \Delta^2)^{1/2}} & (E > \Delta) \\ 0 & (E < \Delta) \end{cases} \tag{2.19}$$

Where $\Delta = \Delta(0)$, and is shown in Fig. 2.4 (b). The above states that as temperature increases, or the electrons are excited with electromagnetic radiation with $\omega \geq \omega_c$ the number

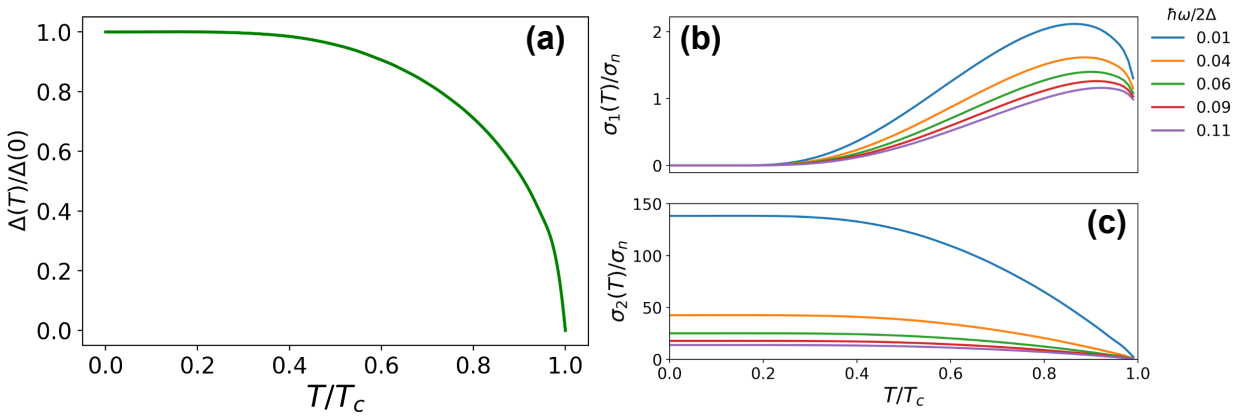

Figure 2.4: **Superconducting gap, and Mattis-Bardeen conductivities**: (a) shows the generalized superconducting energy gap for various values of T/T_c. (b) and (c) show the real and imaginary conductivities, σ_1 and σ_2, for various frequencies ω, and temperatures, as derived from Eqs. 2.23 (a) and (b). These integrals must be numerically computed for different values of T/T_c. Details on the above and superconductors in the presence of high frequency electromagnetic fields can be found in Ref. [64]

of normal state electrons also increases, and now interact with the lattice, leading to normal state resistivity.

So far the microscopic theory we have outlined has been agnostic to whether the electrons are under the influence of a DC or AC field. As previously mentioned, the condensate of cooper pairs do not interact with the lattice. So what does the superconductor do in the presence of an AC field? For that we can look at the predictions made by Mattis-Bardeen (MB) theory [115, 125].

A more natural and intuitive way of thinking about the AC case is to take a more phenomenological approach. Previously we have shown that, at finite temperature there will always be *some* electrons that exist in a normal state. In addition to the *equilibrium* normal state electrons, we can have *non-equilibrium* electrons that form when additional sources of energy cause the breaking of cooper pairs. In our metal, normal-state, which obey Ohm's law, and superconducting electrons will co-exist. In DC, the normal-state electrons are effectively shielded via coulomb screening from the superconducting electrons. In AC, however, the inertia of the cooper pairs leads to a non-zero participation of normal state

electrons. This can be treated as a *two-fluid* model, which can be abstracted as a parallel lossless inductor, and a resistor. The inductance acts in place of the inertia provided by the electrons themselves, giving a conservative force in the presence of an alternating field, leading to some fraction of electrons passing through the resistive path. In our previous 2D circuit model, as shown in Fig. 2.1 (a), this would act as the source of our additional shunt resistance and inductance.

This analogous circuit model now has a complex impedance Z_s. In the two-fluid case, the impedance can be written in terms of a real conductance, σ_n, which is just the normal state conductance, and a complex conductance, σ_s, which contributes to the reactance of the Z_s due to the inertial effects of the cooper pairs. It can be shown that in the two-fluid model the impedance can be written in terms of conductance as [134]:

$$Z_s = \sqrt{\frac{i\omega\mu_o}{\sigma_n - i\sigma_s}} \tag{2.20}$$

From the first London equation we can show that $\sigma_s = (\omega\mu_o\lambda_L^2)^{-1}$, where we have used the definition of the London penetration length, $\lambda_L = \sqrt{m/\mu_0 n_s e^2}$, where n_s is the number of paired electrons. σ_n is just the normal state conductivity. In the case of our superconductor, $\sigma_s \gg \sigma_n$. This allows us to take Eq. 2.20, and make the approximation:

$$Z_s \approx \sqrt{\omega\mu_0}\sqrt{\frac{i\sigma_n}{\sigma_s^2} - \frac{1}{\sigma_s}} \tag{2.21}$$

We can recast the terms inside of the right square root as $x = \sigma_n/\sigma_s^2$ and $y = 1/\sigma_s$, noting that $x \ll y$. Taylor expanding about $x = 0$, and truncating at leading terms in x. Substituting back σ_s and canceling terms gives us the expression:

$$Z_s \approx \frac{1}{2}\sigma_n\omega^2\mu_0^2\lambda_L^3 + i\omega\mu_0\lambda_L \tag{2.22}$$

We can see from above $R_s = \sigma_n\omega^2\mu_0^2\lambda_L^3/2$ and $X_s = \omega\mu_0\lambda_L$.

There are a couple of important takeaways from Eq. 2.22. Z_s is dominated by the residual surface resistance R_s, which is quadratically dependent on the frequency ω, and on the normal state conductivity σ_n, which is proportional to the number of normal-state electrons. We also note that the resistance is now dependent on the material specific London-penetration length, λ_L. We also know, from Eq. 2.5, that the internal quality factor, Q_0 is inversely proportional to the parallel resistance R_s. This means that, as the number of normal-state electrons increase (and cooper pairs are broken) the quality factor of our cavity must decrease.

The two-fluid model, and the approximation for Z_s, gives us important intuition about the performance of a superconductor in an AC field, and is a good approximation for $\omega \ll \omega_c$. In reality, to get an accurate approximation for R_s, especially in materials with smaller gaps, such as aluminum, we must include the predictions made by BCS theory. A full treatment of the two-fluid model using BCS theory was done by Nam in 1967 [125].

Nam rewrote the complex conductance as a ratio between $\sigma_s/\sigma_n = \sigma_1 - i\sigma_2$. Here we take $\sigma_n \propto 1/\lambda_L$ with $\lambda_L = ne^2/m = \frac{2}{3}N(0)e^2v_0^2$, with $N(0)$ and v_0 being the density of states and the mean velocity of normal state electrons at the Fermi surface. Nam showed that the conductivities can subsequently be written, following the form used in [149] as:

$$\sigma_1 = \frac{1}{\omega}\int_{\Delta}^{\infty} \frac{(\epsilon+\omega)\epsilon + \Delta^2}{\sqrt{(\epsilon+\omega)^2 - \Delta^2}\sqrt{\Delta^2-\epsilon^2}}\left[\tanh\frac{\epsilon+\omega}{2\tau} - \tanh\frac{\epsilon}{2\tau}\right]d\epsilon \tag{2.23a}$$

$$\sigma_2 = \frac{1}{\omega}\int_{\Delta-\omega}^{\Delta} \frac{(\epsilon+\omega)\epsilon + \Delta^2}{\sqrt{(\epsilon+\omega)^2 - \Delta^2}\sqrt{\Delta^2-\epsilon^2}}\left[\tanh\frac{\epsilon+\omega}{2\tau}\right]d\epsilon \tag{2.23b}$$

Where $\tau = \hbar/k_bT$ and we have used the identity $\tanh(x) - \tanh(x-a) = 2(1/(e^x+1) - 1/(e^{x-a}+1))$ and $\tanh(x) = 1 - 2/(e^x+1)$ to rewrite the Fermi-Dirac distributions in the integrand. These two integrals must be solved numerically, with the addition of $\Delta(T)$ described in Eq. 2.17.

Nam also gave an approximate form of Z_s in terms of the normal impedance Z_n using σ_1 and σ_2 gives:

$$Z_s \approx Z_n(\sigma_1 - i\sigma_2)^{\nu} \tag{2.24}$$

Where ν can either be $-1/2$ for the dirty (Pippard) limit, or when the mean-free length $l \ll \xi_0$, where ξ_0 is the coherence length,[1] giving us an equation resembling Eq. 2.21, or $-1/3$ for the clean limit, or when $l \gg \xi_0$.

With the full BCS description of Z_s in hand we can now determine a range of important parameters for our cavity systems. Following a similar approach to what was done for the simple two-fluid approximation, below T_c we can take $\sigma_2 \gg \sigma_1$. Following Ref. [149] we can show that in this limit we can approximate R_s and X_s as follows:

$$R_s \propto R_n \sigma_1 \sigma_2^{\nu-1} \tag{2.25a}$$

$$X_s \propto R_n \sigma_2^{\nu} \tag{2.25b}$$

For our parallel inductor-resistor model, and remembering back to Sec. 2.2.2, the quality factor can be written as $Q_{mag} = X_s/R_s$, where $X_s = \omega L$, or the reactance of our equivalent lossless inductor in our two-fluid model. Substituting Eqs. 2.25a and 2.25b we can write Q_s as

$$Q_{mag}(T, T_c, \omega) \propto \frac{\sigma_1(T, T_c, \omega)}{\sigma_2(T, T_c, \omega)} \tag{2.26}$$

It is important to note that the value Q_{mag} gives us the surface quality of our superconductor. This value implies that, regardless of the geometry of our 3D (or 2D) resonator, the

1. or the length scale in which the two electrons making up the cooper pair are correlated

material is the only contributing factor in determining the quality factor of our resonator. As we will see in later sections, however, geometry *does* in fact play a large roll in the measured quality of the cavity that we wish to make. This is because only part of the magnetic energy participates in the lossy surface of our superconductor, with the other fraction living in the lossless vacuum.

2.2.5 Dielectric and two-level system loss

In sec. 2.2.4 we discussed the intrinsic loss of the superconducting materials we will use to make our cavities. This loss, however, only interacts with the surface currents, and therefore interacts with the magnetic field of the electromagnetic mode. For a normal metal conductor, a good approximation is to take $\sigma \to \infty$ at the metal surface, meaning that the $\hat{n} \times \vec{E} = 0$ at the wall. This assumes that the vacuum/metal interface is "clean", however, most metals—and in particular the one's we will discuss in later chapters—have native oxides that live at the surface. These oxides can have dielectric qualities, as is the case with aluminum oxide, Al_2O_3, or can be made up of a number of stable oxides that have dielectric, lossy dielectric/conductive, and even superconducting properties at our operating temperatures, as is the case for Nb_2O_5, NbO_2, and NbO respectively. In the dielectric case, energy can be dissipated through dipole relaxation and unbound charge carriers on the surface. In addition to this, the surface oxide is amorphous. This high disorder leads to localized electronic states with non-zero tunneling potential which can couple to the electric field and act like two-level system (TLS) [138, 139, 73].

In a perfect dielectric, the individual dipoles will be perfectly, and instantaneaously, polarized in the direction of the field. In reality, there is a lag in the polarization, with a characteristic *relaxation time* scale τ_0 caused by inertial effects. By analogy, this is similar to the hysteresis seen in magnetic domains, which lead to hysteretic losses in inductive AC systems. Subsequently, the dielectric of a material can be written in terms of a real

component ϵ' and a reactive, or complex, component ϵ'':

$$\epsilon = \epsilon' + i\epsilon'' \tag{2.27}$$

This loss can be modelled as an effective series resistance in our microwave circuit, which gives a quality factor:

$$Q_\epsilon = \frac{\epsilon'}{\epsilon''} = \frac{1}{\tan(\delta)} \tag{2.28}$$

Where the expression for Q_ϵ is inversely proportional to the *loss tangent* $\tan(\delta)$. The above picture, however, presumes that the effective dielectric loss is frequency independent, but for relaxation processes, this is not the case. In the presence of a varying field, $E(\omega) = E_0 \sin(\omega t)$, it can be shown that [94]:

$$\epsilon'(\omega) = \epsilon(\infty) + \frac{\epsilon(0) - \epsilon(\infty)}{1 + \omega^2\tau_0^2} \tag{2.29a}$$

$$\epsilon''(\omega) = \frac{(\epsilon(0) - \epsilon(\infty))\omega\tau_0}{1 + \omega^2\tau_0^2} \tag{2.29b}$$

Which gives a frequency dependent loss tangent:

$$\tan(\delta(\omega)) = \frac{(\epsilon(0) - \epsilon(\infty))\omega\tau_0}{\epsilon(0) + \epsilon(\infty)\omega^2\tau_0^2} \tag{2.30}$$

Where $\epsilon(0)$ and $\epsilon(\infty)$ are the effective dielectric constants at zero and infinite frequency respectively. The result of this is that the dielectric loss tangent is not just material and environmentally dependent for determining τ_0, $\epsilon(0)$ and $\varepsilon(\infty)$, but also frequency dependent. Eq. 2.30 shows a maxima when $\omega = (\varepsilon(0)/\varepsilon(\infty))^{1/2}\tau_0^{-1}$, while the imaginary part of the dielectric constant in 2.29b exhibits a maximum when when $\omega = 1/\tau_0$. In reality, for inorganic dielectrics, $\tau_0 \geq 10^{-12}$s, meaning that for low-frequency dielectric constants of $\varepsilon \sim 10 - 35$, the loss-tangent maxima occurs at $\omega > 10^{12}$Hz [161, 23]. This also pre-

sumes that the material is crystalline with no defects or impurities, which may lead to local variation in relaxation time. This relaxation time is also temperature dependent, following an Arrhenius relation of $\tau \propto \tau_0 \exp\{(E_r/k_b T)\}$, where E_r is the energy associated with the relaxation phenomena.

For aluminum's native oxide, Al_2O_3, which exhibit anisotropic dielectric constants in it's pure crystalline phase, the values have been thoroughly investigated [85]. In addition to this, loss tangents for pure sapphire have also been investigated [105], however amorphous aluminum oxide has not been well characterized at low temperatures or microwave frequencies. Inversely, Nb_2O_5, niobium's primary native oxide, has not been well characterized in the microwave regime, however some investigations of amorphous thin film oxides have yielded $\epsilon' > 30$ for frequencies above 1MHz, with relaxation being attributed to uncorrelated charge carrier hopping [71].

Besides relaxation phenomena, the surface dielectric can host local defects that can couple to the electromagnetic environment and act like two-level systems. These defects, while not microscopically well defined, exhibit discrete energy spacing and dipole coupling to their environment. Refs. [139, 138] created the first complete model, known as the standard tunneling model (STM), to describe the dynamics of these defects. These TLS have been experimentally observed in a wide range of superconducting devices, from Al/AlOx/Al junctions [114] and 2D resonators [74, 73, 136], to SAW resonators [5].

In the STM, the eigenstates of the local charge states live in an assymetric double-well potential. The Hamiltonian is given as:

$$\hat{H} = \frac{\hbar}{2} \begin{pmatrix} \delta & \Delta \\ \Delta & -\delta \end{pmatrix} \tag{2.31}$$

Where δ is the double well energy asymmetry, and Δ is the tunneling element. Diagonalizing Eq. 2.31 and solving for the energy eigenvalues of $\epsilon_\pm = \pm\sqrt{\delta^2 + \Delta^2}$. Rewriting our

Hamiltonian gives:

$$\hat{H}_0 = \frac{\hbar}{2}\epsilon\sigma_z \tag{2.32}$$

In the presence of an electric field, the interaction Hamiltonian can be written as:

$$\hat{H}_{int} = \left(\frac{\delta}{\epsilon}\sigma_z + \frac{\Delta}{\epsilon}\sigma_x\right)\vec{d_0}\cdot\vec{E} \tag{2.33}$$

Phillips, in his derivation, makes the case that barrier height is largely unaffected by the interaction of the field, with only perturbations in δ giving rise to the transition from eigenstate ψ_- to ψ_+ [138]. From this we can deduce that the first and second terms are the transition $\vec{d}\,'$ and static, $\vec{d}$, dipole moments respectively.

By making a correspondence between the above interaction picture to that of a spin-1/2 system in the presence of an oscillating perturbative magnetic field, Gao was able to show that, at low temperature and short time scales, resonant interactions with the TLS were dominant. In [136] it was shown that, for cases where the Rabi-drive strength, Ω is small, and $\Omega^2 T_1 T_2 \ll 1$, where T_1 and T_2 are the TLS energy and phase coherence relaxation time scales, the effective loss tangent of the TLS becomes:

$$\tan\left(\delta_{TLS}(T)\right) = F\tan\delta^0_{TLS}\tanh\left(\frac{\hbar\omega}{2k_BT}\right) \tag{2.34}$$

Where F is the filling factor, or the ratio of the participation of the field at the surface vs the volume of our resonator. As in Eq. 2.28, $Q_{TLS} \propto \tan\left(\delta_{TLS}(T)\right)^{-1}$. In the case of 3D microwave resonators, this limit is a good approximation, thanks to the fact that the energy participation of the electric field at the surface compared to the volume is low, meaning that, at typical cavity drive strengths of mean photon number $\bar{n} \sim 1-10^{10}$, $|\vec{E}_{surf}| \ll |\vec{E}_c|$, where $|\vec{E}_c|$ corresponds to the electric field strength where all of the TLS are being driven resonantly, and $|\vec{E}_{surf}|$ is the strength of the field at the surface. Inversely, in planar resonators, where

the electric field interacts more strongly at the surface, the weak-field limit does not hold.

In the strong-field limit, the Rabi frequency, and subsequently the procession of the TLS are far faster than T_1 and T_2, meaning that mean occupation of the TLS in either ground or excited state is 50%. Once $|\vec{E}| = |\vec{E}_c|$ however, all of the TLS are saturated, and minimum loss tangent has been reached. Because $\Omega \propto |\vec{E}|$, the term $\Omega^2 T_1 T_2 \sim |\vec{E}/\vec{E}_c|^2$, which can be shown gives:

$$\tan\left(\delta_{TLS}(T, |\vec{E}|)\right) = \frac{F \tan\left(\delta^0_{TLS}\right) \tanh\left(\dfrac{\hbar\omega}{2k_B T}\right)}{\sqrt{1 + |\vec{E}/\vec{E}_c|^2}} \tag{2.35}$$

Above we introduced the idea of a filling fraction F, and previously we have also discussed how intrinsic dielectric and conductive loss on the system, and subsequently the measured quality factor, is dependent on the amount of the field's stored energy lives in the surface. We have also mentioned that the magnetic losses in Q_{mag} discussed in sec. 2.2.4 are dependent on the effective participation of the electric and magnetic fields with the aforementioned surface loss mechanisms. Because of this, geometry plays an equally important roll in quality factor as surface quality, as will be discussed in the next sections.

2.3 3D resonators

Up to this point we have discussed resonators that are made up of discrete reactive components, and in the abstract in terms of their loss mechanisms. This analogy is informative in describing the transient response of a 3D resonator, but fails to capture the way in which the electromagnetic energy of the 3D system is distributed, or how geometry effect things like the mode dispersion of our resonator. The distribution of energy in our system plays an important roll both in coupling of our non-linear elements, which in our case are transmon qubits, but also in investigating the way that this electromagnetic energy couples to potential sources of loss. Just as the energy is distributed, so is our dissipative element. Before

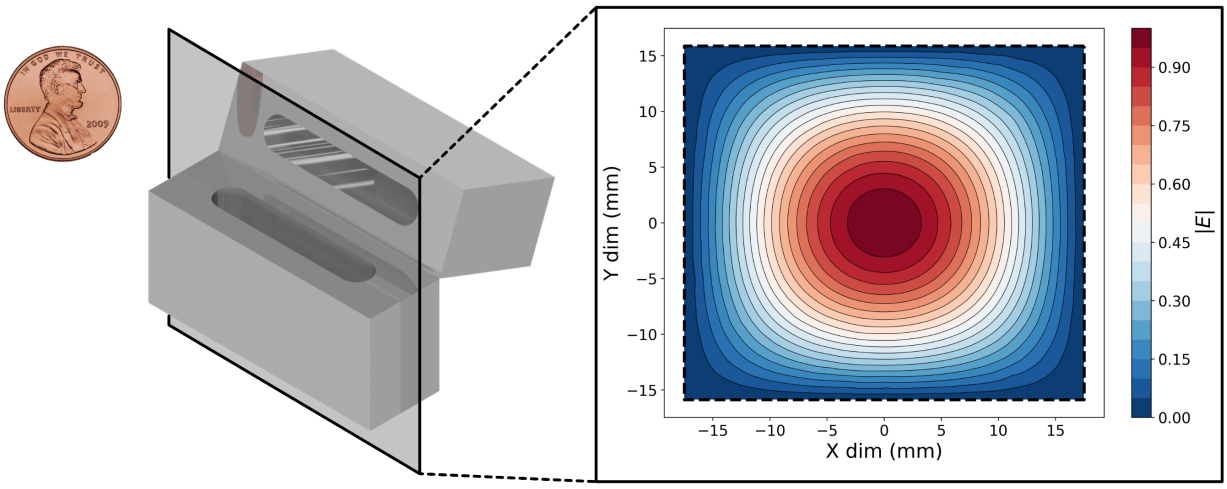

Figure 2.5: **Rectangular waveguide cavity cross-section and TE_{101} E-field distribution**: A render, presented to scale with a penny, of a typical two-piece rectangular waveguide cavity, with a cross-section of the $|E|$ profile for the fundamental TE_{101} mode, showing the characteristic "drum-head" profile, with maximum electric field magnitude at center.

discussing the geometric dependence of this loss, we first need to discuss the basics of *how* this energy is distributed in 3-dimensions.

A natural place to begin is the shunted rectangular waveguide cavity, both because it has a straightforward analytic solution, but also because it will act as the basis for the flute cavity discussed in ch.4.2. It will also play host to some more detailed analysis of loss in later sections. This, and the following discussions, should be taken as an exercise for the purpose of intuition; in reality, most electrodynamics problems are intractable analytically, with no nice closed-form solutions. Even for the coaxial $\lambda/4$ and tapered TE_{101}, there is no nice solution for the electromagnetic fields. In the next chapter we will discuss the use of FEM and automated numerical techniques for simulating and understanding these systems in depth, but for now, lets continue with our simple rectangular waveguide example. Thorough derivations of some simple and relevant geometries may be found in Refs. [144, 50, 149], among others.

The rectangular cavity is equivalent to a rectangular waveguide with either end shorted. With this in mind we can begin with the generic wave equation for an electric field moving

along an axis of the cavity geometry.

$$E(x, y, z) = E_{x,y}(A^{+}e^{-j\beta_{nm}z} + A^{-}e^{j\beta_{nm}z}) \tag{2.36}$$

By convention we will call this axis z, and for the lowest frequency mode, corresponds with the longest axis of the cavity. We have also separated the electric field profile in x and y and left them as $E_{x,y}$ for the time being.

We have introduced a few important variables above. The first is the *propagation constant* or wavenumber, β_{nm}. In reality, the propagation constant of a wave is a complex number with a real amplitude α, known as the attenuation constant, and complex amplitude β or the phase constant. For a wave moving through a lossless medium, $\alpha = 0$, and only the complex phase constant remains. In the case of a waveguide where the transverse fields, $E_{x,y} \neq 0$, and the propagating wave $E_z = 0$, then to satisfy the Helmhotz wave equation, all propagating fields must have a frequency higher than a certain frequency, known as the *cutoff frequency*, ω_c, which has a corresponding wavenumber k_c. The propagation constant in the case of a TE waveguide mode for our rectangular waveguide cavity can be rewritten as:

$$\beta_{nm} = \sqrt{k^2 - k_c^2} \tag{2.37}$$

For our rectangular waveguide the propagation of TE_{nm} and TM_{nm} modes only occurs when:

$$\omega_0 \mu\varepsilon > \left(\frac{m\pi}{l_{cav}}\right)^2 + \left(\frac{m\pi}{l_{cav}}\right)^2 = k_c \tag{2.38}$$

To determine the scattering coefficients in Eq. 2.36 we note that at one wall ($z = 0$), $E(x, y, z)$ is zero, which gives $A^{+} = -A^{-}$. From a physical perspective this makes sense; a perfectly conducting wall should reflect back 100% of it's incident energy. Substituting this

into Eq. 2.36 we get:

$$E(x,y,z) = E_{xy}A^{+}(e^{\beta_{nm}z} - e^{-\beta_{nm}z}) = -2iA^{+}E_{xy}sin(\beta_{nm}z) \tag{2.39}$$

At the second wall, where $z = l_{cav}$, we also set $E(x,y,z) = 0$. We note for this condition to be true, then $\beta_{mn}l_{cav} = l\pi$. This basically says that only waves with a half wavelength that is a integer multiple of l_{cav} can exist in the cavity. This gives us the resonant conditions:

$$k_{mnl} = \sqrt{\left(\frac{m\pi}{h_{cav}}\right)^2 + \left(\frac{n\pi}{w_{cav}}\right)^2 + \left(\frac{l\pi}{l_{cav}}\right)^2} \tag{2.40}$$

And subsequently a resonant frequency:

$$f_{mnl} = \frac{c}{2\pi\sqrt{\mu\varepsilon}}\sqrt{\left(\frac{m\pi}{h_{cav}}\right)^2 + \left(\frac{n\pi}{w_{cav}}\right)^2 + \left(\frac{l\pi}{l_{cav}}\right)^2} \tag{2.41}$$

For typical rectangular cavity designs the TE mode is used. With the condition that $w_{cav} < h_{cav} < l_{cav}$, the TE_{101} mode is the dominant, and lowest frequency mode, with next higher frequency modes following TE_{10l} until the wavelength of the next highest mode is less than the $m = 2$ mode.

Before we continue to describe both the electric and magnetic fields of the TE_{10l} mode, lets take a second to think about what this looks like physically. By virtue of the boundary conditions we have created, the field profile goes to zero at either end. For $l = 1$ these are the only two nodes, but for $l = 2$ we get an additional node at $l_{cav}/2$, and for $l = 3$ we get 2 additional nodes at $l_{cav}/3$ and $2l_{cav}/3$, or $n_{node} = l - 1$ excluding the two at either side. At these points the electric field goes to zero. We will see that the spacial location of these nodes have implications in the design of these cavities for real-world applications, as the rate of coupling is determined by the field strength at a given point. In a single mode system this is not an issue, but for multimodal systems, where coupling to *all* modes is important careful design considerations must be made, as we will see.

Knowing the mode frequency is important, and as we have seen, to first order, we can even estimate the electric field strength along the length of the cavity, to understand where the stored magnetic and electric energy lived we need a complete description of the fields. First, we can solve for E_{xy}. From the Helmholtz equations for a rectangular waveguide, we see that the term E_{xy} for the TE mode is:

$$E(x,y) = A_{xy} \sin \frac{m\pi x}{h} \cos \frac{n\pi y}{w} \tag{2.42}$$

Plugging this back into Eq. 2.36, and looking at the TE_{10l} mode we get:

$$\vec{E}_y = E_0 \sin \frac{\pi x}{h_{cav}} \sin \frac{\pi z}{l_{cav}} \hat{y} \tag{2.43}$$

Where we have replaced the generic scattering amplitude (and other coefficients) with E_0, and noted the vector in which this field is pointing, in this case, as we have defined our coordinates, in the $+\hat{y}$ direction. This is the magnetic field profile, but we also need to know the electric field profile as well. Using Maxwell's equations we can solve for the other field components. We note that for TE waves $E_z = 0$ and $B_z \neq 0$. Solving for the other field components with this condition in mind we get:

$$H_x = \frac{-jE_0}{Z_{TE}} \sin \frac{\pi x}{h_{cav}} \cos \frac{l\pi z}{l_{cav}} \hat{x} \tag{2.44a}$$

$$H_z = \frac{j\pi E_0}{k\eta h_{cav}} \cos \frac{\pi x}{h_{cav}} \sin \frac{l\pi z}{l_{cav}} \hat{z} \tag{2.44b}$$

From Eqs. 2.43 and 2.44 (a) and (b) we now have a complete description of *how* the electromagnetic energy is distributed in our cavity. As we can see, instead of discrete circuit elements like our initial model, both magnetic and electric energy is distributed according to the boundary conditions that we have imposed.

We will use the TE_{101} cavity in our exploration of how geometry effects the losses of our cavity system. The above is also just one example of a 3D resonator. In this thesis we will discuss modified TE_{101}, or tapered, rectangular waveguide resonators, along with cylindrical (and the similar elliptical), and coaxial resonators, however, any geometry with the correct boundary conditions can and will have resonant modes. It is because of this that special care has to be taken in simulating the complete electromagnetic system. The techniques for doing so will be discussed in the next chapter.

2.3.1 Cavity geometry optimization

In the previous sections we have outlined the primary sources of loss in our microwave resonators. We have also discussed a simple 3D resonator geometry and how the electric and magnetic field profiles are distributed in space. In this section we will quantify the effects these loss mechanisms have on the measured quality factor in these 3D resonators but discussing how the electromagnetic energy of such a cavity is distributed. This section will discuss the analytic methods for determining this ratio, before discussing how this is done in practice for the resonators discussed in later sections. The analytic derivations heavily barrow from refs. [144, 149].

To begin we can discuss the electromagnetic energy stored in any volume material is just given by:

$$U_{vol} = \int_V \varepsilon|E^2| + \mu|H^2|dV \tag{2.45}$$

To determine the participation of energy that participates in a sub-volume of the resonator we can just take the ratio of the energy in one volume over the total volume:

$$p = \frac{\int_V \varepsilon|E^2| + \mu|H^2|dV}{\int_{V_{tot}} \varepsilon|E^2| + \mu|H^2|dV} \tag{2.46}$$

Where we have called this *energy participation* ratio p. In practice the above integrals are difficult to compute, since you would have to know the values of the electric field profile both in the volume of the cavity, and in the participating material. We can, however, make a few simplifications. Given that our field is sinusoidally varying, we can look at the time averaged electric and magnetic stored energy independently:

$$W_e = \frac{1}{2}\int_V \varepsilon |\vec{E}|^2 dV \tag{2.47a}$$

$$W_m = \frac{1}{2}\int_V \mu |\vec{H}|^2 dV \tag{2.47b}$$

Where over one period $W_m = W_e$. Rewriting Eq. 2.46 as:

$$p_e = \frac{\int_V \varepsilon |\vec{E}|^2 dV}{\int_{V_{tot}} \varepsilon |\vec{E}|^2 dV} \tag{2.48a}$$

$$p_m = \frac{\int_V \mu |\vec{H}|^2 dV}{\int_{V_{tot}} \mu |\vec{H}|^2 dV} \tag{2.48b}$$

In the case of a cavity, the volume associated with the numerator of Eqs. 2.48 a and b is at the surface, with the thickness of the participating volume set by either the london penetration length, λ_L (or skin depth for normal-state conductors), or the dielectric thickness, t_{ox}, of the participating surface oxide.

We can now rewrite Eqs. 2.48 a and b in terms of these material specific parameters, but first, we can note that, for a superconductor, the tangental E-field is zero, meaning that only the perpendicular field is present. We can also note that the displacement field $\vec{D}$ in our dielectric is continuous with the field in the volume of our cavity at the dielectric boundary. Because of this we can state that the electric field in the dielectric surface should scale with the volumetric field as $E_{surf} = E_{vol}(\varepsilon_0/\varepsilon_{ox})$, where we have preemptively used

the vacuum permittivity ε_0 to denote the dielectric constant of our mode volume, and the permittivity of our oxide ε_{ox} for our participating surface. These two approximations give us the approximate form for p_e:

$$p_e \approx \frac{t_{ox}}{\varepsilon_{r,ox}} \left(\frac{\int_S |\vec{E}_{vol}|^2 dA}{\int_{V_{tot}} |\vec{E}_{vol}|^2 dV} \right) = \frac{t_{ox}}{\varepsilon_{r,ox}} s_e \tag{2.49}$$

Where we have assume that the oxide permittivity is isotropic throughout the amorphous oxide. The ratio of the integrated surface and volume electric field magnitude in parenthesis is often called the *surface participation ratio*, which will be denoted as s_e for shorthand. These integrals can be difficult to solve analytically except for cases in which cavity symmetry can be used to easily compute the surface and volume integrals. In reality, and for the rest of this thesis, s_e, and the corresponding magnetic surface participation s_m, will be computed numerically using FEM software. The equivalency between these two methods will be demonstrated for a simple geometry in the next section.

Now that we have computed s_e, we can also compute s_m. Instead of doing what we did for s_e by starting with eq. 2.48, the more intuitive way to derive s_m is to begin by looking at the amount of energy that is *dissipated* in the conducting layer. For a cavity the pertinent value is the power dissipated per unit area at the surface, or dP_{res}/ds. This power is just:

$$\frac{dP_{res}}{ds} = \frac{1}{2} R_s |\vec{H}|^2 \tag{2.50}$$

We can easily solve for P_{res}, or the resistive power dissipation by integrating over ds to get:

$$P_{res} = \frac{1}{2} R_s \int_S |\vec{H}|^2 dA \tag{2.51}$$

Here P_{res} is the power loss in our system. We can then use the definition of Q_0 from eq. 2.3 and eq. 2.47 (b) while remembering that over one period $W_m = W_e$ to get:

$$Q_{res} = \frac{2\omega_0 W_m}{P_{res}} = \frac{\omega_0 \mu}{R_s}\left(\frac{\int_V |\vec{H}|^2 dV}{\int_S |\vec{H}|^2 dA}\right) = \frac{\omega_0 \mu}{R_s}\frac{1}{p_m} \tag{2.52}$$

We can see from eq. 2.52 that we have once again introduced a surface participation factor, however unlike in eq. 2.49, the coefficient is $1/s_m$. To get a unitless participation p_m we can note that the volume associated with the interacting magnetic field is set by the penetration length scale of our current. As previously mentioned, for normal state conductors this scale is set by δ, or the skin depth, while for our superconductor this length scale is determined by the London penetration length λ_L. We can substitute this length into eq. 2.52 to get:

$$Q_{res} = \left(\frac{\omega_0 \mu \lambda_L}{R_s}\right)\left(\frac{1}{\lambda_L p_m}\right) \tag{2.53}$$

And subsequently:

$$p_m = \frac{\lambda_L \int_S |\vec{H}|^2 dA}{\int_V |\vec{H}|^2 dV} \tag{2.54}$$

The deliberate choice of writing eq. 2.53 in this way is to highlight the relationship that p_m has with the concepts that have been previously discussed. The denominator of the left parenthesis is simply the reactance X_s of our superconducting impedance Z_s as defined in eq. 2.22 for our two-fluid model. With that in mind we can also see that the left half of eq. 2.53 is really equivalent to Q_{mag} as defined in eq. 2.26 once the BCS conductivities are included. What this shows is that Q_{res} as defined above is really just Q_{mag} as defined in our two-fluid model with a proportionality constant of p_m^{-1}, which itself is a geometric dependent value.

Similarly the Q associated with our dielectric loss tangent for our cavity can be expressed as:

$$Q_{diel} = \frac{\varepsilon_{r,ox}}{t_{ox} S_e}\frac{1}{\tan\delta} \tag{2.55}$$

We can now see the relationship between the two-level system filling factor F in eq. 2.35 with 2.55 above, with $F = p_e t_{ox}/\varepsilon_{r,ox}$ for our resonator.

The above shows that, in addition to the quality of the material we use, both in terms of its conductive and dielectric properties, geometry also plays a big part in determining the quality factor of our resonant cavities. From eqs. 2.53 and 2.55 we see that *minimizing* the amount of field at our surface is critical to good performance. Inversely, in systems where low Q is desirable, like cavity attenuators, then *maximizing* the surface participation s_e and s_m may be desirable. In the next section we will explore these ideas further in discussing how in practice one can compute the geometric values from simulation, and how to use them to assist in extracting material parameters.

2.3.2 Seam loss

The above sections outlined the most relevant material and geometric factors in the minimization of loss as they pertain to the devices discussed in later sections of this thesis, however for completeness, they are not the *only* sources of loss, either.

In our discussion of cavity geometry we have assumed contiguous superconducting surfaces enclosing our resonant cavity. This, however, is only a theoretical treatment of an ideal cavity. In reality, to make the enclosed volume, access has to be created to mill out the material. The most common way this is done is by splitting the cavity volume in half, creating a seam. In later sections we will discuss another technique which avoids this seam (and its associated loss) in exchange for the use of holes which are only evanescently coupled to the cavity mode. The derivation of this seam loss is taken from Refs. [31, 148].

The loss at the seam is due to a combination of resistive material (such as a thin oxide layer) or caused by surface imperfections that reduce the overall contact area. Either way, the loss can simply be taken as a boundary with total finite conductance G, over a total contact surface. The path length of that surface allows us to get a conductance per unit

length $g_{\oslash} = G/l$ which gives us:

$$dG = g_{\oslash} dl \tag{2.56}$$

The power dissipated along a finite length $d\vec{l}$ is given by an ohmic relation:

$$dP_s = \frac{1}{2}\frac{(dI)^2}{dG} \tag{2.57}$$

The finite current element dI can be written in terms of the surface current density $I = |\vec{J}_{\perp} \times \vec{l}|$. This current density, which is perpendicular to the contacting surfaces, is a result of the normal component of the parallel magnetic field profile, giving:

$$\vec{J}_{\perp} = \hat{n} \times \vec{H}_{\parallel} \tag{2.58}$$

Substituting in this definition and Eq. 2.56 into Eq. 2.57 gives us:

$$dP_s = \frac{1}{2}\frac{|\vec{J}_{\perp} \times \hat{l}|^2 dl}{g_{\oslash}} = \frac{1}{2}\frac{|\vec{H}_{\parallel}|^2 dl}{g_{\oslash}} \tag{2.59}$$

Integrating both sides of Eq. 2.59 gives us:

$$P_s = \frac{1}{2g_{\oslash}} \int_l |\vec{H}_{\parallel}|^2 dl \tag{2.60}$$

The above represents the dissipated loss. Using Eq. 2.3 and the time averaged magnetic energy from Eq. 2.47 (b), we arrive at:

$$Q_{seam} = \omega \mu g_{\oslash} \frac{\int_V |\vec{H}|^2 dV}{\int_l |\vec{H}_{\parallel}|^2 dl} \tag{2.61}$$

As before, Eq. 2.61 can be further separated into a loss factor $g_{\oslash}$ and a geometric participation factor:

$$Q_{seam} = \frac{g_{\oslash}}{y_{seam}} \tag{2.62}$$

Where we have defined the geometric term in Eq. 2.61 as:

$$y_{seam} = \frac{\int_l |\vec{H}_{\|}|^2 dl}{\omega\mu \int_V |\vec{H}|^2 dV} \tag{2.63}$$

Which is often called the seam admittance.

As stated, the effect of seam loss is only a problem in cavities that require construction from multiple parts. With the above in hand, however, it is possible to design the cavity in a manner that minimizes y_{seam} by choosing a line that runs parallel with $\vec{H}_{\|}$ to minimize the surface current density running across the seam.

2.3.3 Radiative and evanescent loss

In later sections we will introduce the flute cavity design, which is intrinsically seamless. It does, however, require that overlapping holes be drilled into opposing sides of the cavity. These holes act like waveguides, coupling electromagnetic energy into the external environment. Naively, one may expect that a "holy" cavity would be an awful one, but in reality this loss can be suppressed. To see how, lets take a single hole connected to a resonator with a resonant frequency f_0. This problem has even more general implications, especially with concern to coupling to a cavity mode via an external coupler, which often requires the input of a coaxial line (TEM waveguide mode) coupling to the resonator via propagating TM or TE modes.

To begin the discussion of how a hole introduces loss into the system, we can begin by looking at the basic principles of the propagation of electromagnetic energy through circular waveguides [2]. Circular waveguides support both TE and TM modes. To begin, a TE_{nm}

2. For a more thorough description of waveguides see Refs. [144, 50]

mode for a circular waveguide has a propagation constant β_{nm} given as:

$$\beta_{nm} = \sqrt{k^2 - \left(\frac{p'_{nm}}{r}\right)^2} \tag{2.64}$$

Where p'_{nm} is the $m-th$ root of $J'_n(x)$, or the derivative of a Bessel function of the first kind. Here $k = \omega\sqrt{\mu\epsilon}$, where ω is a waveguide mode. We see that for values of $\omega < \omega_c = p'_{nm}/r\sqrt{\mu\epsilon}$, then the propagation constant is imaginary. The value for ω which corresponds to this is given above as ω_c, or the *waveguide cutoff frequency*. To see the effect that this has on our system we can look at the effect β_{nm} has on the azimuthal and radial electric fields of our TE_{11} waveguide mode:

$$\vec{E}_\rho = \frac{-i\omega\mu}{k_c^2\rho}(A\cos\phi + B\sin(\phi))J_1(k_c\rho)e^{i\beta_{11}z}\hat{\rho} \tag{2.65a}$$

$$\vec{E}_\phi = \frac{i\omega\mu}{k_c}(A\sin\phi + B\cos\phi)J'_1(k_c\rho)e^{i\beta_{11}z}\hat{\phi} \tag{2.65b}$$

And magnetic components:

$$\vec{H}_\rho = \frac{i\beta_{11}}{k_c}(A\sin\phi + B\cos\phi)J'_1(k_c\rho)e^{i\beta_{11}z}\hat{\rho} \tag{2.66a}$$

$$\vec{H}_\phi = \frac{i\beta_{11}}{k_c^2\rho}(A\cos\phi - B\sin\phi)J_1(k_c\rho)e^{i\beta_{11}z}\hat{\phi} \tag{2.66b}$$

$$\vec{H}_z = A\sin\phi J_1(k_c\rho)e^{i\beta_{11}z}\hat{\phi} \tag{2.66c}$$

Both $\hat{\rho}$ and $\hat{\phi}$ components of $\vec{E}$ and $\vec{H}$ have independent amplitudes A and B, with either term being a valid solution due to the azimuthal symmetry of the cylindrical waveguide. In

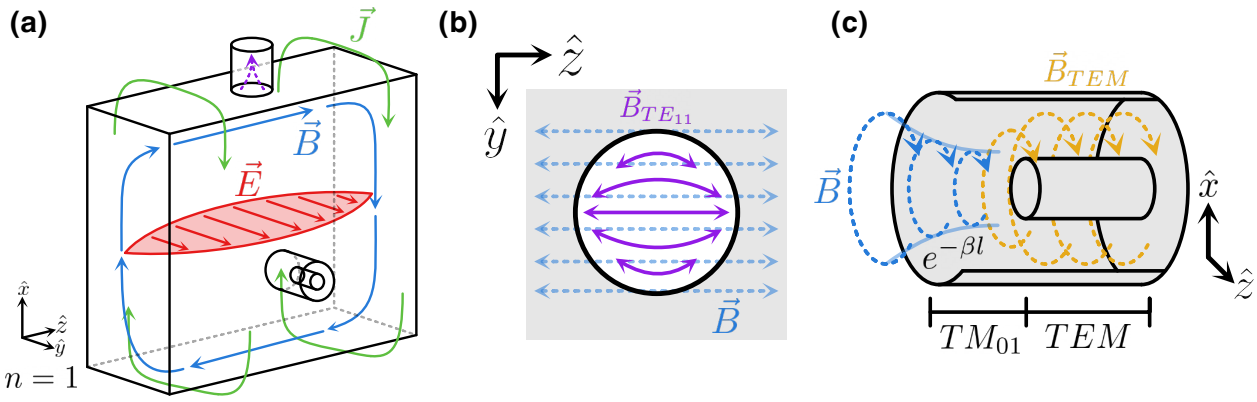

Figure 2.6: **Field profiles for rectangular cavities and evanescent coupling holes**: (a) shows the magnetic and electric field vectors, and the surface current density, for the TE_{10n} modes of a rectangular waveguide cavity. (b) shows the magnetic dipole coupling between the circular waveguide mode of an evanescant hole and the TE_{10n} modes of the cavity. (c) shows the electric dipole coupling between the input/output coupler pins. This mode decays and couples into the TEM mode of a coaxial pin, with the $Q_{ext} \propto l$

practice we set either $A = 0$ or $B = 0$ depending on the nature of the excitation of the waveguide mode. We also see that, for frequencies below the cutoff frequency, when the propagation constant is imaginary, the electric field profile decays exponentially. We can go one step further and see what the effect of this is on the energy dissipation of a cavity mode with $\omega_0 < \omega_c$ for a given waveguide mode. For an aperture, the coupling of the waveguide to the cavity mode was described by Bethe as a dipole coupling [24]. For the flute cavity discussed later, we can determine that the waveguide mode excited is the aforementioned TE_{11} through an effective dipole coupling between resonant mode's $\vec{E}$ field, as shown in Fig. 2.6.

To determine the Q_{ext} from this waveguide coupling, we want to determine what the dissipated power is for a circular waveguide terminated by a shunt resistance R_s. We know that the power through this resistance is given by Eq. 2.51. For a load at the end of the hole we can write this as:

$$P_l = \frac{R_s}{2} \int_S |\vec{H}_\rho|^2 + |\vec{H}_\phi|^2 ds \tag{2.67}$$

Since $\beta_{11} = i\,\mathrm{Re}(\beta_{11})$, we can rewrite Eqs. 2.66 accordingly and substitute it into the equation above. This gives a value for P_l at the end of the waveguide of:

$$P_l = \frac{R_s \beta_{11}^2 |A|^2}{4k_c^4}({p'_{11}}^2 - 1)J_1^2(k_c r_0)e^{-\beta_{11} l_{hole}} \tag{2.68}$$

Where we have substituted z with the l_{hole}, or our hole length. From Eq. 2.68 we see that, as the length increases, the power loss exponentially *decreases*. To get the effect that this has on our total quality factor, assuming that $Q_0 = Q_{ext}$, we can can substitute P_l into Eq. 2.52 in place of P_{res} to get:

$$Q_{ext} = \frac{2\omega_0 W_m}{P_l} = \frac{2\omega_0 e^{2\beta_{11} l_{hole}}}{R_s}\left(\frac{\epsilon_0 hwl}{4\beta_{11}^2 r_0^4}\right)\left(\frac{E_0^2}{|A|^2}\right)\left(\frac{{p'_{11}}^4}{({p'_{11}}^2 - 1)J_1^2(p'_{11})}\right) \tag{2.69}$$

The two main takeaways from Eq. 2.69 are that the participation at the end of the hole falls exponentially with its length, meaning that Q_{ext} *increases* exponentially with length. The second factor has to do with the resistance R_s. In the definition we have used, this value for R_s is effectively acting like a series resistance in our circuit, however, in practice, the path acts like a parallel current shunt, meaning that as the effective resistance of our end hole *increases*, then the effective series resistance R_s *decreases*. This means that ideally the effective impedance at the hole boundary would be infinite, meaning that, from the perspective of the cavity mode, 100% of the power incident on the port would be reflected. This also shows that, from the perspective of the evanescent mode, a termination with an infinite impedance is equivalent to a hole with infinite length. Both of these properties are demonstrated in Fig. 2.7.

The other factor in Eq. 2.68 is the fraction of the electromagnetic energy transmitted from the mode into the port, as given by the fraction $E_0^2/|A|^2$. For our hole located at $\pm x$ we can derive the strength of the coupling to the magnetic field, following Ch.6 from Ref. [50]

(or Ch.4 from [41]), which has a horizontal electric dipole of:

$$\vec{m} = -\frac{4}{3} r_0^3 \vec{H}_t = -\alpha_m |\vec{H}_0| \delta^3(x - x_0)\hat{x} \tag{2.70}$$

Where $\alpha_m = 4r_0^3/3$ and $|\vec{H}_0|$ is the amplitude of the cavity field at the wall on resonance.

For a TE_{101} mode in a rectangular cavity, the magnetic field components are:

$$\vec{H}_z = \frac{-iE_o}{Z_{TE}} \sin\frac{\pi h}{h_{cav}} \cos\frac{\pi l}{l_{cav}} \hat{z} \tag{2.71a}$$

$$\vec{H}_x = \frac{i\pi E_o}{k\eta h_{cav}} \cos\frac{\pi h}{h_{cav}} \sin\frac{\pi z}{l_{cav}} \hat{x} \tag{2.71b}$$

$+x = a$ and $-x = a$ then $\vec{H}_z = 0$, so only the magnetic field in the $\hat{x}$ component will couple into our waveguide. This gives us the magnetization density from Eq. 2.70:

$$\vec{M} = \alpha_m \frac{i\omega\mu_0 \pi E_0}{k_0 \eta h_{cav}} \sin\frac{\pi z}{l_{cav}} \delta^3(x - x_0)\hat{x} \tag{2.72}$$

To determine the effect of the apertures magnetic dipole moment on the cavity field, an expansion must be made:

$$\vec{H} = \sum_n h_n \vec{H}_n \tag{2.73}$$

Determining the expansion coefficient h_n can be tricky, but Collins [50] showed that, for the resonant case, for h_n corresponding to ω_n, detuned from the cavity resonant frequency ω_0, their components would be suppressed, leading to the simplification of Eq. 2.73, and leading to an approximate form for h_n:

$$h_n \approx -i\epsilon_0 \omega \frac{\int_V \vec{J}_m \cdot \vec{H}_n dV}{k_n^2 - k_0^2\left(1 + \dfrac{1-j}{Q_0}\right)} \tag{2.74}$$

In the case of a coupled drive port, the value for h_n would give the strength of cavity drive on $\vec{H}_n$ corresponding to our resonant mode. Subsequently, the Q_0 value in Eq. 2.74 would be dependent on the external coupling Q_{ext}, since this value of Q_0 is really the loaded Q of the system. In our case, we have set the boundaries of our simulated cavity to be perfect conductors, meaning that $Q_0 = Q_{ext}$. In a real-world cavity, with holes that have a strong evanescent decay, $Q_{ext} \gg Q_0$, meaning that this value will be independent. Next, we must compute the numerator, remembering that:

$$\vec{J}_m = i\omega\mu\vec{M} \tag{2.75}$$

We can substitute in Eq. 2.72 into Eq. 2.75 and solve for the numerator of Eq. 2.74. To begin, we can not that the TE_{11} waveguide mode excited has a magnetic field component:

$$\vec{H}_{11} = H_\phi\hat{\phi} + H_\rho\hat{\rho} \tag{2.76}$$

We can then take the dot-product of Eq. 2.70 and Eq. 2.76, returning:

$$\int_V \vec{J}_m \cdot \vec{H}_{11} = \int_V [H_\rho J_m \cos\phi - H_\phi J_m \sin\phi] dV \tag{2.77}$$

Plugging in Eq. 2.72 and Eqs. 2.66 (a) and (b) into Eq. 2.77 and computing the integral. This integral involves the evaluation of first kind Bessel functions (and their derivative) from Eq. 2.66, which give can be difficult to evaluate in closed form, however since we only have to evaluate to p'_{11}, or the first root of the J_1 Bessel function, we can treat that as an integration constant C_ρ with units of $1/m^2$. The final result of this is the following:

$$|h_0| = \varepsilon_0\omega\left(\frac{4}{3}r_0^5|E_0|\right)\left(\frac{C_\rho\beta_{11}^2}{p_{11}'^2\eta k_{\text{cav}}h_{\text{cav}}}\right)\left(\frac{Q_{ext}}{\sqrt{2}k_0^2}\right) \tag{2.78}$$

Where $|h_0|$ in units of A/m. If we divide both sides by E_0 we arrive at a ratio of the driven field to the cavity field $|h_0|/|E_0|$, which is analogous to the ratio of $|E_0|/|A|$ in Eq. 2.69,

where $|A| = |h_0|$. With this in mind, we can plug into Eq. 2.78, into Eq. 2.69:

$$Q_{ext} = C_{|h|} \frac{e^{2\beta_{11} l_{hole}}}{\beta_{11}^6 r_0^{14} Q_{ext}^2} \tag{2.79}$$

Where we have collected additional constant terms into the coefficient $C_{|h|}$. We have left terms that are dependent on r_0 above, with $\beta_1 1 = \sqrt{k_0^2 - k_{11}^2}$ where $k_{11} = p'_{11}/r_0$ is the wavenumber for the TE_{11} mode of the circular waveguide, and $p'_{11} \sim 1.841$ is the first root of the derivative of a first kind Bessel function. Solving Eq. 2.79 for Q_{ext} gives us the form:

$$Q_{ext} \propto \frac{e^{2\beta_{11} l_{hole}}}{\sqrt[3]{\beta_{11}^6 r_0^{14}}} \tag{2.80}$$

It shows that for the evanescent holes located as shown, in Fig. 2.7 that not just does the field decay exponentially with length, but also in the radius of the hole as well, with an additional factor of $r_0^{8/3}$. This relationship is shown in Fig. 2.7 (a) in the form of an analytical fit, using the simulated cavity parameters with only a fitting parameter for $C_{|h|}$ provided, which acts as an offset in the figure shown. For the same offset parameter Eq. 2.80 shows very good agreement with the simulated data over a number of values of l and r_0, with deviation occurring for $Q_{ext} \gtrsim 1 \times 10^{14}$ due to limits in the numerical precision of the finite element simulation.

The above treatment of evanescent holes has also been done for coupling the cavity mode to a TEM mode of a coaxial coupler evanescently through a sub-cutoff circular waveguides TM_{01} mode by Brecht and Reagor [31, 149]. In that analysis it was electric dipole coupling with the transverse electric field of the cavity that provided the coupling to the antenna. Interestingly, for that particular coupling, the mode exhibited an $Q_{ext} \propto r_0^{-6}$ relation. This shows that, even for evanescent holes, considerations should be made for the specific orientation if minimizing loss through such holes is a priority. As we will see, for the monolithic flute cavity, the suppression of radiative leakage is sufficiently minimized.

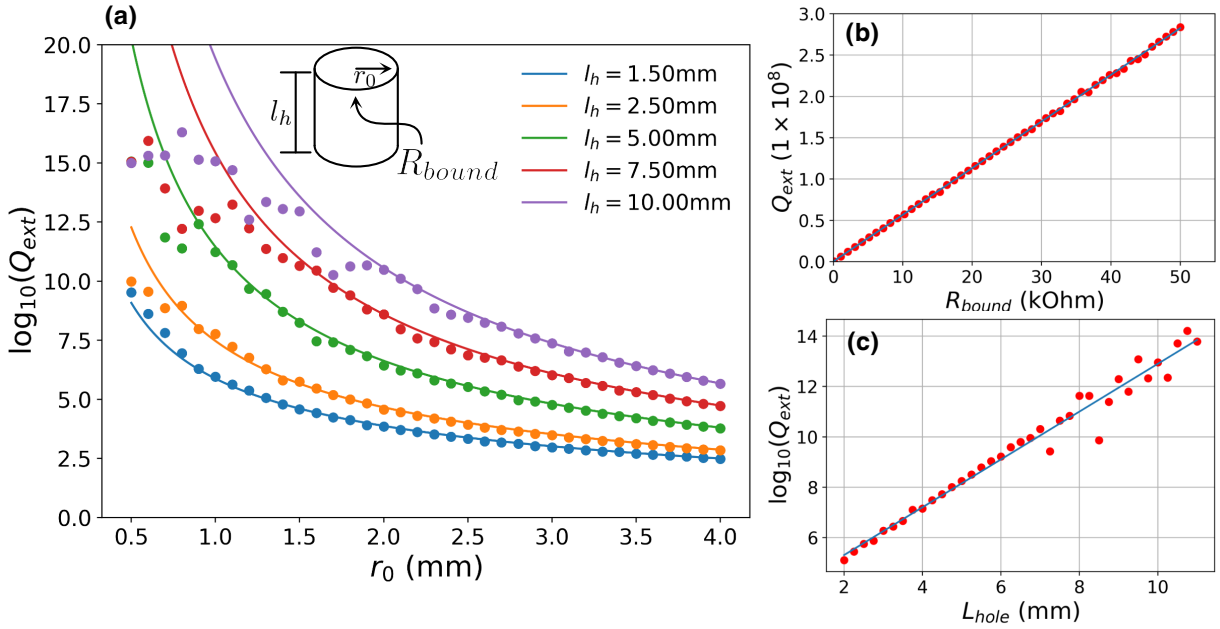

Figure 2.7: **Scaling of external Q for evanescent hole**: (a) shows the effect of changing the radius of the hole, r_0, for various hole lengths, showing the non-trivial scaling of Q_{ext} with hole radius, as described in Eq. 2.80. The data is derived from parametric simulations of an $f_0 \sim 7.5$GHz rectangular waveguide cavity. (b) shows the change in Q_{ext} vs boundary impedance, with a linearly increasing Q with resistance, showing the fact that the hole acts as a parallel current path. (c) shows Q_{ext} vs evanescent length l_{hole} for a fixed radius, showing the exponential increase in Q_{ext} vs length, as described in Eq. 2.67. In this simulation $Q_0 = Q_{ext}$ since all other boundaries are set to be perfect conductance.

CHAPTER 3

3D CAVITY SIMULATION AND DESIGN

3.1 Introduction

In the previous chapter we discussed the basic principles of electromagnetic resonators, with the majority of the time dedicated to loss and geometric considerations for 3D resonant cavities. While we described a simple example in the form of a TE_{101} cavity, we also saw that for analytically describing even a simple coupling of a *single* hole to that cavity was a tedious undertaking. For many, sticking to simple geometries with easy to evaulate closed form solutions to their electromagnetic field is good enough, but for complex geometries such as the multimode cavity, which has no compact analytic description for its resonant properties, we must utilize simulation tools. This is compounded when we introduce additional complications, such as non-linear resonant circuits, or couplings between adjacent cavities.

In this chapter we will discuss some basic simulation principles, and show a myriad of techniques for solving almost all of the aforementioned surface participation ratios and more. More importantly, this chapter will be an instruction manual for how to use two very powerful tools for cavity design (PyInventor) and simulation (PyHFSS, PyEPR)[122], and discuss the tool chain for simulating a few simple geometries. In addition to this we will introduce some methodologies for simulating the effects of adding a non-linear resonator (in the form of a qubit), into our system. This won't be an exhaustive overview of all these two pieces of software have to offer, but along with the complete demonstration code [1], should be enough to begin designing and simulating 3D cavities.

1. You can find the packages and code at github.com/AndrewOriani for complete examples of both PyInventor and PyEPR demonstrations

3.2 PyInventor

In 2D superconducting circuits, a design can be made with almost any vector graphics program, as long as the final filetype can be interpreted by the fabrication tools used to pattern the design. In reality, for complicated parametric designs, code is used to generate the individual vector elements and layers before exporting into a compatible format. In the past, custom libraries were used for this purpose[2]. More recently, standalone SC-circuit design tools (KQCircuit)[90] and full-stack design and simulation tools (Qisket Metal)[123] have given scientist powerful open-source design tools for 2D quantum circuits. This makes sense given the increasing complexity and level of integration that modern 2D designs require for state of the art quantum experiments. Normally, 3D cavity geometries are quite simple, made up of shapes that require little or no parameterization, and could be made simply by normal hand-operated machine tools. As 3D quantum experiments become more complicated, however, similarly powerful design tools will be needed.

3D designs, however, are much more difficult to implement directly in code given the additional spatial relations that must be considered. While there are a few open sourced Python packages that exist for generating .STEP, .STP, and .IGS files, their featureset, documentation, and flexibility are limited. Standalone software, like AutoDesk Inventor and Dessault Systems' Solidworks, have become industry standards for 3D parametric design, with 30+ years of continuous development. These standalone software however have little in the way of automation built-in, and lack the integration with existing simulation tools like Ansys HFSS (microwave simulation) and Maxwell (Electromechanical) or Comsol (Multiphysics) out of the box. Another option is to use the design tools featured in these simulation software packages. This, however, is of limited utility, since these packages lack the features of the aforementioned 3D parametric design software, but also their designs cannot be exported

2. For the specific packaged used by the SchusterLab, see: https://github.com/SchusterLab/MaskMaker and https://github.com/SchusterLab/maskLib

into files that those software can interpret (although they *do* support importation of most standard 3D design files). With this in mind, what is really needed is a code package that can automatically generate designs in one piece of software, before exporting that design to the desired simulation software. The second part of the equation already mostly[3] exists in the form of the excellent PyEPR Ansys API [122]. The first part, concerning the automated design, is the discussion of the following section.

3.2.1 PyInventor: Fundamentals and setup

This section will discuss the fundamentals of the Autodesk Inventor API written to solve the aforementioned problem. PyInventor[4] is a open-source python package which is capable of generating individual part files (which contain the full parametric design) and exporting them into filetypes that can be used to generate toolpaths for automated computer numerical control (CNC) machining, or into existing simulation software. `PyInventor` works by interfacing with the Autodesk Inventor SDK, which is natively written using Windows Visual Basic for Applications (VBA). Because of this, Inventor (and `PyInventor`) will *not* work on MacOS. To demonstrate the use of PyInventor we will look at a simple example first in the form of the $\lambda/4$ coaxial cavity.

3. I say *mostly* because PyEPR was never intended to be used solely as an Ansys API. Because of this, additional code was written to accomodate the 3D design, and will be discussed in a later section

4. https://github.com/AndrewOriani/PyInventor

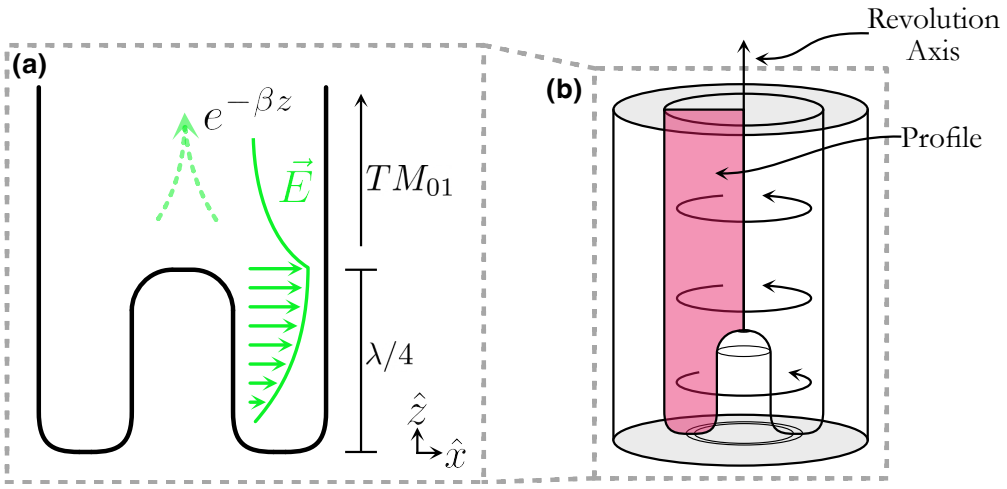

Figure 3.1: **Outline of coaxial cavity design and pyInventor steps for part creation**: (a) shows a cross-section of the $\lambda/4$ coaxial design. This cavity is effectively a coaxial line that has been shorted at one end. The fundamental mode is a TEM mode that has a fundamental frequency set by the pin length. The top section is a circular waveguide that has a cutoff frequency such that $f_c >> f_{\lambda/4}$. (b) shows a schematic of the process in which we generate the part in PyInventor. First we select a plane that bisects the geometry $(\hat{x} - \hat{z})$ and create a sketch that represents half of the geometry. We choose a perpendicular plane $(\hat{x} - \hat{y})$ to create the body of the cavity. Finally, using the `ipart.revolve` in addition with the `operation='subtract'` flag we are able to generate the interior geometry of our cavity

To begin, after importing the necessary PyInventor classes using `from PyInventor import ...`, we will need some basic initial setup for our part:

```
1  com_obj().close_all_parts()
2  units='imperial'
3  overwrite=True
4  fname='coax.ipt'
5  path='C:\\Users\\Public\\Documents\\Demos'
6  part=iPart(path=path, prefix=fname, units=units, overwrite=overwrite)
```

Most of the above snippet is self-explanatory (unit type, filename, path to save etc.), but a few points to highlight. The first is `com_obj().close_all_parts()` on line 1, which will close any opened part files in your Inventor application (including unsaved parts), and the `iPart` class on line 6, which contains all of the part specific functionality for Inventor.

Next we need to define a *work plane* to attach our sketch to. The work plane is the plane

in which we will *sketch* our design onto before conducting 3D operations. For the coaxial cavity, we will start with a sketch that defines the outline of our cavity cross-section, as shown in Fig. 3.1 (b). In addition to the work plane for our cavity profile (`revolve_wp`), we also define a stock workplane (`stock_wp`) in the $\hat{x} - \hat{y}$ plane that will be used for the stock profile:

```
1       revolve_wp=part.add_workplane(plane='xz')
2       stock_wp=part.add_workplane(plane='xy', offset=2)
3       stock_sketch=part.new_sketch(stock_wp)
4       revolve_sketch=part.new_sketch(revolve_wp)
```

Following the definition of our workplanes on lines 1 and 2, we define our two primary sketch objects associated with each workplane on lines 3 and 4. These sketch objects each act like their own independent 2D vectors, but unlike a standard vector they maintain associativity that is a hallmark of parametric CAD.

We can now create our profiles. The process of doing this borrows heavily from the `MaskMaker` library used for DXF creation for photolithography masks to create 2D superconducting circuits. First a `structure` object is created. This is a local dictionary object that contains both the type of vector object one is creating (line, arc, spline, etc.) and the associated points. This differs slightly from DXF software, in that the vector object is associated with a specific sketch object type inside of Inventor, each with their own associativity and constraint rules:

The above code starts by defining some of the pertinent cavity design parameters, which have been hardcoded here for clarity. Line 7 introduces the aforementioned `structure` class, which takes the part object and associated sketch[5]. In addition to those associativities, `structure` also takes the start point, which in our case is the top center corner of our stock

5. A newer version has been written that solves this clumsy inheritance, but since the majority of existing design code is built around this version of `PyInventor`, we will continue using this structure.

```
1  r_cav=0.3 #inches
2  h_cav=1.85 #inches
3  bot_rad=0.094 #inches
4  pin_height=0.231 #inches
5  pin_rad=0.0625 #inches
6  stock_height=2.0 #inches
7
8  s=structure(part, revolve_sketch, start=(0,stock_height))
9
10 s.add_line(r_cav, 180)
11 s.add_line(h_cav, 270)
12 s.add_line_arc(start_angle=180, stop_angle=0, radius=bot_rad, flip_dir=True,
   ↪  rotation=0)
13 s.add_line(pin_height, 90)
14 s.add_line_arc(start_angle=0, stop_angle=270, radius=pin_rad, flip_dir=False,
   ↪  rotation=180)
15 s.add_line(r_cav-2*bot_rad-pin_rad, 0)
16
17 rev_path=s.draw_path(close_path=True)
```

since we have set up this design to be radially symmetric about the $\hat{z}$-axis

Lines $9-15$ are our sketch operations. Each operation adds a sketch object dictionary entry to the `structure` object. At this point the structure dictionary (`s.obj_dict()`) is a locally saved variable. To tell Inventor to actually *draw* this set of line objects, we use `draw_path`, which compiles each entry of the dictionary object into a set of connected Inventor line objects.

Two important things to note with `draw_path` are the `close_path` flag, and the output. The `close_path` flag tells inventor that the first and last points should be connected. If the two points do not have the same coordinate, then the structure object will simply add a line that closes the path. This is important, as if the path is open, it will not be possible to do 3D operations. At a later stage. The second aspect worth noting is the returned object, which in this case is called `rev_path`. This is an Inventor API object known as an object collection, and allows us to treat this structure as a single entity that can be operated on.

We will see how this works at a later stage.

First, though, we need to define our stock. Here we have a cylindrical profile, which, when projected onto the $\hat{x} - \hat{y}$ plane is a simple circle:

```
1  stock_bot_circle=part.sketch_circle(stock_sketch, center=(0,0), radius=.5)
```

Here we see a different way of creating a sketch object. In the case of the cavity profile, where the shape was irregular, we had to use the `structure` object to define the contour, however in the case of the stock profile, which is a circle, we can directly tell inventor to sketch the desired shape, in this case using `sketch_circle`[6], with a center point (`center`) and radius as inputs.

Up to this point we have only worked in two dimensions. To create our 3D cavity we need to conduct 3D operations *with* these 2D geometries. Here the order of operations is important, and reflect the order in which a part would actually be made. The first step is to generate the *stock* of the cavity, i.e. the outer form of the part. The second step is to then *remove* the desired material to form the interior cavity geometry, which is analogous to physically machining out the material:

```
1  stock_solid=part.extrude(stock_sketch, thickness=stock_height,
   ↪  obj_collection=stock_bot_circle, direction='negative', operation='join')
2
3  cav_vol=part.revolve_full(revolve_sketch, axis='z', obj_collection=rev_path,
   ↪  operation='cut')
```

Line 1 shows the creation of the stock geometry, using a simple extrusion. Note that the extrusion always moves perpendicular to the plane that the sketch is located in. In addition to this, the extrusion also allows for Boolean operations in addition to `'join'`. The

6. In addition to circles, `PyInventor` also has rectangles (center and two-point), and slots for closed regular geometries

`direction` flag is with reference to the coordinate system directions. We also see the use of the `obj_collection` defined previously that represents our stock profile.

For the cavity volume, the process is a bit different. Taking advantage of the radial symmetry of the profile we had generated (`rev_path`), we can simply rotate that profile about the $\hat{z}$-axis. Since we are removing material, we use the `'cut'` operation. In this case, the revolution axis was a cardinal axis for our design, however, can be designated by a line or work axis. At this point the design specified in Fig. 3.1 has been realized. Additional steps, such as using `part.save()` to save the `.ipt` file and also saving an `.stp` copy for export using `part.save_copy_as(copy_name='coaxial cavity.stp')`.

We have shown the basics of how `PyInventor` can be used to generate a simple cavity design. In reality, the above cavity is not the best demonstration of what `PyInventor` can do from a parametric design standpoint. Tapered cavity geometries, like the multimodal flute cavity discussed in a later section, or cavity lattices like that used in Refs. [131, 132] are perhaps more compelling uses. The code for creating a similar lattice to that used in the aforementioned references can be found at the `PyInventor` GitHub[7] along with many other examples.

3.3 HFSS and Electrodynamic Simulations

The previous section demonstrated how we can automate the cavity design process, but once a design has been made, how do we determine the pertinent properties, i.e. resonant frequency and participation ratios? As we saw in Ch.2, it is possible to calculate these values analytically, but in practice most geometries deviate from the ideal. In this section we will discuss how we extract these important values from simulation, and use it as an introduction to the techniques that will be used in later sections for simulating the interaction of these cavities with our qubits.

7. https://github.com/AndrewOriani/PyInventor/blob/main/_Tutorial_Notebooks/PyInventor%20Demo.ipynb

3.3.1 HFSS and PyHFSS

:

Ansys HFSS (High-Frequency Structure Simulator) is an electrodynamics FEM software package used for simulating microwave resonators. Along with Q3D, an electrostatics modeller, and Maxwell, for electromechanical systems, HFSS is one of the primary tools used in simulating resonators along with coupled microwave systems. As with Inventor, HFSS also has a robust API which is natively written in Python. `PyHFSS` is a sub-package of the larger PyEPR package [122] developed by Zaki Leghtas and Zlatko Minev, with the HFSS interface being a continuation of PyHFSS originally written by Phil Reinhold[8]. In later sections we will discuss the use of `PyEPR` and its static circuit solver, but in this section we will just look at solving for the electromagnetic properties of our cavity that we created in the previous section. One note is that the version of code that will be presented in part in this section uses a fork of PyEPR that can be found on my Github[9].

Most of the simulation code concerns definition of variables, geometry and boundary condition settings, and finally simulation settings. I won't go in depth into all of those parts here, but would like to first point out a few important things when setting up the simulation.

After instantiating HFSS Desktop, creating a project, and adding a design file to that project, the first important step, and one unique to my branch of code (as of writing this thesis) is the 3D importation function. In the previous section we exported the cavity design as an `.stp` file, and the first step in our simulation is *importing* of that design as follows:

Normally we would set boundary conditions by selecting the surfaces and assigning a material property to it, however since we wish to study mode frequency, field participation, and the effects of evanescent loss, we can set these boundaries to perfect conductance. A fast way of doing this is to do a Boolean subtract on the interior cavity volume by first creating

8. https://github.com/PhilReinhold/pyHFSS

9. https://github.com/AndrewOriani/Automated-RF-Design-Demo-MASTER/tree/main/Updated%20pyEPR%20Files/pyEPR

```
path=path+'\\'+'coaxial cavity.stp'
EM_design.make_active()
model=HFSS.HfssModeler(EM_design)
model.import_3D_obj(path)
```

a stock solid form. The following code describes this process:

```
Stock_L=EM_design.set_variable('Stock_L', '%.3fin'%(stock_L))
Stock_W=EM_design.set_variable('Stock_W', '%.3fin'%(stock_W))
Stock_H=EM_design.set_variable('Stock_H', '%.3fin'%(stock_H))
cav_dims=[Stock_L, Stock_W, Stock_H]

box=model.draw_box_corner([-Stock_L/2, -Stock_W/2, 0], cav_dims)
objs=model.get_object_names()
obj_name=model.subtract(blank_name=objs[1], tool_names=[objs[0]])
```

This process results in the automatic assignment of the interior volume as *vacuum*, which HFSS treats as already having perfect conductor boundaries. In our coaxial cavity design, the cavity top is open to the environment, with the TEM mode of the pin coupled to a sub-cutoff TM_{01} mode of a circular waveguide, leading to evanescent decay. One aspect in our design worth considering is how that evanescent field interacts with the top of our cavity. A natural place to start is to set this to vacuum boundary impedance $Z_0 = \sqrt{\mu_0/\epsilon_0} \sim 377\Omega$ as follows:

```
top_hole_loc=[0,0,Stock_H]
top_hole_fid=model.get_face_id_by_pos(obj_name, pos=top_hole_loc)

top_imped=model.assign_impedance(377, 0, obj_name, top_hole_fid, name='Top_Imped')
```

For now we will skip the rest of the intervening steps, which include the creation of a hole for our qubit to be coupled into the cavity, which also has an assigned boundary impedance, and go to the simulation setup. In our case we want to know the eigenmodes of

our cavity. If we remember from Ch. 2, the analytic process of doing this is by solving for the resonant conditions that satisfy Maxwell's equations given our boundary conditions. In effect, HFSS is doing the same, except it attempts to numerically compute this condition through the simultaneous solving of Maxwell's equations over many finite elements. These surface elements, known as a *mesh*, have a characteristic length scale, which is made continuously smaller until a convergence criteria is met in solving for our desired resonant condition. To first order, this criteria is simply the eigenfrequencies of our system. This convergence can also include the loss as well, or the imaginary component of the returned eigenfrequencies. To initially set these parameters we type the following:

```
EM_setup=EM_design.create_em_setup(name='Test_EM',
                                   min_freq_ghz=2,
                                   n_modes=5,
                                   max_delta_f=0.1,
                                   max_passes=20,
                                   min_converged=1,
                                   converge_on_real=True)
```

The `min_freq_ghz` key acts as a means to tell HFSS how to seed the initial mesh. `n_modes` sets how many higher-order modes we wish to compute, and the additional keys set the convergence criteria, with `converge_on_real` telling HFSS to only converge on the frequency, instead of the frequency and loss. After running `EM_setup.analyze()`, HFSS will run the simulation until convergence is met, or the maximum passes have been reached. Following this we can return our solution data and eigenfrequencies:

```
solns=EM_setup.get_solutions()
eigen_real=solns.eigenmodes()[0]
eigen_imag=solns.eigenmodes()[1]
```

3.3.2 Computing Surface Participation

As we saw in Ch. 2, to characterize the loss we must first know the amount of electric and magnetic energy is stored at the source of our loss, i.e. the surface, versus the losses vacuum. To do this we need to compute P_e and P_m as written in Eqs. 2.49 and 2.54 respectively. We *could* export the vector fields and do these integrals numerically in Python, however we can use HFSS, and PyHFSS, to do the work for us. There are two ways we can do this in an automated fashion, but the first way is to use API function calls *directly* to the calculator. To access the built in field calculators we can use the `CalcObject` class as follows:

```
calcobject=HFSS.CalcObject([], EM_setup)
```

Now, before we continue, we should talk about the way we construct an expression in HFSS. The syntax, known as *reverse polish notation* (RPN), has the operand *proceed* their operators. As a simple example, $2 + 2 \rightarrow 22+$. For most who have never used RPN, seeing this for the first time is disorienting, but with this in mind, lets remind ourselves what we we are trying to do. Using P_e as an example, from Eq. 2.48 (a), we are taking the volumetric integral of the quantity $|E|^2$. In RPN this would look like $\vec{E}\vec{E}^* \cdot \mathrm{Re}$. For the denominator of P_e we can express this as:

```
solns.set_mode(0)
vecE = calcobject.getQty("E")

A = vecE
B = vecE.conj()
A = A.dot(B)
A = A.real()
A = A.integrate_vol(name=volume)

lv = self._get_lv(variation)
vol_int=A.evaluate(lv=lv)
```

Line 1 tells HFSS the mode of the associated field we want to work with. Line 10 returns a quantity in this case in Vm^2 in this case. We can swap out line 8 for `A.integrate_surf(name=surface)` to get the numerator, where `volume` and `surface` variables are the identifiers of the desired surface or volume we wish to integrate[10]. For P_m we can change line 1 and instead use `calcobject.getQty("B")`.

The above is how we can use the API to directly calculate these quantities. We can also pre-compile this code into code that HFSS can directly compile, and save it to the native `.clc` file. To load and use a compiled `.clc` we instead use:

```
1  solns.set_mode(0)
2  calc=HFSS.HfssFieldsCalc(EM_setup)
3  calc.clear_named_expressions()
4  Se_calc_path=path+"\\E_energy_S_to_V.clc"
5  Sm_calc_path=path+"\\H_energy_S_to_V.clc"
6
7  Se_name=calc.load_named_expression(Se_calc_path)
8  Sm_name=calc.load_named_expression(Sm_calc_path)
9
10 calc.use_named_expression(Se_name).evaluate()
11 calc.use_named_expression(Sm_name).evaluate()
```

One reason we chose the coaxial geometry for this instance is to compare the simulation to an analytic approximation. Since the coaxial cavity can be approximated as a TEM mode, where we consider the evanescent section above to have minimal contribution to the participation of the fields. From [144] we can show the electric and magnetic field is:

$$\vec{E}_\rho = \frac{V_0 e^{-i\beta z}}{\rho \ln b/a}\hat{\rho} \tag{3.1a}$$

10. In this case the identifier is the name assigned by the user or HFSS for the solid and are represented in our function as strings. One can also select specific surfaces or volumes which can be independently assigned or are assigned by default with an ID number. You can find the ID number by position using `model.get_face_id_by_pos(obj, pos)`

$$\vec{H}_\phi = \frac{V_0 e^{-i\beta z}}{2\pi\eta\rho} \tag{3.1b}$$

From the above we can solve the magnetic and electric participation using Eqs. 2.49 and 2.54). We will only look at the geometric participation, s_e and s_m respectively, and ignore the dielectric thickness and constant and the London penetration length for the time being. Normally, for a coaxial line, and because of the symmetry of the TEM mode, the surface to volume participation of both fields would be the same. In our case the bottom surface also participates, but because the electric field goes to zero at the boundary we can still calculate p_e as:

$$s_e = \frac{1}{\ln b/a}\left(\frac{1}{a} + \frac{1}{b}\right) \tag{3.2}$$

For p_m, we must also calculate the participation of the bottom surface. We do this by taking the surface integral of the tangential field, similar to how we calculated the end of the evanescent hole participation in Eq. 2.67. Taking both contributions gives us:

$$s_m = \frac{1}{\ln b/a}\left(\frac{1}{a} + \frac{1}{b}\right) + \frac{1}{l_{pin}} \tag{3.3}$$

Here, $a = 1.58$mm is the pin radius, while $b = 7.62$mm is the total cavity radius, while $l_{pin} = 9.84$mm is the pin length. We can compute $s_e = 477m^{-1}$ and $s_m = 627m^{-1}$ from the above. Now lets compare to the simulated results. For the same values in dimensions, the simulated results yield $s_e = 392m^{-1}$ and $s_m = 548m^{-1}$, which is $\sim 17\%$ lower in s_e and a $\sim 13\%$ lower in s_m from the analytic expression. One reason is that the simulated geometry is not a simple shorted coax, with additional curvilinear surfaces at the top and bottom of the pin[11]. The main reason, however, is that the evanescent field increases the

11. The reason for this will be explained in a later section, but in short the removal of sharp edges was

effective volume by a noticeable amount, increasing the denominator of both integrals. We could have included this deviation in our analytic solution, but the coupling of the TEM to the evanescent TE_{11} mode greatly complicates the above expressions. Either way we see that the accuracy of the simulation in capturing the predicted behavior of the system is in line with expectations. In addition to this we can also look at the exact participation of specific parts of the geometry independently, which is useful when correlating the presence of particular impurities with types of loss, as we will see in later chapters.

3.4 Q3D and static qubit simulations

In the previous sections we introduced and discussed the creation of a 3D coaxial cavity and the process of extracting both its eigenfrequencies, quality factor, and electric and magnetic geometric participations. For cavities this represents the extent of the analysis one can do. Ultimately we would like to integrate a qubit into this geometry and study how coupling our resonator to a non-linear LC circuit changes the energy spectrum of our system. In a later section we will use Blackbox [127] and `PyEPR`[122] to extract the dispersive coupling rates of our combined qubit-cavity system, but first we need to solve the energy spectrum of our transmon qubit. In this section we will discuss how we can determine the circuit parameters using our automated simulation code and solve for the energy eigenvalues of the CPB Hamiltonian to determine the transmon's energy eigenvalues.

3.4.1 3D transmon circuit analysis

In Ch we introduced the transmon qubit and some of the properties that make it particularly attractive for quantum information, namely its suppression of dephasing thanks to an exponential suppression of charge dispersion. With this in mind it is important that our qubit is both in the transmon limit, or rather $E_j \gg E_c$[104, 164], and has the correct frequency to

done for more predictable etching

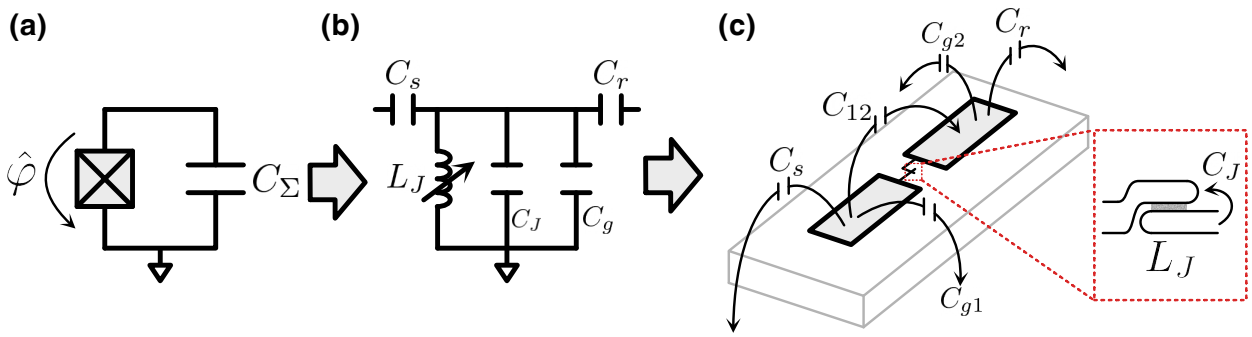

Figure 3.2: **Circuit model and capacitance of CPB**: (a) shows the layout of our cooper pair box (CPB) as a canonical circuit diagram. The inductive energy introduced by the Josephson junction sets the E_j of our circuit, while the charging energy, E_c is associated with the amount of energy required to move e worth of charge across our junction and onto our island. This charging energy is set by the total capacitance C_{Sigma} in the above circuit. (b) shows the real-world layout of our 3D transmon circuit. Here the total capacitance is the effective capacitance of each pad to ground and the cross capacitance of each pad to one another.

be in the desired dispersive coupling regime. In addition to this, it can be shown that the charging energy is equal to the anharmoncity, $\alpha = -E_c$, for the transmon regime. We will discuss how to determine the correct E_j for a given E_c and desired ω_{01} frequency later in this section, but first lets discuss what E_c is.

Fig. 3.2 (a) shows the effective circuit diagram for our 3D transmon. We can note that, for any cooper pair box, $Ej = e^2/2C_\Sigma$ [53]. To get the *total* capacitance C_Σ, we can note that C_{g1} and C_{g2} are in series where as C_{12}, or the cross-pad capacitance is in parallel with our junction. This leads to an effective capacitance of:

$$C_\Sigma = C_{12} + \frac{C_{g1}C_{g2}}{C_{g1} + Cg2} \tag{3.4}$$

Now that we know what C_σ is equivalent to, we need to determine their values. To do so, we will be using Ansys q3d, a quasi-static field extractor. In the next section we will be showing how we implement that, but first lets discuss what q3d actually solves for. Q3d returns a reduced capacitance matrix. The capacitance matrix is obtained by solving:

$$Q = \mathbf{C}v \tag{3.5}$$

Where Q is a column vector of the charge on each island, or in our case, pad, and v is the corresponding voltage on each island. The matrix C is our reduced capacitance matrix. For our example, lets start with the total charge on each conductor as a function of the voltages on each conductor:

$$Q_1 = C_{g1}v_1 + C_{12}(v_1 - v_2) \tag{3.6a}$$

$$Q_2 = C_{g2}v_2 + C_{12}(v_2 - v_1) \tag{3.6b}$$

From the above we see that we can regroup the voltages to get to recover the reduced capacitance matrix.

$$\begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} = \begin{bmatrix} C_{g1} + C_{12} & -C_{12} \\ -C_{12} & C_{g2} + C_{12} \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} \tag{3.7}$$

Q3d will return the above capacitance values that solve the above system. From those values we can determine the effective capacitances and plug them back into Eq. 3.4. In the next section we will setup this simulation and show how by directly solving the CBP hamiltonion we can recover an effective E_J that will give us the desired $g - e$ frequency ω_{01}.

3.4.2 q3d implementation

:

The q3d simulation process mirrors that of the HFSS simulation setup, with the primary difference being the way in which we import the object and assign both boundaries and the surfaces to be simulated. We will skip most of these steps, and jump to the importation and

surface assignments.

PyHFSS as implemented in PyEPR version 0.8.5.5, does not have a built in .dxf importation functionality. This is not strictly necessary as the one could use the built in 2D design functionality of Ansys HFSS and q3d, however, as with the 3D design, having the ability to export already compiled designs streamlines the workflow. In the branched version of PyEPR used in this thesis this functionality was added. The following is a brief snippet to show its use:

```
1  layer_group=model.import_DXF(chip_path, self_stitch=True)[0]
2  layer_group=model.unite(model.get_objects_in_group(layer_group))
3
4  layer_group=model.separate(names=layer_group, split_plane='XY')
```

In line 1, import_DXF imports all layers, returning a list of layer handles in the design tree. The key self_stitch ensure all surface polygons are combined into closed forms. This allows for line 2 to be used to unite independent objects into a single layer, before each individual component, i.e. the pads and readout resonator, are separated into surfaces that can be independently assigned boundary conditions. In addition to this we must conduct a Boolean operation to turn the now separated 2D forms into sheets that can be assigned as thin conductors:

```
1  cond_plane=model.draw_rect_corner([chip_x_origin,chip_y_origin,chip_z_origin],
   ↪  x_size=chip_x_dim, y_size=chip_y_dim)
2
3  layer_group=model.intersect(names=[cond_plane, layer_group])
```

After this we can create the pad objects, first by using their object identifiers to create a model object from an unassigned sheet, before then assigning the thin conductor properties and finally making the pads nets:

```
#detach and make independent face objects
pad_1_obj=model.create_objects_from_faces(obj=layer_group, faces=pad_1_id)
pad_2_obj=model.create_objects_from_faces(obj=layer_group, faces=pad_2_id)

#assign finite conductivity and sheet thickness
film_thickness=q3d_design.set_variable('pad_t', '%.4fnm'%50.0)
pad_1_obj.make_thin_conductor(name=None, material='Copper',
↪  thickness=film_thickness, direction='positive')
pad_2_obj.make_thin_conductor(name=None, material='Copper',
↪  thickness=film_thickness, direction='positive')

#make nets
pad_1_obj.make_net()
pad_2_obj.make_net()
```

The process of making the thin-conductor pads nets tells q3d to add the conductive surface to the capacitance matrix. Finally we can setup the simulation:

```
q3d_setup=q3d_design.create_q3d_setup(name="Ec_Test",
                                      min_passes=1,
                                      max_passes=10,
                                      min_converged=1,
                                      pct_refinement=30,
                                      pct_error=1,
                                      soln_order='High',
                                      save_fields=False)
```

As before, most of the setup above concerns the convergence criteria. The GC solver for q3d, which is what is being used, is capable of DC and AC measurements, however for capacitance (and conductance) DC is used, so no additional frequency information has to be given. Finally we can extract the reduced capacitance matrix and its units:

```
cap_matrix=q3d_setup.get_matrix_dict()['matrix']['capacitance']
cap_unit=q3d_setup.get_matrix_dict()['units']['capacitance']
```

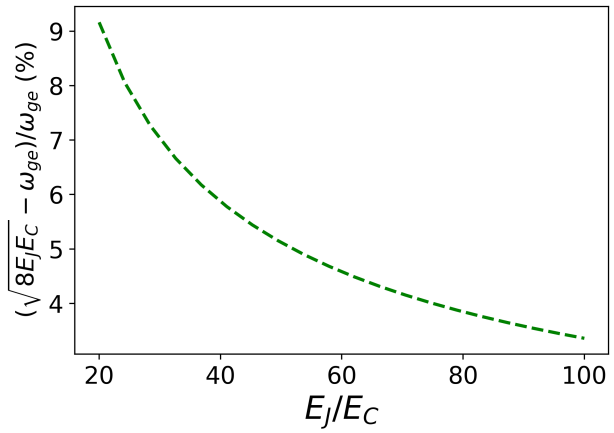

Figure 3.3: **Deviations in ω_{ge} from analytic and numeric approximations**: The figure shows deviations in the numerical solution for the ω_{ge} frequency from the transmon approximation in Eq. 3.8. A complete numerical treatment of the CPB Hamiltonian found in Appendix. A.

With the returned matrix, we can solve for C_{g1}, C_{g2}, and C_{12} using the derived for of the reduced capacitance matrix we derived in Eq. 3.7 before plugging these values into the derived form of C_Σ into Eq. 3.4.

3.4.3 Qubit energies from CPB hamiltonian

With an estimate for E_c in hand we can move on to discuss other properties of our transmon qubit, particularly the estimation of E_j, for a given frequency, as well as some estimations for the charge dispersion limited $T_{2,}$. In Ch. 1 we discussed the unique properties of the CPB and the exponential suppression of charge noise with anharmonicity changing weakly as a power of E_j/E_c. In the transmon regime where $E_j/E_c \gg 1$, the $|g\rangle - |e\rangle$ transition frequency, ω_{ge} can be taken as:

$$\omega_{ge} \approx \sqrt{8E_jE_c} \tag{3.8}$$

The above approximation is accurate within $\sim 10\%$ for typical $E_j/E_c \sim 50 - 100$ values used in the experiments discussed. For more accurate determinations of ω_{ge}, and subse-

quently determinations of E_j, a full diagonalization and solving of the CPB Hamiltonian energy eigenvalues is necessary. A plot showing the discrepancy from the numeric calculation and the transmon limit approximation is shown in Fig. 3.3

Another quantity worth getting at this point is the charge dispersion. Even though the sensitivity to charge noise is exponential suppressed with larger E_j/E_c values, for typical E_j/E_c values, it can still play a noticeable role. The effective dephasing time of a qubit can be written as:

$$\frac{1}{T_{2,eff}} = \left(\frac{1}{2T_1} + \frac{1}{T_{2,\epsilon_n}}\right) \tag{3.9}$$

Where $T_{2,_n} \sim 1/\epsilon_n$ is the dispersion limited dephasing time for the n^{th} level of our CPB. We can reiterate ϵ_n below [104]:

$$\epsilon_n \simeq E_c \frac{2^{2n+5}}{n!}\sqrt{\frac{2}{\pi}}\left(\frac{E_j}{E_c}\right)^{n/2+3/4} e^{-\sqrt{8E_j/E_c}} \tag{3.10}$$

The above is computing the change in energy over one charge period ($\Delta n_g = 1$). For a modest value of $E_j/E_c \sim 60$, from the above we compute a $T_{2,\epsilon_1} \sim 800\mu s$ for the first transition ($n = 1$), which forms the basis for our spin-1/2 system. This may sound like a lot, but for a qubit with $T_1 \sim 100\mu s$, this would lead to a $\sim 20\%$ *decrease* in the maximum effective dephasing time, assuming other sources of dephasing are small.

Before we continue to simulating and extracting parameters describing the full coupled qubit-cavity system, we will note that, while it is an informative exercise to solve this system by oneself, the Python package `SCQubits` as well as Qisket Metal [123] can do the same basic analysis and much more.

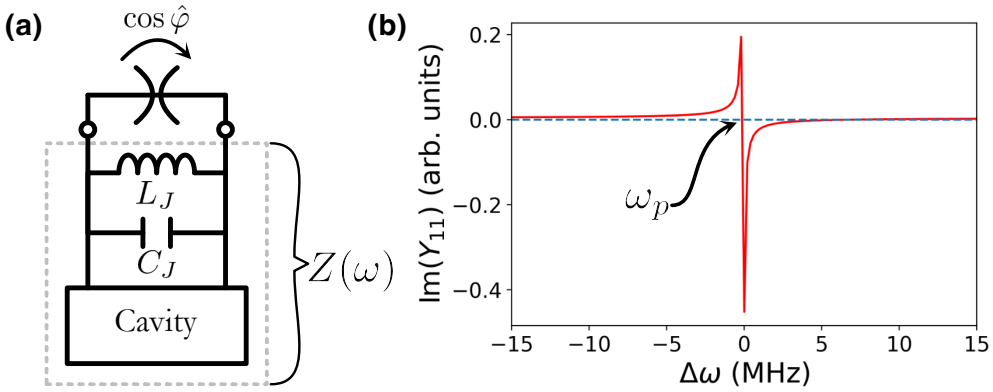

Figure 3.4: **Black-box equivalent circuit and cavity admittance**: (a) A schematic representation of the blackbox quantization process. A qubit is modelled as a linear RLC circuit in parallel with a non-linear Josephson element, coupled to a cavity with impedance $Z(\omega)$. From the perspective of the non-linear element. (b) A plot of the cavity admittance seen from the perspective of the junction. The phase $\hat{\phi}$ across the non-linear element can be described in terms of the imaginary part of the derivative of the admittance, $\mathrm{Im}[Y'(\omega_p)]$, where ω_p can be determined by the zero-crossing of $\mathrm{Im}Y(\omega)$. For small fluctuations in phase, one can expand the $\cos \hat{phi}$ to fourth order and add it as a perturbative correction

3.5 Coupled qubit-cavity system simulation

The static circuit quantization gives us solutions to the energy spectrum of the cooper-pair box in the absence of an external cavity field. In a cavity, however, through the coupling described by the Jaynes-Cummings hamiltonian in Ch. 1, a cavity field will hybridize with the qubit. In 2D this can be done by relating the junction phase to the zero-point voltage fluctuations of the resonator [164]. With discrete elements the coupling can be computed directly, however in 3D, where our electromagnetic energy is distributed, we have to turn to an alternative approach to determining the field participation in the junction. In this section we will discuss the two methods used in this thesis to conduct these calculations and discuss the advantages and disadvantages of both.

3.5.1 Black-box circuit quantization

The first method, known as Black-Box quantization [127], works by splitting the qubit into a linear component with the non-linearity added back into the system perturbatively. Fig. 3.4 (a) shows the equivalent circuit, where we have split the system into an equivalent circuit that describes the cavity modes that is in parallel with the linear component (L_j and C_j) of our qubit and a nonlinear part due to the additional josephson inductance. From the perspective of the qubit we can write the impedance of the Foster equivalent circuit as:

$$Z(\omega) = \sum_{p=1}^{m} \left(j\omega C_p + \frac{1}{j\omega L_p} + \frac{1}{R_p} \right)^{-1} \tag{3.11}$$

At $\omega = \omega_p$, where ω_p is the resonant frequency of the mode, our impedance looks like a pole, whereas the admittance, $Y(\omega) = Z(\omega)^{-1}$ crosses zero, which we will use later to find ω_p values. It can be shown that, from Eq. 3.11 we can calculate the values for C_P, L_p, and R_p in terms of the admittance as [127]:

$$C_p = \frac{1}{2}\mathrm{Im}[Y'(\omega_p)] \tag{3.12a}$$

$$L_p = \frac{1}{2\omega_p^2 \mathrm{Im}[Y'(\omega_p)]} \tag{3.12b}$$

$$R_p = \frac{1}{\mathrm{Re}[Y(\omega_p)]} \tag{3.12c}$$

We can write a Hamiltonian in terms of the conjugate variables flux, $\hat{\phi}_p$ and charge, $\hat{q}_p$ for our circuit in terms of the equivalent circuit parameters:

$$H_0 = \sum_{p}^{m} \frac{\hat{q}^2}{2C_p} + \frac{\hat{\phi}^2}{2L_p} \tag{3.13}$$

We can see that Eq. 3.13 reduces to the quantized form $\sum \hbar\omega_p(a_p^\dagger a_p + 1/2)$. We can rewrite the flux in terms of the creation-annihilation operators. We can also write the phase across the junction as $\varphi = \sum \phi_p/\phi_0$, where $\phi_0 = h/2e$ is the flux quantum. Combining these two gives:

$$\varphi = \frac{1}{\phi_0}\sum_p^m \sqrt{\frac{\hbar}{2}Z_p}(a_p + a_p^\dagger) \tag{3.14}$$

Where $Z_p = \sqrt{L_p/C_p}$. Now we have a value for the phase φ across the junction from the contributions of each mode. Assuming the phase is small we can expand the non-linear part of our CPB Hamiltonian in orders of φ:

$$H_{nl} = E_j(1 - \cos\varphi) - \frac{1}{2}E_j\varphi^2 = E_j\bigg(- \frac{\varphi^4}{4!} + \mathcal{O}(\varphi^6)\bigg) \tag{3.15}$$

We can plug in Eq. 3.14 into Eq. 3.15, careful to discard non-conserving combinations in a_p and $a_p^\dagger$ in the fourth order expansion, to get an effective non-linear Hamiltonian:

$$\sum_p \hbar\Delta_p n_p + \frac{1}{2}\sum_{p,p'} \hbar\chi_{pp'} n_p n_{p'} \tag{3.16}$$

In the above $n_p = a_p a_p^\dagger$ is the number operator on mode p, while Δ_p is a linear correction to frequency due to the lamb-shift. The second sum on the right contains the term $\chi_{pp'}$, which represent the self-Kerr ($p = p'$) or cross-Kerr ($p \neq p'$). These can be written as follows[12]:

$$\chi_{pp} = -\frac{L_p}{L_j}\frac{C_j}{C_p}E_c \tag{3.17a}$$

12. It should be noted that, by convention, $\chi_{pp'}$ is equivalent to the dispersive shift 2χ. This is the same in the later `PyEPR` methodologies.

$$\chi_{pp'} = -2\sqrt{\chi_{pp}\chi_{p'p'}} \tag{3.17b}$$

Eqs. 3.17 (a) and (b) give us a wealth of information about our system, since they describe the effective dispersive interaction between the modes and the qubit, as well as with neighboring modes, in the case of a multimodal system. It is easy to look at Eq. 3.17 and wonder what the point of this exercise was, but we should remember that the circuit parameters are derived *from* an exact electromagnetic description of our cavity from the admittance of the modes at the qubit, including distortions in the electric and magnetic field profile from the addition of participating metal and dielectric surfaces on the chip. With the above we can also compute the inherited anharmonicity of each mode, α_p, which is equivalent to the self-Kerr χ_{pp}, and the vacuum Rabi-rate g_p for each mode:

$$g_p = \sqrt{\chi_{qp}\Delta_{qp}\frac{\Delta_{qp}+\alpha_q}{\alpha_q}} \tag{3.18}$$

Where we use the subscript q to denote the interaction of a mode with the equivalent linear contribution of the qubit. Here Δ_{qp} represents the detuning of the qubit and target mode frequency, $|\omega_q - \omega_p|$, and α_q is simply the qubit anharmonicity.

While this is all interesting in an abstract sort of way, it isn't all that helpful without a way of accurately simulating the cavity-qubit system and extracting the admittance. In the next section we will briefly go over the important aspects in this process, and discuss how we can use `PyHFSS` again to extract this data.

3.5.2 Black-box in PyHFSS

The last section showed how one could use an equivalent circuit model and a measurement of the electromagnetic environment experienced by the qubit to approximate the qubit-cavity

interaction. Because this technique is dependant on a simulated result, the quality of the information extracted is only as good as the model being used to simulate the system. So far in our discussion of using HFSS (and `PyHFSS`), we have only discussed simulating cavities and extracting participation. Black-box differs in two ways from those simulations. The first is that the admittance is a function of frequency as defined by Eq. 3.11, or more accurately, the inverse. This means we need to drive the system at different frequencies and get back the response function. The second important factor goes back to the way we defined the problem in Fig. 3.4 (a), namely that we want to get the admittance of the cavity modes at the qubit. To do this we need to create a *driven modal* simulation where the port that we want to drive *is* the qubit itself.

Much of the setup is the same as with the eigenmode and q3d extractor setup in previous sections. Here however we select the driven model solution type when initializing our design by using the `new_dm_design` function in the `project` class of `PyHFSS`. We will skip the basic setup, and jump to the setup of our driven model port, which in our case, from Fig. 3.4 (a) is the qubit. We first create a box that is a non-model object that spans the two pads:

```
1  model.assign_perfect_E([pad_1, pad_2], name='TPads')
2  Junc_y=DM_design.set_variable('junc_y', Pad_y+Pad_L)
3  Junc_x=DM_design.set_variable('junc_x', Pad_x+(Pad_W-Junc_W)/2)
4
5  junc_loc=[Junc_x, Junc_y, Pad_z]
6  jj=model.draw_rect_corner(pos=junc_loc, x_size=Junc_W, y_size=Pad_gap, name='jj')
```

In the above we have assigned the pads to be perfect conductors, and set variables that represent the pad locations for easy reference, and use `draw_rec_corner` to create the rectangle. Next we need to tell HFSS that this box represents a *lumped port* in which we are driving through.

Lines $1-3$ on Pg. 84 are simply meshing operations to help in simulating the small

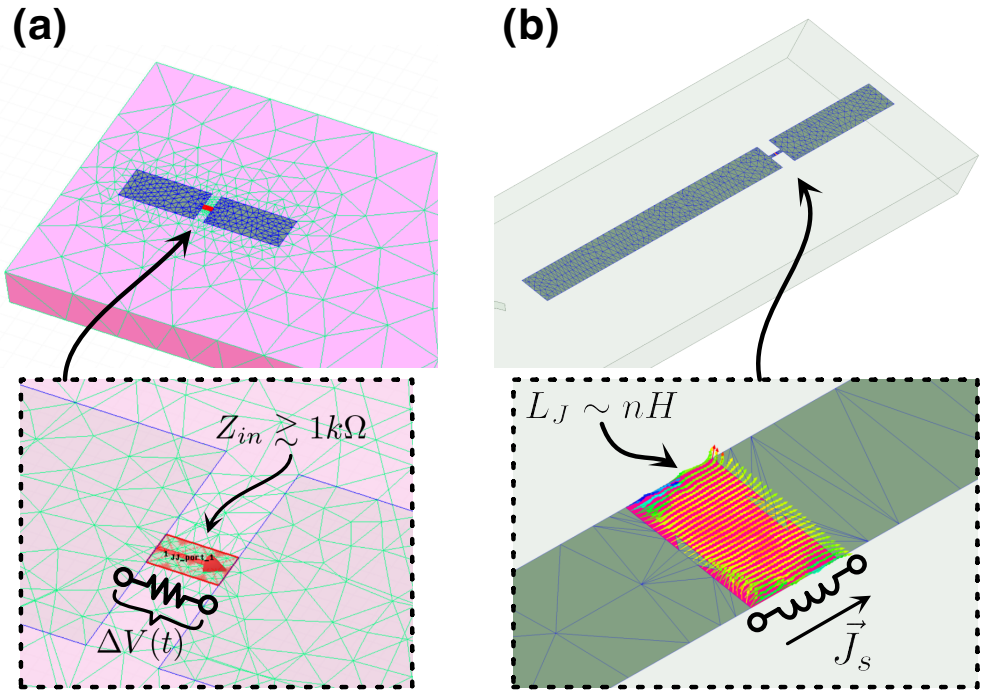

Figure 3.5: **Black-box and EPR simulation setup comparison**: The above figure shows the qubit where the junction area has been assigned as a high-impedance port. Also shown is the mesh operation of the junction, pads, and chip to better capture the dynamics of the system. The inset shows a zoomed in picture of the junction with the port assignment. The voltage integration line is draw across the nodes seen in Fig. 3.4 (a). For comparison, (b) shows the eigenmode setup for the electric-participation ratio (EPR) techniques, where instead at looking at the time dependent voltage across the nodes, instead one integrates the current density to extract the participation of a mode across the inductive element of our qubit. For this, the junction is assigned as an RLC lumped-element, with inductance equal to L_J corresponding with E_J and E_C to achieve a desired $_{ge}$. Similarly, an integration line is made across the pads, which correspond with the nodes across the linear inductor shown diagramatically in the inset.

```
1  model.mesh_length(name_mesh='jj', objects=[jj], max_length='%.3fum'%20)
2  model.mesh_length(name_mesh='pads', objects=[pad_1, pad_2], max_length='%.3fum'%100)
3  model.mesh_length(name_mesh='substrate', objects=[chip], max_length='%.3fmm'%1)
4
5  jj.make_lumped_port(axis='Y', z0="1000ohm", name="JJ_port_1")
```

features of the qubit more accurately, while line 6 is the creation of the lumped port. The `axis` flag in line 6 represents the integration line of the port. This axis is in reference to

the rectangle object made in the previous step. Using the lower level `_make_lumped_port` function can allow for the assignment of arbitrary axis if the object assigned does not sit along a cardinal axis. The impedance, Z_0 is set to be high since it acts like a parallel resistance path for current to pass, meaning a low impedance can artificially load the cavity.

Once the setup is complete, the final steps are setting up the simulation. As before, the driven modal setup is primarily required for designating convergence criteria. In addition to this, there is also a sweep setup that sets the frequency sweep and the type of solver to be used. Below is an example of this setup:

```
DM_setup=DM_design.create_dm_setup(freq_ghz=1,
                                   name="QTune_DM",
                                   max_delta_s=0.1,
                                   max_passes=10,
                                   min_passes=1,
                                   min_converged=1,
                                   pct_refinement=30,
                                   basis_order=-1
                                   )

DM_sweep=DM_setup.insert_sweep(start_ghz=5,
                               stop_ghz=7,
                               count=10000,
                               step_ghz=None,#<--If you rather use frequency step
                               type="Fast",
                               save_fields=False)
```

The maximum number of sweep points is 25000, meaning for multimodal systems which cover large bandwidths (and may have narrow features) it may be required to split the range into smaller sweeps and recompile the full spectrum after importing the sweep data.

Once analyzed, HFSS can return either scattering, impedance, or admittance matrices. Given that our device is one port, we will only look at the Y_{11} values. We can return both the real and imaginary components separately using the following:

Using the returned arrays of frequency and $\mathrm{Im}Y_{11}(\omega)$ data, we can go on and compute

```
Y11_report=DM_sweep.create_report(name=DM_sweep.name,
    expr="im(Y(JJ_port_1,port_1))")
Y11_array=DM_sweep.get_report_arrays(expr="im(Y(JJ_port_1,JJ_port_1))")

freq=Y11_array[0]
im_Y11=Y11_array[1]
```

the quantities listed in the previous sections. The required derivative $d(\mathrm{Im}Y_{11}(\omega))/d\omega$ can be computed numerically through finite difference, while the roots, which correspond with ω_p for each mode, can be determined through iterative means as well. Once these values are obtained χ and the equivalent circuit parameters can be computed analytically.

3.5.3 Electric-participation ratio (EPR) quantization

In the previous section we discussed how we can determine the phase across the junction based on the admittance of qubit at the resonant frequency of a target mode. This technique, while relatively straightforward, requires that we extract data from driven modal simulations, which have limited precision. This introduces more uncertainty into the final calculations. In addition to this, simulating multiple qubits coupled to a single resonator is challenging with the black-box technique.

EPR attempts to solve these problems by relating the phase of the junction instead to the amount of the total stored electromagnetic energy in the junction versus in a given target mode [122]. This participation ratio, which is similar to the ones we have seen previously in our discussion of cavity simulations and loss mechanisms, is easily determined through eigenmode simulations, which are more straightforward to compute, and have precision only set by the mesh size (and computer memory). EPR also has the advantage of being able to compute the expected dispersive coupling between multiple modes and multiple qubits.

To begin we can define the participation ratio as the time averaged inductive energy stored in the junction (qubit) versus a target mode m:

$$p_m = \frac{\langle\psi_m| \frac{1}{2} E_j \hat{\varphi}_j^2 |\psi_m\rangle}{\langle\psi_m| \frac{1}{2} \hat{H}_{lin} |\psi_m\rangle} = \frac{E_j \hat{\varphi}_{jm}^2}{\frac{1}{2}\hbar\omega_m} \tag{3.19}$$

Where state ψ_m can represent either a Fock or coherent state of the resonator. Here, E_j is the Josephson energy of the junction, $\hat{\phi}_{jm}$ is the phase across the junction caused by target mode m, and ω_m is the angular frequency of the target mode in question. The inductive energy of the junction is stored in a geometric inductance, and a kinetic inductance associated with the junction, $W_{ind} = W_g + W_{kin}$. The contribution of the geometric inductance of the qubit is part of the measured inductive stored energy of the mode. The time averaged inductive energy is equal to the capacitive energy. The value of p_m is therefore just the ratio of the junction's kinetic inductive energy and the inductive energy of the system. With all of this in mind we can write p_m in terms of the stored energy in the magnetic and electric fields:

$$p_m = \frac{W_e - W_m}{W_e} = \frac{\int_V \epsilon |\vec{E}|^2 dV - \int_V \mu |\vec{H}|^2 dV}{\int_V \epsilon |\vec{E}|^2 dV} \tag{3.20}$$

The above energies can be extracted using the same techniques described in Sec. 3.3.2 directly from the eigenmode simulation. With p_m in hand, we can solve for the junction phase induced by target mode m using Eq. 3.19. With $\hat{\varphi}_{jm}$, we can make a similar perturbative approximation for the non-linear contribution of the junction phase as we did in Eq. 3.15. This is for the general single junction case, however, for multiple junctions the kinetic inductance is split amongst the junctions. For this we need to directly compute the inductive energy at the junction by looking at the current across a junction with index j with a linear inductcance L_j:

$$W_{kin,j} = \frac{1}{2} L_j I_j^2 \tag{3.21}$$

Where the current can be determined by integrating the current density across junction j due to mode m over the cross section representing the junction S_j:

$$|I_{mj}| = \frac{1}{l_j} \int_{S_j} |\vec{J}_{jm}| ds \tag{3.22}$$

We introduce the length of the sheet representing our junction l_j. The new participation for junction j is simply:

$$p_{jm} = \frac{L_j \int_{S_j} |\vec{J}_{jm}|^2 ds}{2l_j^2 \int_V \epsilon |\vec{E}|^2 dV} \tag{3.23}$$

Following the perturbative expansion of the non-linear Hamiltonian, we can write the effective dispersive interaction as we did in Eq. 3.16, with χ_{qq} and χ_{mm}, which are the qubit and cavity self-Kerr as:

$$\chi_{qq} = -p_q^2 \frac{\hbar \omega_q^2}{8E_j} \tag{3.24a}$$

$$\chi_{mm} = -p_m^2 \frac{\hbar \omega_m^2}{8E_j} \tag{3.24b}$$

Which gives us the qubit-cavity cross-Kerr as:

$$\chi_{qm} = -p_q p_m \frac{\hbar \omega_q \omega_c}{4E_j} = -2\sqrt{\chi_{qq}\chi_{mm}} \tag{3.25}$$

We see that the right side of Eq. 3.25 is equivalent to Eq. 3.17 (b). It should be noted that, through energy conservation, $p_q + p_m = 1$. It should also be noted that we have imposed the sign of the dispersive shift accordingly.

Even though the way we extracted the effective junction phase differed from black-box, the end result, namely calculating the dispersive interaction between the modes of our cavity and the qubit, was the same. In the next section we will show the setup and implementation

of the above in practice. We will see in the next section a short introduction into how we use the above in python to simulate this interaction.

3.5.4 PyEPR simulations

Zlatko Minev and Zaki Leghtas, along with others, developed the EPR method with a corresponding package which incorporates all of the tools necessary to compute the full interaction matrix for a given qubit-cavity system. They have also provided an excellent set of basic tutorials in how to use `PyEPR` that[13], along with the paper (Ref. [122]), are excellent companions for anyone attempting to use the `PyEPR` package. Because of this, I will not go into much detail, but for completeness I will discuss briefly the key aspects in the simulation setup.

As shown in the previous section, EPR computes the participation ratio p_m by integrating the electric field and current density across the junction. All of these values are obtained from an eigenmodal simulation, as opposed to black-box's driven modal simulation. While the `PyEPR` package handles all of the non-linear Hamiltonian calculations for us, under the hood it is doing the same set of operations that we used in Sec. 3.3.2 to compute the integrals seen in Eqs. 3.20 and 3.23. With this in mind, the ultimate accuracy in the simulation hinges on the convergence of the field-profiles in simulation. Because the junction represents such a minuscule volume of the total mode volume, it is important to setup the simulation in a way that accurately models the field at the junction without unnecessarily wasting computational resources meshing the entire cavity surface. In the black-box setup we demonstrated this by meshing the pads and junction box with much finer resolution, with the maximum mesh length no greater than the characteristic length of the object, as seen on Pg. 84.

The one difference between black-box and EPR is the way we add the linear inductance of the junction into the simulation. Because `PyEPR` is written in a general form to include

13. https://github.com/zlatko-minev/pyEPR/tree/master/_tutorial_notebooks

both single and multiple qubit designs, the form of p_m used is that of Eq. 3.23. Because of this the linear part of the Josephson inductance is required, along with the integration line in which we want to compute the current:

```
1 jj=model.draw_rect_corner(pos=junc_loc, x_size=Junc_gap, y_size=Junc_width, name='jj')
2 L_jj=EM_design.set_variable('Lj_1', '%.3fnH'%Lj_val)
3 jj.make_rlc_boundary(axis='X', r=0, l=L_jj, c=0, name='jj1_RLC')
4 jj_line=model.draw_polyline(jj.make_center_line(axis='X'), closed=False, name='jj_line')
```

On line 1 we repeat the process of defining the junction box as we did on Pg. 83. However, unlike black-box, we need to include the linear inductance *in-situ*. In line 2 we define our variable containing a target inductance value associated with an E_j for our desired ω_{ge} transition frequency, which we can get from the calculated using Eq. 3.8. Line 3 sets the junction boundary as a lumped element RLC circuit with the target inductance. To draw the integration line, which needs to be passed separately into the EPR calculator later, we need to define a variable for the line l_j, which is done in line 4.

The eigenmode setup is in effect identical to that shown on Pg. 68. Because we have parameterized L_j we can do an optimetric sweep to optimize the target qubit frequency as follows on Pg. 90, Where we define some variable to step the L_j value over.

```
opti_setup.create_setup(variable=L_jj,
2                       swp_params=('%.2fnH'%(Lj_val-num_step*step),
                        ↪  '%.2fnH'%(Lj_val+num_step*step), '%.2fnH'%step),
3                       name='Lj_Sweep',
4                       solve_with_copied_mesh_only=False
5                       )
```

To finish the extraction of the dispersive coupling rates one can follow the routine outlined in the aforementioned `PyEPR` tutorial notebooks. The returned values include the Kerr

and self-Kerrss for all of the simulated modes, which can be designated in the eigenmode simulation setup as seen on Pg. 68, and every variation from the optimetrics sweep.

In summary, the preceding sections have shown the full-stack of simulation techniques for extracting design specific cavity and qubit parameters. In the intervening chapters we will see how to use this extracted data to inform our design to minimize loss mechanisms and control the cavity and qubit spectrum and coupling rates. In the next chapter, as we discuss the development of seamless flute cavities with non-trivial geometries, and discuss the coupling dynamics between a single qubit to multiple modes, we will see more examples of how the preceding is used.

CHAPTER 4

THE SEAMLESS FLUTE CAVITY

4.1 Introduction

By this point we have gained a theoretical and practical framework for what goes into making a good cavity for 3D cQED systems. To summarize, a good cavity should minimize residual material losses and the geometric electric and magnetic participation, and should ideally have no lossy seams for supercurrents to cross. That last part, the seams, may seam to be the most obvious and straight-forward problem to remove, but have been a large part of the discourse of superconducting cavity performance. The reason is that, for practical reasons, splitting a cavity geometry in two halves makes manufacturing *significantly* easier. Because of this, a lot of work has been done to create seams with as little loss as possible. In elliptical niobium cavities used for particle acceleration, extensive resources are spent on producing high-purity autogenous welds using vacuum electron-beam welding [166]. In 3D cQED applications, and in particular TE_{101} rectangular waveguide cavities, losses are minimized through geometric optimization of the seam location [31], and the creation of low-loss indium joints through the cold-welding of indium bump bonds [110].

The other route is to remove the seam altogether. Because the cavity volume still has to be made through the removal of bulk material[1], through machining operations requiring tool access, a cavity can be designed where the tool access creates a hole that has a waveguide cutoff frequency above the fundamental mode of the cavity. The coaxial $\lambda/4$ cavity which we have previously discussed, is based off of this idea. As long as the frequency, which is set by the pin height of the cavity, is lower than the TE waveguide mode, then the loss from the mode can be exponentially suppressed with length. One annoying consequence

1. It should be noted that cavities made with additive manufacturing have been developed in both aluminum and niobium, however the 3D printing process produces material which has worst bulk performance when compared to high-purity wrought base material [109, 117, 54, 70]

of this, especially for very high quality factor cavities, is that the aspect ratio of the hole length to the pin height can become quite large, making tool access and machining of the pin feature difficult. One other drawback of coaxial cavities is that, while they have higher order modes, they are typically far detuned from the fundamental mode. This is an advantage in situations where parasitic coupling to lossier higher-order modes is a problem, but means it is more challenging to create multimodal designs. Inversely, rectangular waveguide TE_{101} style cavities are intrinsically multimodal, and, as we will see, can be designed in a way to control their mode dispersion. With this mind, can we adapt the ideas of the coaxial cavity into a rectangular waveguide cavity, machining all of the volume through only evanescent access holes?

4.2 The flute cavity

The basic idea of the flute cavity is quite simple; by drilling two sets of opposing holes, such that the center-to-center distance of the opposing holes is *less* than the hole diameter, it is possible to effectively "core" out a cavity volume. To most observers, this design may seem doomed from the start; how can a cavity with so many holes be any good? Intuitively we can see that the holes, which are much smaller in diameter than the wavelength of the mode, effectively screen the internal cavity field, much like a Faraday cage on the inside of a microwave door. The more exact description, which was discussed in some detail in Sec. 2.3.3, showing that the external coupling of a sub-cutoff mode through a waveguide is exactly described as being exponentially suppressed with hole length. There, we only looked at how one hole scales. In that case, we set the hole location to be centered on an antinode of the target mode, where the coupled field, described by $\vec{H}_x$, was at a maximum, at $z = l_{cav}/2$, and $\sin(\pi z/l_{cav}) \to 1$. This means that, in addition to length, and hole radius, the coupling strength is *also* dependent on the location along the $\vec{z}-$axis. Of course, the above spatial variation of the field profile is only true for the fundamental mode, with higher order modes

having similar modulation, with different periodicity, and also having frequencies closer to the cutoff, leading to slightly higher propagation. We will quantify this in more detail in the next couple of sections.

4.2.1 Quantifying loss and flute scaling

Since the quality factor of any given mode is based on the distribution of field energy, which is itself not uniform, the exact scaling of the quality factor for a given number of holes is non-trivial for a given hole diameter and length. We will, however, briefly discuss the scaling of the quality factor and frequency of a sub-cutoff cavity mode as a function of increasing the number of evanescent holes in a TE_{101} cavity. In chapter 3 we showed that the addition of any evanescent hole increases, albeit slightly, the effective mode volume. In addition to this, the coupling Q to the hole is dependent on the radius, length, and location.

In Fig. 4.1 we see that the decrease in mode frequency is monotonic and linear with number of holes. We can see why if we look at the radiated power through the waveguide for it's propagating TE_{11} mode:

$$P_{\text{rad}} = \int_s \vec{E} \times \vec{H}^* \dot{ds} \tag{4.1}$$

The term in the integrand is the Poynting vector of the waveguide mode, representing the direction of energy flow down the hole. From Eqs. 2.65 and 2.66 we can compute the radiated power, which give us the expression:

$$P_{rad} = i\frac{\pi\omega\mu|A|^2\beta_{11}}{4k_c^4}\left(p_{11}'^2 - 1\right)J_1^2(k_c r_0) \tag{4.2}$$

We see that, because our complex Poynting vector is imaginary due to the imaginary propagation constant of our sub-cutoff mode, the radiated energy is also imaginary. This means that our radiated power out of the hole is zero (but *not* the electric and magnetic

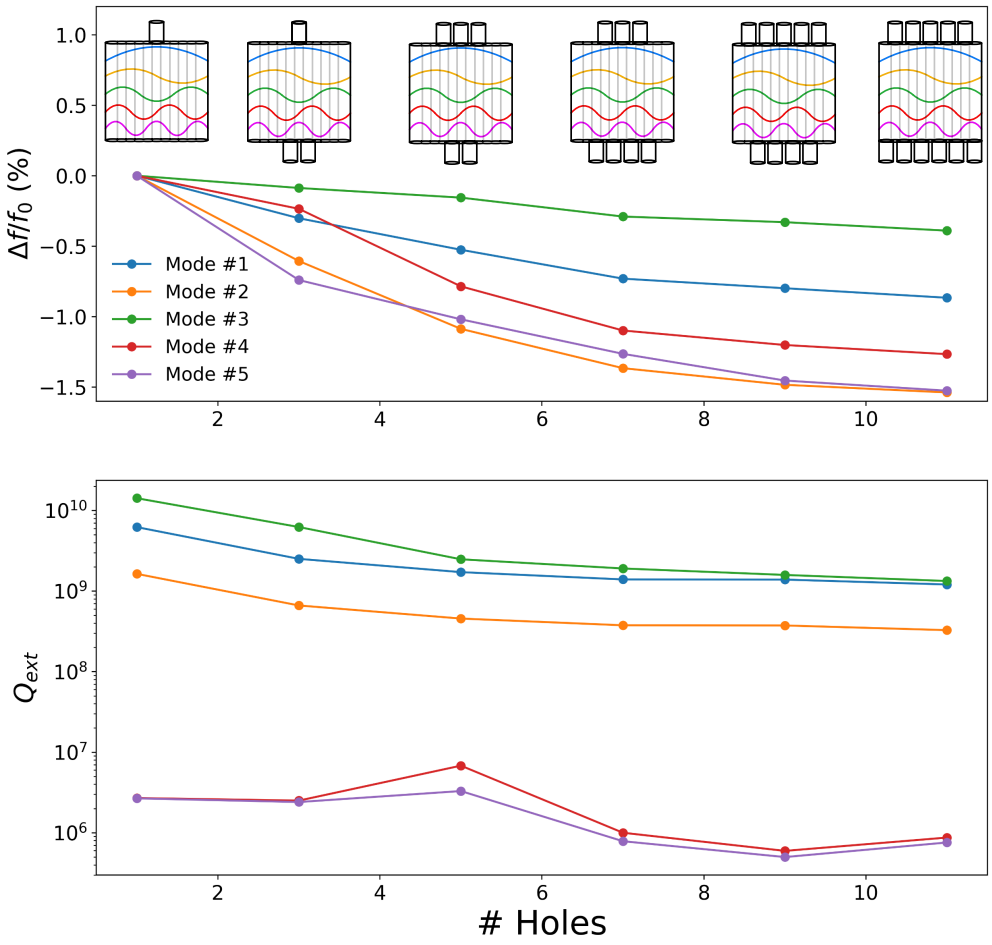

Figure 4.1: **Scaling of Q due to evanescent loss of a single-mode flute cavity**: The above shows a simulated scaling in the quality factor for the first 5 modes of a ∼6.3GHz TE_{10n} "pan-pipe" geometry. We see a monotonic decrease in the frequency as expected from the theoretical treatment in Ch.2. The quality factor scaling also decreases for all modes as the hole number increases, however if we inspect the relative quality factors between modes we begin to see the spatial dependence of the field.

field participation at the end of the hole), but the hole *does* add additional reactance to the mode, changing the frequency. The decrease in frequency can be seen as an effective increase in the mode volume [149]. The above wave amplitude A is a product of the dipole coupling of the TE_{101} mode to the hole using the Bethe coupling equations, which was previously derived, and as with Q_{ext} is dependent on the hole diameter, the location of the hole along the length of the cavity, and the mode's electromagnetic field profile. This overall change

is shown to be as much as 1.5%, with the complex convolution of spatial and frequency dependence. For most single mode, and even multimode, cavity designs, this perturbation is not of major concern, but for precise tuning of the cavity frequency, these additional effects should be considered when designing the cavity, as even an 1.5% changes, as is the case in the above example, is a change in frequency of nearly 84MHz, with larger cavities with more holes experiencing even greater deviations. These deviations, however can be compensated for by decreasing the characteristic dimension of the cavity, although the change may impact different modes differently.

The Q_{ext} variation across modes for a few holes is largely dependent on the specific location of the hole with respect to the field profile. For a flute cavity, where the holes are used to make the cavity volume, the added dissipation scales as the length and radius of the hole.

4.2.2 Multimode flute cavity design

As we have seen, any cavity design is inherently multimodal in nature, however, as discussed in chapter 2, the electric field profile of the cavity varies depending on the mode. In chapter 3 we discussed the ways that we can compute the coupling of our transmon qubit to the cavity modes. We saw in Sec. 3.5.3, Eq. 3.20, that the scaling of our effective dipole-like coupling and subsequent dispersive shift is dependent on the magnitude of the time averaged electric and magnetic field at the location of the junction. We also saw in chapter 2 that, for TE_{10n}, the mode number determines the location of nodes in the field distribution, meaning that, while the location of a qubit for a given mode may lead to strong coupling, a higher mode may have little to no coupling if a node in the field happens to occur at the qubit location.

In addition to this, if we look at the spacing of frequencies for a standard TE_{10n} cavity, we see from Eq. 2.41 that for modes of the type TE_{10n} that the spacing of the lowest order modes, such that $l/l_{cav} \leq m/h_{cav}$, the frequency of neighboring modes has non-linear

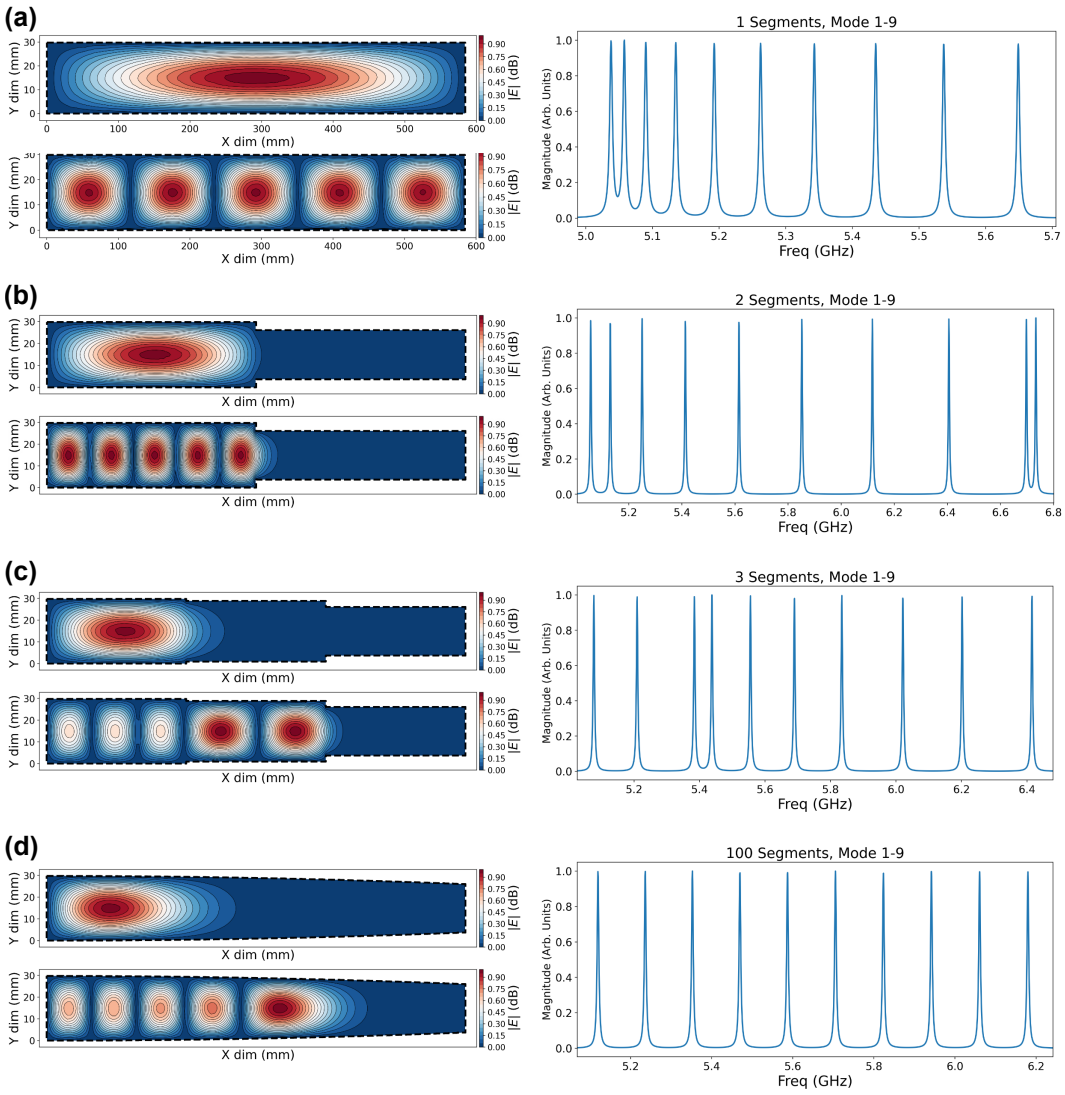

Figure 4.2: **Electric field distribution and mode dispersion for tapered flute**: (a) shows the typical mode distribution and profile for the first and 9^{th} mode of a multi-modal flute cavity, along with the spectral distribution of mode frequencies, highlighting the quadratic nature of modes of small n. (b) shows what happens when half the cavity is made slightly evanescent to the lowest laying mode. The imaginary propagation leads to decay of the field into the second half, with the mode largely contained within the left half, although the effective increase in mode volume only leads to a slight increase in the mode frequency. More importantly, the field, which is now confined to one half of the cavity, has been "lensed" for all modes, meaning that the qubit can now have relatively strong coupling to all of the modes. (c) shows a full taper, such that the modulation of the height is quadratic in nature. We see now that for mode numbers with small n, the spacing of frequencies is close to linear.

spacing, where at higher values of l the spacing should become linear. In practice, that will occur after higher order modes that are transverse to the height of the cavity, of the form TE_{20n} are excited. One point that can be made is how the length and height of the cavities factor into the mode spacing of the system. The first is that, for a given cavity height, increasing the length decreases the mode spacing, meaning more accessible modes can be placed within an operable bandwidth for a given qubit frequency and desired dispersive shift ($\chi \propto 1/\Delta$). Changing the length also changes the frequency of the fundamental mode, however if $l_{cav} \gg h_{cav}$, then the cavity height can be adjusted to compensate for this shift without greatly impacting the mode spacing at higher mode numbers.

With these two technical issues in mind, the nonlinear mode spacing and the period distribution of electromagnetic energy in our cavity, is there a way to design a cavity that forgoes these problems? To adjust the field profile, a first order correction is to take the cavity and make half of the length a different height. If the corresponding cutoff frequency of this new section is less than the first half, then the field will only evanescently decay into the second section. As we have seen with the hole example, this new evanescent volume will add reactance to the mode and increase the mode volume. The effect is to confine the mode to one half of the cavity while also counteracting an increase in frequency since the effective propagating length is now shorter.

To create more evenly spaced lower n modes is a bit more subtle. It may seem that one could come up with a closed-form solution to describe the propagation of a mode, and its eigenfrequencies, for a waveguide of varying height over its length. Naively, one would think describing it would require the simultaneous solving of i equations with $i+1$ boundary conditions, where i is a discrete number of elements in your approximate model, however describing the field *at* those boundaries is tricky, since the scattering of the field off the step-interface leads to interference between reflected and transmitted modes, with some being evanescent. The technique to describe these types of discontinuities is known as the *mode*

matching-eigenmode expansion method, and is a numerically intensive procedure that was worked out in full by WC Chew [111, 42, 41].

Another way of approaching this problem is to use existing finite-element modelling and parametric design to "tune" the cavity geometry to get the desired profile. With the parametric design and simulation tools described in Ch.3 we can generate cavity profiles that give us the desired mode dispersion. One way of approximating the correct profile is to think about the effective length of the cavity from perspective modes. What we want is a profile in which the n^{th} mode has a propagating length that varies like $\sqrt{n}$. This can be seen in the small n limit since the expression f_{nml} from Eq. 2.41 becomes:

$$f_{lmn} \approx \frac{w_{cav}}{2h_{cav}} + \frac{w_{cav}h_{cav}}{4l_{cav}^2}n^2 \tag{4.3}$$

In the above this condition is met when the profile follows a parabolic change in height along the length, since it, in effect, changes the length of the waveguide where the n^{th} mode's propagation constant is real. We can see that this is similar to the method discussed in the previous paragraph, except now it can *also* be used to tune the frequency as well! It is a bit serendipitous that we can solve *both* the mode spacing *and* field profile problems at the same time, but it is a convenient result nonetheless.

4.3 Flute cavity construction

4.3.1 Cavity fabrication and machining

The flute cavity is machined from monolithic blocks of high purity ($> 99.999\%$) aluminum[2]. A schematic of this process is illustrated in Fig. 4.3. The process—from left to right—is as follows: in Fig. 4.3 (a) we begin by drilling undersized holes that are smaller than the desired diameter ($d = 4.76$ mm) by $\sim 50 - 100$ μm, along the top of the stock at our desired spacing

2. Sourced from Laurand and Associates

($l_s = 6.03$ mm). We use a standard uncoated jobbers style drill bit with parabolic flutes for reduced galling and wear. After the top holes are drilled to their desired depth, the stock is flipped and the process is repeated with the holes displaced by half the center-center hole distance from the top holes, as shown in Fig. 4.3 (b). The overlap of these holes form the cavity volume, depicted in white. An additional reaming step is performed to bring the holes to their final dimension, and to ensure surface uniformity and hole straightness, as shown in Fig. 4.3 (c). This is repeated for both top and bottom holes. Finally a ball hone, made of silicon carbide abrasives, is used to create a uniformly smooth surface and remove internal burrs that may form during drilling and honing, as shown in Fig. 4.3 (d). This process is repeated three times using hones made of successively finer abrasive media. The ductility and galling properties of the aluminum causes the pores of the abrasive to "fill", leading to a drop in honing efficacy after many successive holes. It is thus recommended to inspect the surface finish regularly at this point, and to use a lubricant to reduce heat caused by friction. This is repeated for the top and bottom holes. Special care should be taken to remove any hanging burrs that may be left after machining.

4.3.2 Cavity etching and surface preparation

The machining of the cavity volume introduces defects and impurities into the surface. Additional defects can decrease the effective mean-free path of cooper pairs, leading to longer London-penetration length and more participation of the supercurrent with lossy magnetic defects [175]. To combat this, the top $\geq 100\mu$m of the machined surface is removed via a chemical etching process.

The makeup of the chemical etch[3] is a combination of phosphoric (H_3PO_4), nitric (HNO_3), and acetic (CH_3COOH) at between $25 - 50°$C. The etch works by attacking the oxide instead of the base metal [194, 147]. Hot phosphoric acid etches aluminum oxide, ex-

3. Transene Aluminum Etch A

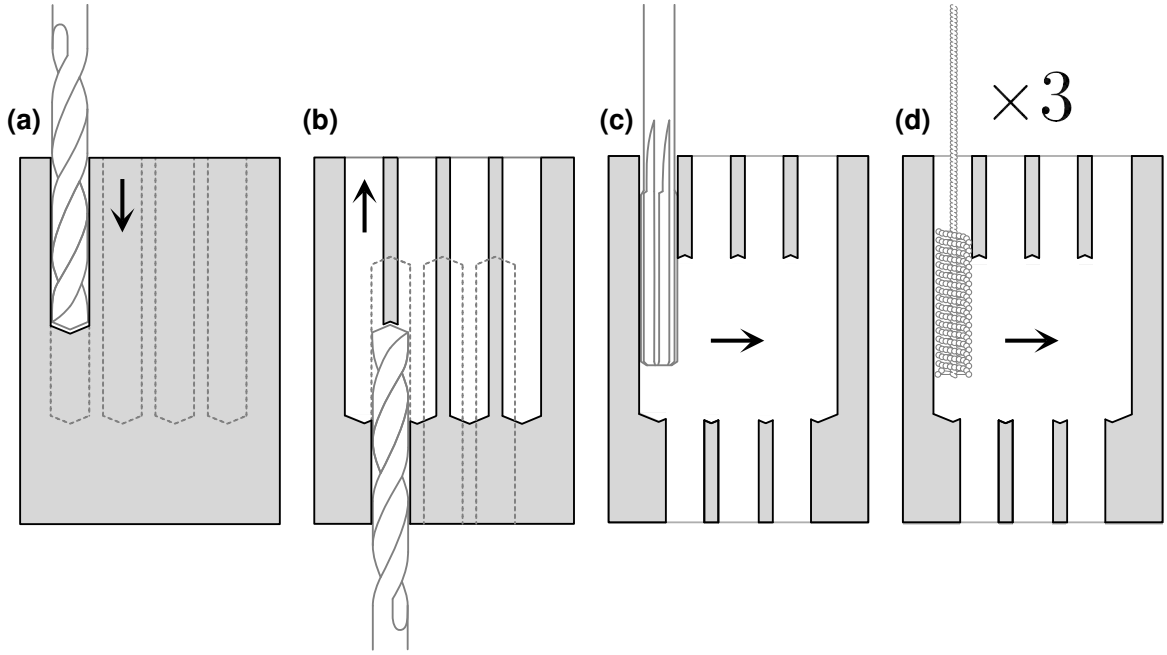

Figure 4.3: **Flute cavity production steps**: The steps involved in making a flute cavity. An initial drilling step involves drilling holes on either side, where the center-center distance of the opposing holes is less than the hole diameter. A subsequent reaming step brings the hole to its final size while also removing deeper surface imperfections from the initial drilling. Next, a series of homing steps smooth the surface further and soften sharp edges while removing burrs.

posing the base material, before the now clean surface is reoxidized by nitric acid. The cycle will continue until the amount of dissolved reaction products in solution slows the reaction kinetics.

The reaction is exothermic, with the rate of reaction being dependent on the heat of solution, with agitation also speeding up the reaction rate. In addition to this, the total amount of generated heat also scales with the surface area. Because of this, in our experience, larger cavities can be etched with lower applied temperature at the hot plate. A picture of a typical etching setup is provided in Fig 4.4.

Over the course of this, and later investigations, the aluminum etching process has been refined. In the following data, all cavities were etched on a hot plate set at 50°C, however later etching for the next generation multimode cavities showed that the applied heat and

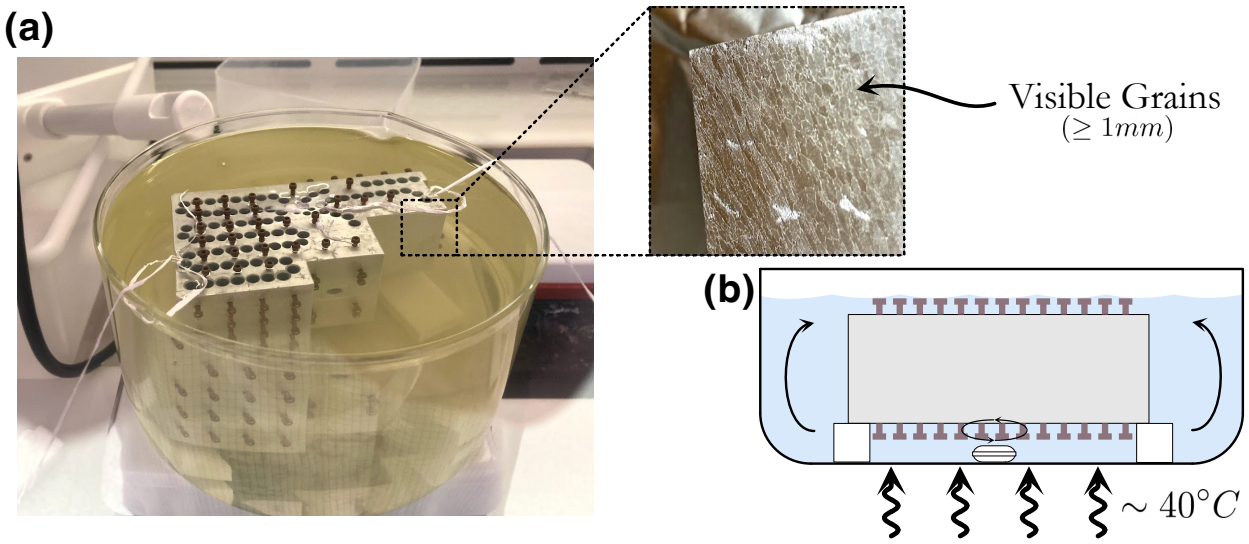

Figure 4.4: **Aluminum cavity etching setup**: (a) shows the newer manipulate-multimode cavity undergoing etching. The etching setup consists of a magnetic hot plate to provide heat and constant agitation. The specified rate of etching is between $30 - 80$Å$/s$ between 30 and $40°C$. The brown screws are made of a PEEK material to protect the delicate threads of the mounting holes from being etched. In the older MM1 and MM2 cavities, these screws weren't present leading to degradation of these small threads. (b) shows a schematic of the whole setup for clarity. The inset shows the sign of a successful etch, with large $>$ 1mm grains visible.

actual heat deviated substantially. In reality, for the vessel size used in many of the etching processes, the bath temp was closer to 30°C.

Following a thorough degreasing and TAMI[4] clean under a range of sonication frequency to remove any residual oils and foreign particulate, the cavities were placed in the etching bath. The following cavities that will be discussed were all etched for 4 hours presuming an etch rate of ~ 80Å$/s$, to give a total etch of $\sim 110\mu$m. The etchant was replaced either every two hours, for the smaller cavities, or every 30min after the first two hours for the larger multimodal cavities. It was seen that coherence was in many cases enhanced by conducting two full scale etches, suggesting a deeper damaged layer. In follow up studies to the etch rate, it was found that the actual etch rate was slower than predicted, further corroborating both the temperature deviation and the performance improvements after follow-up etches.

4. Toluene, acetone, methanol, and IPA

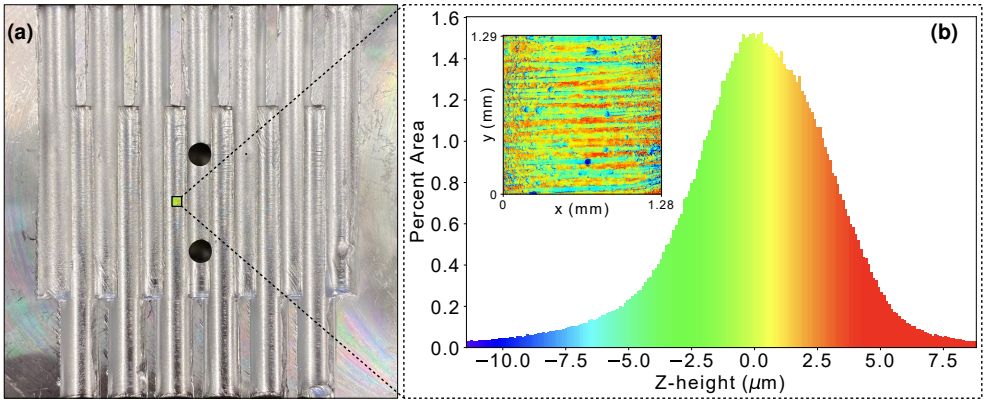

Figure 4.5: **Surface roughness of pan-pipe flute cavity after etch**: (a) shows the characteristic grains of a high-purity aluminum cavity after a successful $> 100\mu$m etch showing very clear grains at the outer surface. (b) shows the inner surface of a pan-pipe style flute cavity. An Olympus OLS LEXT confocal microscope was used to interrogate surface roughness, showing an RMS surface roughness $S_a \sim 2 - 3.6\mu$m over five locations. The etchant does not dramatically smooth over larger defects caused by machining, however does produce a smoother finish than an unetched surface. Since the size of these defects are fare smaller than the characteristic wavelength of the mode, and make only a small fraction of the mode volume, it is not expected that the surface roughness at this scale makes a large contribution to the mode coherence. With the rougher surface, however, we do expect slightly more surface to volume participation, which could enhance existing surface losses.

Fig. 4.5 shows the surfaces of the flute cavity following etching. The inner surface still shows bands characteristic of tooling marks, however finer surface features do seem to be smoothed out by the etching process. The post-etch roughness is far smaller than the wavelength of the mode, meaning that deviations do little to perturb the frequency or induce scattering, however the increased surface area can also increase the surface to volume participation of the modes

4.4 Flute cavity measurements

4.4.1 Overview of cavities tested

For determining the effects of manufacturing, cavity type, and material quality, a series of single-mode cavities were made. All of the cavities made an tested are listed Fig. 4.6. Three

rectangular cavities were made. Two were simple single-mode TE_{101} cavities (pan-pipes) made of 5N (99.999%) pure aluminum (R5N), and on made of 6N (99.9999%) pure aluminum (R6N). In addition to this, an additional 6N cavity was constructed and measured with a qubit. This cavity had a storage resonator with a frequency of ∼6GHz and a readout cavity of ~ 8.1GHz.

To test the efficacy of the flute technique on other geometries, a cylindrical cavity was constructed out of 6N aluminum (Cyl6N). The construction of this cavity had the holes drilled radially along the circumference instead of vertically. This cavity was coupled to the TM_{010} fundamental mode of the cavity.

In addition to the flute style cavities, a traditional 5N aluminum coaxial cavity was also tested with a center frequency of ∼6.6GHz. This cavity was used as a baseline and as a point of comparison between the flute geometries.

Two different multimodal flute geometries were constructed (MM1 and MM2). The primary difference between the two was the inclusion of a 3D readout made also using the flute method on MM2, whereas MM1 was originally intended to have a stripline readout located on the qubit chip itself. MM2 was also constructed in-house alongside the other flute style cavites using the steps illustrated in Fig. 4.3 whereas MM1 was machined by an external shop.

4.4.2 Single-mode resonator measurements

To determine the effect that the design and manufacturing processes had on intrinsic quality factor, the single-mode cavities (with the exception of PP(6N)) were measured using frequency-domain network analysis. A diagram of the measurement setup can be seen in Fig. 4.7 (a). The setup allowed the cavities to be measured in both reflection (S_{11}) and transmission (S_{21}), allowing for more flexible determination of the cavity frequency f_0 and to extract both coupling Qs and internal quality factor. The frequency response of the cavity

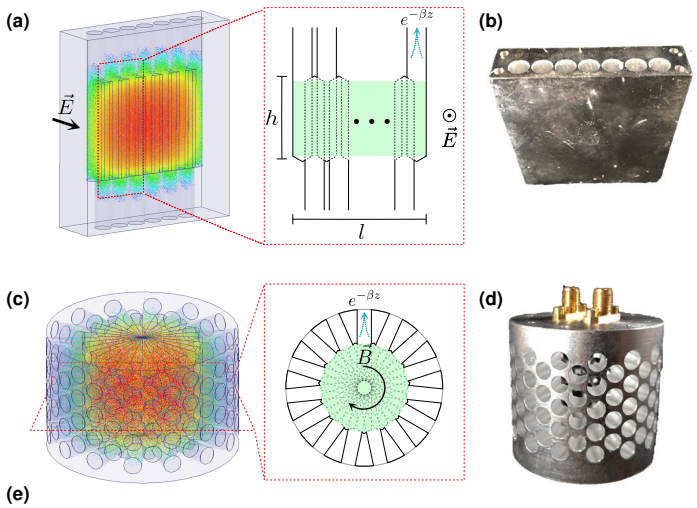

(e)

Device	Mode	f_0(GHz)	Q_{int}	S_e(1/m)	S_m(1/m)
R(5N)	TE101	6.84	1.15×10^7	409	505
R(6N)	TE101	6.75	9.73×10^6	"	"
P(6N)	TE101	6.011	2.1×10^7	420	550
MM1(5N5)	TE10n	$7 - 8.8$	$2.8 - 6.2 \times 10^7$	$420 - 422$	$533 - 543$
MM2(5N5)	TE10n	$5.4 - 6.9$	$6.5 - 9.5 \times 10^7$	$398 - 403$	$513 - 537$
Cyl(6N)	TM010	6.389	2.54×10^7	64	187
Coax(5N)	$\lambda/4$	6.602	9.79×10^7	577	490

Figure 4.6: **Outline of seamless *flute* cavity design**. (a) An FE model showing the $\vec{E}$-field magnitude for the TE_{101} mode of a rectangular waveguide cavity. (inset) A side-view cutaway of the flute design highlighting the overlapping holes, with the effective mode volume highlighted in green. The evanescent decay through the holes is also shown, where β is the waveguide propagation constant for the TM_{0m} modes of the hole. (b) A picture of the R(5N) cavity. (c) An FE model of a cylindrical style flute cavity showing the $\vec{E}$-field magnitude for the fundamental TM_{010} mode. (inset) A top-view cutaway showing the effective mode volume created by the hole overlap. (d) A picture of the C(6N) cavity. (e) A table outlining the performance of various cavity geometries, highlighting the internal quality factors (Q_{int}), and the magnetic (S_m) and electric (S_e) participation ratios from FE simulations.

in reflection (S_{11}) and transmission (S_{21}) are given by Eqs. 2.12 and 2.14 respectively. We can extract the both coupling rates κ_1 and κ_2 and the internal loss rate κ_{int} by doing simultaneous reflection and transmission measurements and fits. We can subsequently compute the coupling ($Q_{ext} = \omega_0/\kappa_{1,2}$) and internal ($Q_{int} = \omega_0/\kappa_{int}$) quality factors.

Fig. 4.7 (b) shows a survey of measured cavity quality factors. The measured values, shown by a red circle, are compared with their *theoretical* performance assuming the intrinsic residual resistance, R_s, and dielectric loss tangent δ, of the cavity under investigation is equivalent to Coax5N's. This is represented by a blue bar, where the range represents if the

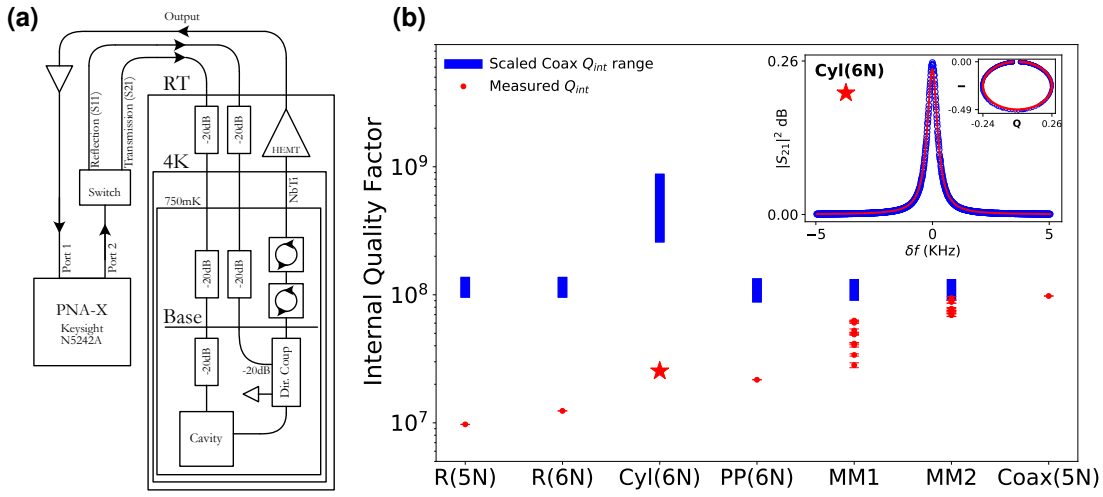

Figure 4.7: **Linear resonator measurement setup and measured Q comparison**: (a) A depiction of the measurement setup for the measurement of cavity quality factors without a qubit. S_{21} and S_{11} measurements were done concurrently by using a directional coupler on the output line of the cavity. (b) A comparison of quality factors for various flute geometries showing measured quality factors (red circles) and the theoretical quality factor range (blue bars) as determined by the scaled coaxial cavity [Coax(5N)] quality factor by the respective values for S_e and S_m values. A plot showing the cylindrical flute cavity [Cyl(6N)] resonator spectroscopy in transmission (S_{21}) (inset).

loss was wholly attributed to dielectric, magnetic, or some combination of the two. Cyl6N's spectroscopic line is shown in the Fig. 4.7 (b) inset, showing both $I-Q$ and magnitude data and fits.

Also included in the above is MM1, MM2, and PP(6N), however, it should be noted that those cavities were measured using the qubit, and will be discussed in more detail in a later section.

The above shows relatively high discrepancies between the theoretical performance range and the realized value. Of the cavities measured, the quality factor of the cylindrical flute deviated the most, having the lowest numerically determined magnetic participation ratio for its TM_{011}, leading to a measured Q_i nearly an order of magnitude lower than the predicted Q_i, assuming the same superconducting properties of the Coax5N cavity. Fig. 4.8 shows the temperature dependence of both the internal quality factor and mode frequency as a

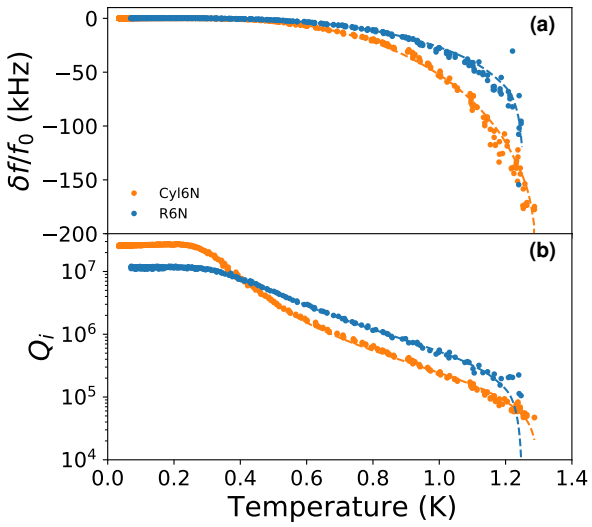

Figure 4.8: **Frequency and quality factor temperature dependence of cylindrical $\mathbf{TM_{011}}$ mode**: (a) The change in the TM_{011} fundamental mode frequency shift as a function of temperature for the 6N cylindrical flute cavity and the 6N rectangular pan-pipe cavity (depicted in Fig. 4.7(b) by a star and R6N respectively). (b) The subsequent internal quality factors of the TM_{011} mode of the cylindrical cavity and TE_{101} mode of the pan-pipe cavity. These fits (dashed lines) give a $T_c \sim 1.31$K and a $T_c \sim 1.25$K for the cylindrical and rectangular cavity respectively. From the frequency fits, we can extract the London penetration lengths for both cavities, with values of 235 ± 3 nm and 37.6 ± 0.9 nm for the cylindrical and rectangular cavity, respectively.

function of temperature for the fundamental modes of both the rectangular *pan-pipe* [R(6N)] and cylindrical cavity [Cyl(6N)]. Both cavities were measured in the same cooldown using separate but identical measurement chains as outlined in Fig. 4.7 (a). The large shift in frequency of the cylindrical cavity from $T = 0K$ to T_c indicates a high kinetic inductance fraction and penetration depth. To determine the London penetration depth and scale of scattering to coherence length, the quality factor data shown in Fig. 4.8 was fit using:

$$Q_i(T) = \left(\frac{1}{Q_{i,\max}} + \frac{p_{\text{mag}}}{Q_{mag}(T)} \right)^{-1} \tag{4.4}$$

Where Q_{mag} is given by Eq. 2.54, and p_{mag} is the geometric magnetic participation, which is described in Eq. 2.48 (b). $Q_{i,max}$ is the maximum internal quality factor at $T \ll T_c$

and $\hbar\omega \ll \Delta$, where Δ is the superconducting gap. The frequency shift shown in Fig. 4.8 (b) is given by:

$$\frac{\delta f(T)}{f_o} = p_{\mathrm{mag}} \left(\frac{\delta \sigma_2(\omega, T)}{\sigma_2(\omega, 0)} \right)^{\nu} \tag{4.5}$$

Where the term in parenthesis is taken from the ratio of reactance at $T = 0$ and some finite temp. The BCS reactance was defined in Eq. 2.25b (b). $\sigma_2(T)$ is the imaginary component of the BCS conductivity and ν is a scaling parameter based on the ratios of the mean-free path (l) and the coherence length (ξ_o). For aluminum, the mean free path is typically much shorter than the coherence length ($\xi_o \gg l$), therefore $\nu = -1/3$ (the "dirty" or Pippard limit) [125]. With $p_{\mathrm{mag}} = \lambda_L S_m$, we can extract the london penetration length λ_L form the fit and from the known geometric participation ratio. Using the techniques used in Ch.3, we calculate a S_m for the Cyl6N and R6N cavities of $\mathrm{S_m} = 187\mathrm{m}^1$ and $\mathrm{S_m} = 505\mathrm{m}^1$ respectively.

For the cylindrical cavity, $\mathrm{S_m} = 187\mathrm{m}^{-1}$, and $p_{\mathrm{mag}} = 4.61 \times 10^{-5}$, giving a $\lambda_L = 235 \pm 3$nm. In the Pippard limit, the London penetration length scales as $\lambda_L \approx \lambda_o(1 + \xi_o/l)^{1/2}$, where $\lambda_o = 16$nm and $\xi_o \sim 1600$nm are the intrinsic London and coherence lengths for aluminum, respectively [175]. This indicates that the mean-free path is much lower than the coherence length, meaning that the superconductor is extremely "dirty", leading to more of the supercurrent interacting with magnetic defects of the material. In contrast, doing the same analysis as above for the R6N flute cavity gives us a London penetration depth of 37.6 ± 0.9nm. This value of λ_{L} is nearly an order of magnitude lower than that of the cylindrical cavity, showing that the surface defects are much better mitigated in the 6N rectangular flute design.

While the maximum quality factor ($Q_{i,\mathrm{max}}$ in SEqn. 4) at base temperature—where thermal quasiparticle formation should be zero—is still higher for the cylindrical cavity, the additional field penetration means that the cylindrical cavity is still more susceptible to

loss due to non-equilibrium quasiparticle formation. An additional contribution to $Q_{i,\max}$ is electric field participation with lossy surface oxides. The TM_{011} mode of the cylindrical cavity has a participation ratio of $S_e = 64\ m^{-1}$ compared to an S_e of 409 m^{-1} for the rectangular design. This disparity likely explains the differences in $Q_{i,\max}$ at $T \ll T_c$.

The implications of the above results highlight an important consideration when designing flute cavities. In both cases the overlap of holes create the cavity volume, however in the rectangular cavity the surface created by the drilling operation, and the subsequent honing steps used to smooth that surface, directly participate in the field, whereas the cylindrical cavity's participating surfaces are left untouched by the subsequent smoothing and refining steps, leaving the participating surface with defects. These defects lead to larger surface areas which can affect the etching process. In niobium cavities surface roughness has been shown to negatively impact etching efficacy and quality factor, however in aluminum coaxial cavities the connection between surface roughness and performance after etching is not as strongly correlated [133, 106].

4.5 Multimode Flute performance

Unlike the aforementioned cavities, which used traditional network analysis techniques to interrogate the cavity performance, the multimode flute cavities, MM1 and MM2 had their coherences measured directly using the dispersive interaction between the qubit and cavity modes. In addition to this, PP6N was also measured with a qubit, which is discussed in far more detail in [63, 62]. This section will briefly discuss the process by which the mode coherences were measured in the multimode system, and show the result of the aforementioned engineering in ensuring consistent coupling to the qubit, and the preservation of mode-coherence with a qubit. This section is really an abridged explanation to the work in Ref. [37], and is covered in far more detail in the as yet published thesis by Kevin He, who largely undertook the process of optimizing and applying the photon-blockade technique for

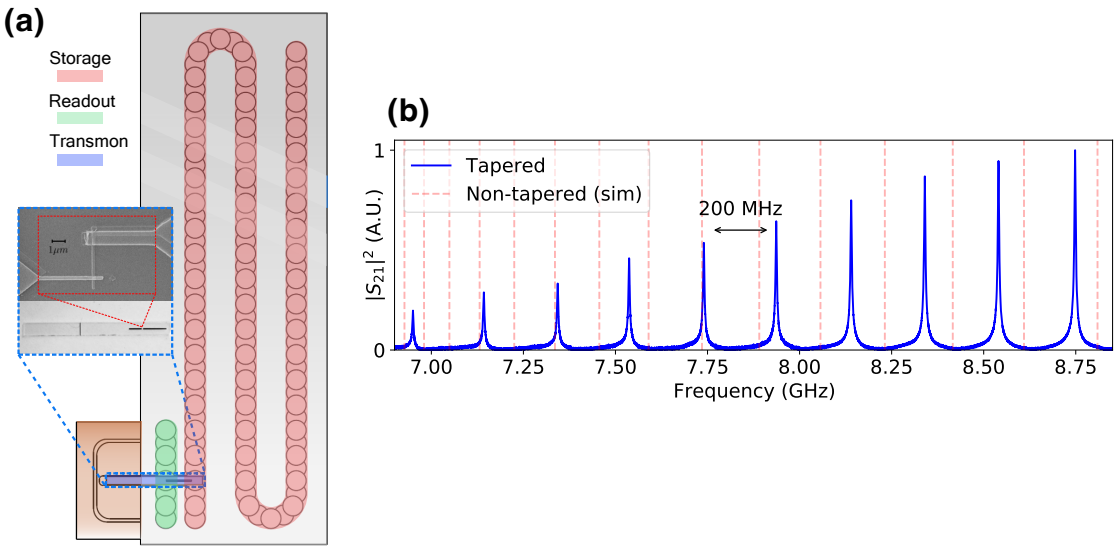

Figure 4.9: **Multimode flute cavity with a transmon**. (a) A schematic of the multimode flute cavity [MM2(5N5)] showing the location of the storage cavity (red), readout cavity (green), and transmon chip (blue). (b) the measured spectrum of MM1[5N5] at room temperature (blue) vs the predicted spectrum for the single-mode case, showing good agreement with simulations from Fig. 4.2. MM2 was designed with a lower mode spectrum as shown in the table in Fig. 4.6 (e)

probing interesting multimodal quantum optics phenomena. With this in mind, the following is provided more for completeness in our discussion of flute cavity performance, and setting up the motivation behind the improved mode coherences discussed in later chapters.

4.5.1 Qubit-cavity device description

The multimode flute consists of a primary storage cavity which follows the mode engineering described in detail in Sec. 4.2.2. A secondary low-Q single-mode flute cavity attached to one side is used as a readout cavity. A 3D transmon qubit (see detail in Fig. 4.9) is made of two capacitor pads spanning across either cavity. Black box techniques described in Sec. 3.4.2 and Sec. 3.5 were used to determine the qubit anharmonicity and individual mode participations and dispersive interaction strengths χ_m for given junction and pad parameters. The qubit and storage cavity are controlled via drives applied through the readout cavity.

The qubit and pads are both made of aluminum that has been deposited via e-beam

Parameter	Quantity	Value
Transmon frequency	$\omega_q/(2\pi)$	4.99 GHz
Storage cavity frequencies	$\omega_m/(2\pi)$	see Fig. 4.6 (e)
Readout frequency	$\omega_r/(2\pi)$	7.79 GHz
Readout dispersive shift	$\chi_r/(2\pi)$	1 MHz
Storage mode dispersive shifts	$\chi_m/(2\pi)$	see Fig. 4.13
Storage mode self-Kerrs	$k_m/(2\pi)$	”
Storage mode cross-Kerrs	$k_{mn}/(2\pi)$	”
Transmon $\|e\rangle \rightarrow \|g\rangle$	T_1^q	86 ± 6 μs
Transmon $\|g\rangle - \|e\rangle$ dephasing	T_2^q	58 ± 4 μs
Readout linewidth	$\kappa_r/(2\pi)$	0.52 MHz
Storage mode relaxation	T_1^m	~ 2 ms, Fig. 4.14
Transmon thermal population	$\bar{n}$	1.2 ± 0.5 %
Storage mode dephasing	T_2^m	$\sim 2-3$ ms, see Fig. 4.14

Table 4.1: **Multimode cQED system parameters**

evaporation in a Plassys evaporator. The evaporation was done on a 430μm thick sapphire wafer cut with C-plane (0001) orientation. Prior to evaporation the wafer underwent an oxygen annealing step at 1200°C for 1.5 hours to recover the sapphire surface to improve the metal-dielectric boundary layer and improve dielectric properties [182]. Following this all wafers were cleaned using a TAMI (toluene, acetone, methanol, and IPA) process under a range of sonicating frequencies ($\sim 30-120$kHz). Following this a 120nm aluminum baselayer was deposited at a rate of $1A^{\circ}/s$. AZ MiR 703 photoresist was applied and the capacitor pads were written with a Heidelberg 408nm UV direct writer and developed for 1 min in AZ MIF 300 1:1 developer. The metalized layer was then etched using chlorine based reactive-ion etching using a plasma made of 30 sccm Cl_2, 30 sccm BCl_2, and 10 sccm Ar. Alignment markers were made by patterning Au alignment crosses, with the resist stripped in 80°C N-Methyl-2-pyrrolidone (NMP) for 4 hours. An MMA-PMMA bilayer resist was then applied before e-beam lithography was used to pattern the junction which was applied

using a Manhattan-style process. This was done using a 100keV Raith EBPG5000 e-beam writer. The resist was developed in an IPA:DI (3:1) mixture for 1.5 minutes. The junctions were evaporated in the Plassys, with an argon milling step used *in-situ* to reduce the oxide layer of the aluminum pads to ensure good galvanic contact. The following 3-step double-angle process was done with a 35nm, 1.0nm/s, 29° deposition followed by an oxidation step with 80:20 Ar:O_2 at 20mBar exposure for 12 minutes, followed by a final Al deposition of 120nm, 1nm/s, with the stage rotated at 29° to be orthornormal to the first angle. The samples were then diced, and the remaining Al and resist layer was lifted off in an 80°C NMP bath for 3 hours.

The qubit is held in place and thermalized using a two-piece copper clamp. A piece of indium wire is used as a compliant washer to tightly hold the chip in place. The readout coupling and drive ports are coupled to the readout via an evanescent hole. The coupling to mediated by a dipole coupling between the TE_{101} mode of the readout and the TM_{01} mode of the circular waveguide to the TEM mode of the coaxial coupling pin. This coupling is different than the coupling of the holes which make up the flute cavity, however the basic scaling of the mode participation with the coupler pin follows the same exponential scaling discussed in Sec. 2.3.3. A detailed discussion of the coupling through this means is provided in Refs. [149] and [31]. As such HFSS simulation is used to predict the coupling to the mode versus pin length by setting the drive port impedance to 50Ω and determining the quality factor. For a simulation with perfect cavity boundaries the $Q_0 = Q_{ext}$, since the port represents the only loss in the system. A full breakdown of the parameters for the system are provided in Table 4.1.

Both readout and storage cavity spectroscopy is done through the readout coupling ports, with the storage cavity measured in transmission. The storage cavity spectroscopy is done by driving at the storage frequencies off-resonantly through the readout cavity and doing a resolved π-pulse on the qubit peak at $|0\rangle$. Populating the storage cavity with $\bar{n} \gg 1$ photons

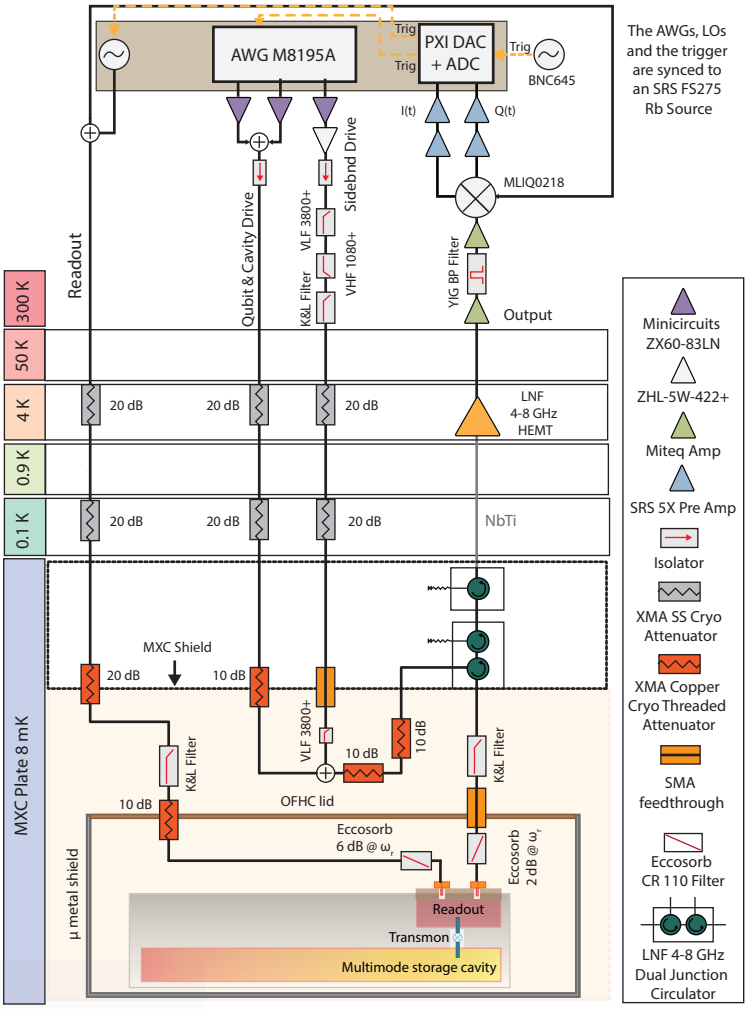

Figure 4.10: **MM2 measurement setup**: A full schematic of the cryogenic and room temperature measurement setup used in the measurement of MM2, diagram courtesy of S. Chakram

would lead to a shift much greater than the qubit linewidth for $2\chi > \Gamma$. The coherent drive strength calibration of the storage modes can then be done by doing number-resolved spectroscopy on the various cavity modes [163]. This will also give the dispersive shift, χ_m, for each storage mode with the qubit.

4.5.2 Measurement setup

The full measurement setup for MM2 is presented in Fig. 4.10. All drives are done through the readout, with each channel being combined from the output of a 64GSa/s Keysight M8195A arbitrary waveform generator (AWG). The return signal is passed through low-loss NbTi cables to 4K, where a HEMT amplifier provides $\sim$ 40dB of amplification. Three stages of isolation at the MXC plate prevents thermal noise from the amplifier effecting the readout thermal occupation. The return signal is further amplified at room temperature before being passed through a narrow band YIG filter to prevent the pre-amp from saturating. The signal $I-Q$ quadratures are obtained via a homodyne measurement setup via signal demodulation through an $I-Q$ mixer and local oscillator (LO) at the readout frequency. This demodulated signal is passed through a final pre-amplifier before being passed into the ADC (Keysight PXI).

The measurement setup used for qubit measurements in the multimode and picollo flute cavities (MM1, MM2, P6N), differed from the setup depicted in Fig. 4.7 (a). Special care has to be taken to avoid excess photon occupation in the readout and storage cavities. Small fluctuations in excess photons in the readout cavity leads to small stark-shifts in the qubit, shifting the frequency by the state-dependent shift. This leads to shot-noise dephasing of the qubit. Because the storage cavity is hybridized with the qubit, it also experiences an effective dephasing rate. This rate can be expressed as follows [48, 153, 149]:

$$\Gamma_\phi = \frac{\gamma}{2}\mathrm{Re}\left[\sqrt{\left(1+\frac{i\chi}{\gamma}\right)^2 + 4i\chi\frac{\mathrm{n_{th}}}{\gamma}} - 1\right] \tag{4.6}$$

Where $\gamma = 1/T_{1,q}$, $\mathrm{n_{th}}$ is the excess or "thermal" photon occupation, and χ is simply the state-dependent dispersive shift between the qubit and the resonator. In the strong-dispersive limit, where $\chi \gg \gamma$ it can be shown that the above expression reduces to $\Gamma_\phi \sim \mathrm{n_{th}}\gamma - \mathcal{O}(\gamma/\chi)^2$.

One may *think* that simply operating at mK temperatures is enough to shield the cavity

from thermal photons, however it is not that simple. It is commonly the naive assumption of quantum physicist that all components are perfectly isothermal with their environment, however constant heat-load from parasitic thermal conduction, and imperfect thermal conduction via highly insulating materials leads to situations where the equilibrium temperature of a component within the measurement chain to be *higher* than expected. A simple example is an attenuator. A typical attenuator is dissipative, meaning energy is dissipated through a resistive network. At low powers and higher temperatures, this heat is minuscule, and readily thermalized to the environment via phonon conduction. For most metals, with Debye temperatures ranging from order $T_D \sim 10^1 - 10^3$K, at temperatures $T < T_D/10$, the phonon conductivity $\sigma_p \propto T^3$, meaning that at mK temperatures, there is little to no phonon occupation [141, 13]. At these temperatures, heat is primarily carried through electron conduction. The relationship between electrical conductivity and thermal conductivity, known as the Weidmann-Franz law, has important implications. For insulators, like sapphire, with $T_D \sim 1050$K, they become thermally insulative at such low temperatures, making the thermalization of non-equilibrium particle in the superconductor of a qubit relatively slow. For a dissipative attenuator, the only connection to ground, and subsequently to the cryogenic environment, is through a resistive element. Under constant or RF power, these attenuators can reach equilbrium temperatures not equal to the base temperature, where $\Delta T = \dot{Q}_{RF}/\sigma_{att}$, where σ_{att} is the conductivity of the attenuator. As such, it is not unreasonable to see how an attenuator, or any RF component, can be at elevated temperature under constant (or near constant) RF load. There will also be higher frequency photons that can cause quasiparticle formation in the superconductor itself. These photons, which have $\hbar\omega \geq \Delta_s$, where Δ_s is the superconducting gap, may come from higher temperature stages. For these, low-pass filters with high-frequency cutoffs in excess of 10s of GHz, are used. These high-frequency adsorbtive materials, namely Eccorsorb, utilize the skin effect or hysteretic loss mechanisms for attenuation.

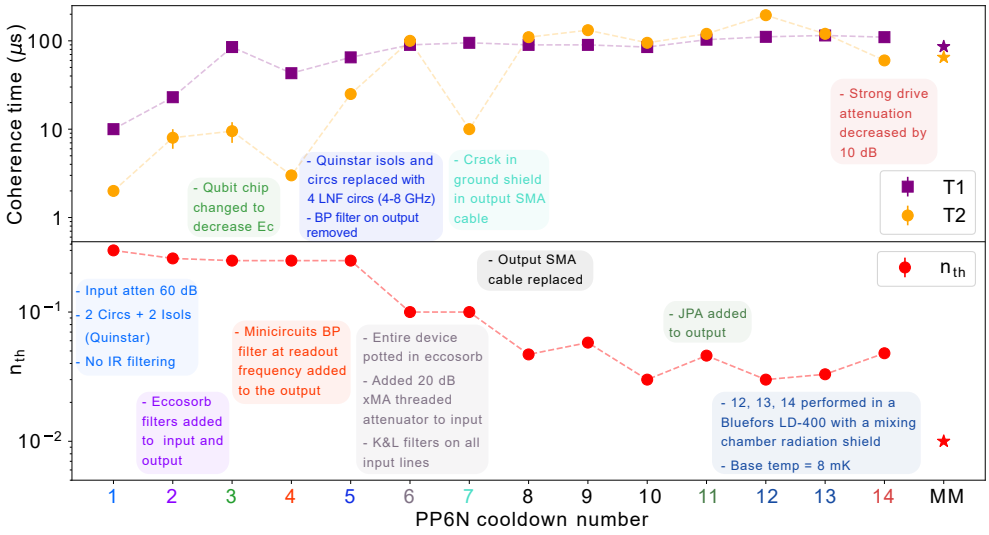

Figure 4.11: **Shielding effects on $T2^*$ of P6N cavity**: An annotated graph showing improvements in qubit thermal occupation (bottom) and qubit coherence time (top). The additional thermal shielding and IR mitigation are listed in the annotations. The largest contributions to improved performance was better IR mitigation using dissipative low-pass filtering (ECCOSORB), more circulators placed inline, and the addition of an 8mK shield. Plot courtesy of S. Chakram. More information on the effects of these improvements in performance can be found in Ref. [62]

For microwave frequencies, an element at 10s of mK can still emit blackbody radiation that is in-band with the device. The exponential aspect of Planck's law means that a small increase in the temperature of any part of the chain can lead to significant changes in the effective power-density spectrum at a given frequency. For example, a black-body source at 100mK will irradiate $\sim$ 1500 times *more* power at 6.5GHz, than at 30mK. Because of this, special care was ultimately taken to shield the cavity from errant IR radiation by using both reflective and dissipative attenuation. A more detailed discussion on the preventative measures used in the P6N cavity can be found in the thesis of Akash Dixit, Ref. [63]. A summary of those additions and their effect is presented in Fig. 4.11, where (a) gives the coherence time T^1 and T^{2*}, while (b) gives the n_{th} cavity occupation.

Many of these lessons were passed onto the MM2 sample, and were incorporated. Primary care was taken to ensure that light leaks through the qubit clamp or the covers on either

side of the cavity were minimized. The addition of a mixing chamber shield at base of the dilution refrigerator, and the modification of the outer mu-metal shield which had light-tight SMA feedthroughs at the top, along with an interior copper liner. This liner was also painted with Berkeley black to act as an IR adsorber.

Thermal population in the qubit is determined by doing length-Rabi measurements on the $|e\rangle - |f\rangle$ transition with and without a π-pulse of the $|g\rangle - |e\rangle$ transition. By taking the ratio of the contrast of the two measurements it is possible to determine the thermal occupation of the qubit. For the contribution readout occupation the qubit T^2 is used in conjunction with Eq. 4.6, with χ and γ being determined independently from resonator dispersive shift—resonator spectroscopy following $|g\rangle - |e\rangle$ π-pulse—and qubit ringdown respectively. The combination of these factors reduced the qubit thermal occupation of the MM2 sample to $\sim 1.2\%$, as shown in Table. 4.1.

4.5.3 Fock state preparation using photon-blockade

For determining the mode lifetimes, donoted as T^1, or the energy relaxation time, a single excitation is placed in the resonator. This Fock-state of the resonator, with an eigenstate denoted as $|1\rangle$, can be generated through the dispersive interaction between the qubit and the cavity modes. In Ref. [39] we demonstrated three different protocols for Fock-state preparation, including selective number-dependent arbitrary phase gates (SNAP)[87], the $|f0\rangle - |g1\rangle$ sideband [158, 137], and photon blockade [37]. For the following discussion, we will discuss the use of photon-blockade for this state preparation.

We can determine the Hamiltonian for our system by expanding on the technique used in Sec. 3.5.1 to determine the participation of an m target mode with the qubit. Expanding the $\cos\varphi$ to fourth order and writing it in terms of the junction phase and writing the Hamiltonian in the dispersive regime ($|\omega_i - \omega_{ge}| > g_i$) gives [124]:

$$H_{mm} = \omega_q |e\rangle\langle e| + \sum_{m=0}^{N-1} \{\omega_m a_m^\dagger a_m + \chi_m a_m^\dagger a_m |e\rangle\langle e| + \frac{k_m}{2} a_m^\dagger a_m (a_m^\dagger a_m - 1)\} + \sum_{n \neq m} k_{mn} a_m^\dagger a_m a_n^\dagger a_n \tag{4.7}$$

Where ω_q is the qubit transition frequency between $|g\rangle - |e\rangle$ states, ω_m is the frequency of target mode m. The second term in the first sum is simply state-dependent shift, or cross-Kerr, between the qubit and target mode, while the last term in the sum is the self-Kerr of each mode, while the last term is the cross-Kerr *between* modes.

We can add a qubit and cavity drive, denoted as $\Omega(t)|g\rangle\langle e|$ and $\epsilon(t)a$ respectively to the above Hamiltonian. It was shown that to block the $|n_m\rangle$ Fock state of target mode m, one can simply apply a qubit drive at $\omega_q + n_m\chi$, which correspond to $|gn_m\rangle - |en_m\rangle$. The result of this is that the target mode sees a splitting of their n^{th} level that is equal to 2Ω. Obviously, the condition only works for drive strengths where $\Omega \ll \chi$. In addition to this there are small dressings to the neighboring modes between the qubit ground and excited states which scale as $\Omega/(\chi(n_n - n_m))$ where $n \neq m$.

The result of this is that our harmonic oscillator's infinite Hilbert space has now been truncated. This added anharmonicity to the target level means that a resonant drive applied to the cavity will only lead to the excitation of $n-1$ levels of our harmonic oscillator. In the case of blockading the $|g2_m\rangle - |e2_m\rangle$ level, and driving the cavity, the cavity state will simply oscillate between $|0\rangle - |1\rangle$ Fock states. By ensuring the Rabi rate of the cavity drive is less than Ω ($\epsilon \ll \Omega \ll \chi$), such that leakage into the blockade transition doesn't occur, and waiting for π-time, it is possible then to prepare a $|1\rangle$ Fock state with high fidelity.

A $|0\rangle - |1\rangle$ Rabi oscillation from a blockade on the $|g2\rangle - |e2\rangle$ transition is shown in Fig. 4.12 (a) for a given drive strength. To conduct a T^1 measurement, a second resolved π-pulse is enacted, swapping the cavity population back to the qubit, before measurement of

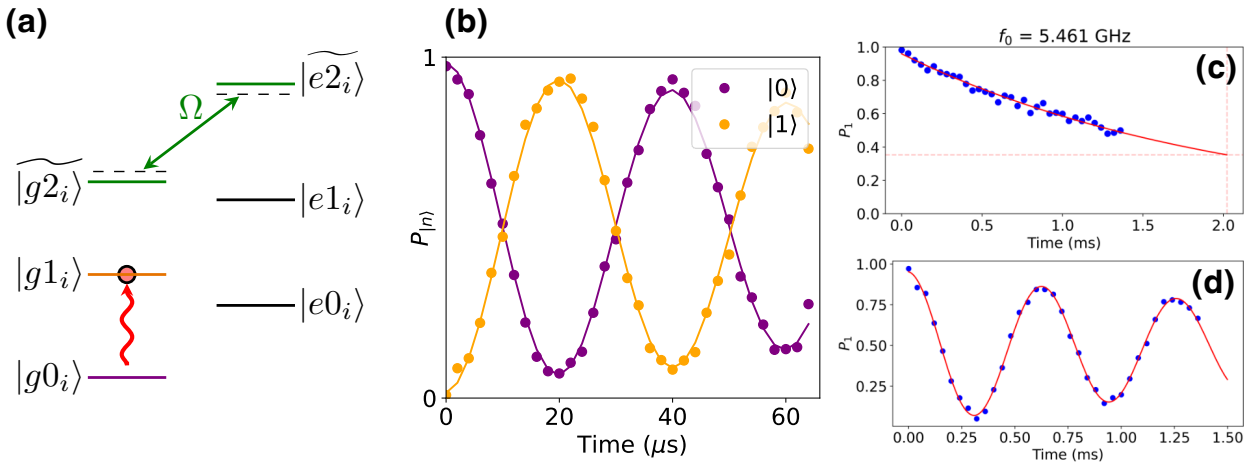

Figure 4.12: **Rabi and coherence using a blockade on $|g2\rangle - |e2\rangle$ transition**: (a) shows a level diagram of the $|g2\rangle - |e2\rangle$ blockade, where a weak drive excludes the occupation of the $|2\rangle$-state of the oscillator. A weak Rabi drive is then applied at the mode frequency, with the drive strength leading to a Rabi oscillation of the number-state, as shown in (b) for the measured $|0\rangle$ and $|1\rangle$ number peaks of the qubit in time. (c) and (d) show T^1 and T^{2*} respectively for mode 1 of the multimode cavity. The values of these coherences are listed in Fig. 4.14 below.

the qubit occurs. For the cavity dephasing rate T^{2*}, the process is similar, with a detuned, $|\omega - \omega_q|$, $\pi/2$-pulse after some time. The energy decay and dephasing time can be extracted to give the characteristic exponential time constants associated with their decay rates.

For this discussion, we will only discuss the creation of single excitation ($|1\rangle$) Fock state on a single mode, however this technique can be more generally applied to multimodal systems to generate entangled states. To see how this is done, see Ref. [37] or see Kevin He's thesis.

4.5.4 Multimode cavity performance

The aforementioned measurement and control scheme were used to determine both χ and the mode T^1 and T^2. As discussed, T^1 and T^2 were determined via photon blockade on the target mode to generate either $|1\rangle$ or $|0\rangle + |1\rangle$ Fock states. χ was determined by doing resonator spectroscopy on the target mode after applying a π-pulse on the qubit, leading to a shift in the cavity resonance that was equal to 2χ. Based on the measured storage mode and

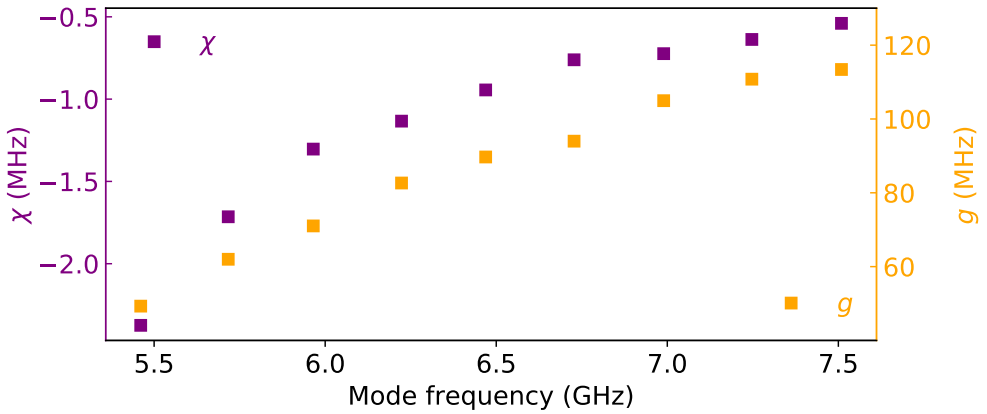

Figure 4.13: χ **and** g **vs mode number**: The dispersive shift, χ, and extracted coupling g, between the qubits and modes. These values are extracted from resolved number splitting of the qubit.

qubit frequencies, and the anharmonicity α of the qubit, it was also possible to determine the vaccum-rabi rate g from Eq. 1.24. The χ and g are plotted in Fig. 4.13. We see that, over the 9 target modes measured, the scaling of the coupling and dispersive shift range from -2MHz (60MHz) and -500kHz (120MHz) for χ and g respectively. The important takeaway from this result goes back to Sec. 4.2.2. One of the reasons for tapering the TE_{10n} cavity was to distribute the E-field such that g, or the qubit and cavity coupling rate would be non-zero (and ideally relatively high). While the data indicates that there is variability across modes, there are no points where the coupling drops precipitously due to the presence of a node in the field profile of any given mode. Here we see that g scales roughly linearly with increasing mode number, while χ follows an almost quadratic scaling, with deviation being due to the increased detuning of the qubit-mode frequencies. Either way, the data indicates that *all* modes are still in the strong-dispersive regime ($\chi > \kappa, \gamma$).

With this in hand, we can also determine the cooperativity of each mode with the qubit. In Fig. 4.14 we plot the aforementioned decay and dephasing (T^1, T^2) of each mode. The blue bands on the T^2 values represent the range predicted by n_{th} and the qubit T^1 to calculate the predicted dephasing rate due to thermal occupation. Over the 9 modes we see a near constant 2ms lifetimes, which correspond to quality factors of $Q_m \sim 70 - 90 \times 10^6$. If taken with the

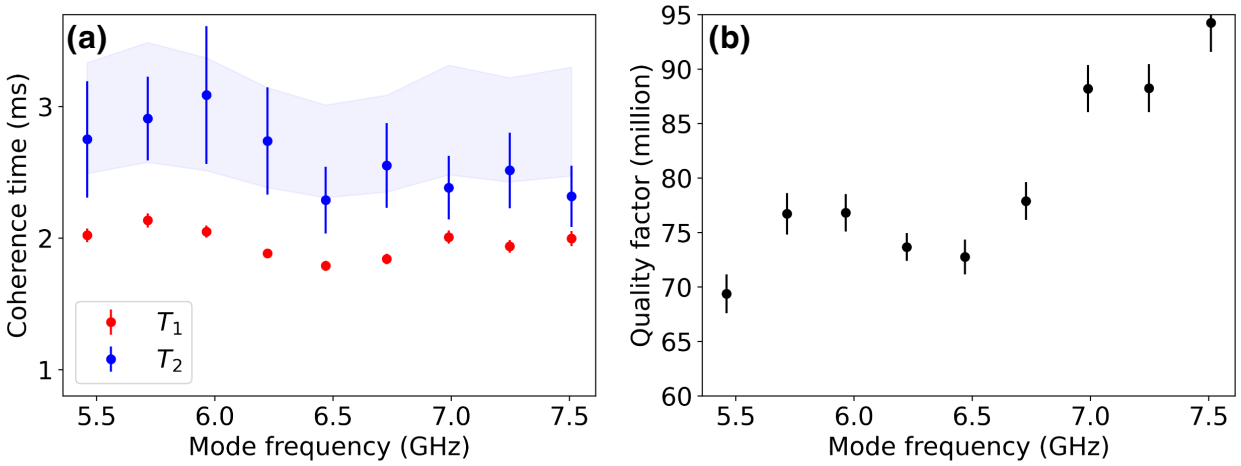

Figure 4.14: **T^1 and T^2 vs mode number**: (a) shows the distribution of T^1 and T^{2*} across the first 9 modes using the blockade technique described in Sec. 4.5.3, while (b) shows the equivalent quality factors for each mode.

measured T^2 values, we can use the mode coherences, with a measured qubit coherence of $T_q^1 = 86 \pm 6\mu$s to get a cooperativity across all modes of $g^2/(2\kappa_m\gamma) \sim 0.5 - 1.5 \times 10^9$.

These results and the discussion of the preceding chapter have showed the power of the flute technique, in particular in the creation of a multimodal 3D cQED system. While we have not gone into great depth discussing the potential for control and preparation of interesting quantum states in this chapter (you will have to wait for Kevin He's thesis to see what can be done with this system, but as a prelude see Ref. [37]), we have taken the next logical step in the refinement of a useful quantum memory or multimode quantum processor. Given that this thesis is invested in the creation of the best quantum memory—single or multimodal—it is only logical that we continue the discussion by addressing the avenues in which we can make this system *even better*.

The two main limitations of this system is that the dispersive coupling between each cavity mode and the qubit ultimately lead to unwanted off-resonant interactions between neighboring modes, i.e. the back-action due to the dispersive shift in the qubit from target mode m occupation, leads to a small shift in non-target mode n, which is captured by an effective intermodal cross-Kerr term as seen in Eq. 4.7 in the form of k_{mn}. This, of course,

leads to a bookkeeping nightmare, with the calibration of these dispersive shifts for any given mode with any given prepared state. At the end of this thesis we will discuss one avenue to solve this by removing the non-linearity, and subsequently the dispersive interaction, altogether. The second point of interest has to do with the scaling of the multimode system. In the next chapter we will begin by outlining this problem more carefully, in discussing the way mode lifetime ultimately effects the number of modes one can utilize for the creation of a useful register, and how building high quality factor cavities can open the door for memories with hundreds, if not thousands of modes.

CHAPTER 5

3D NIOBIUM CAVITIES

5.1 Motivation and Background

5.1.1 Quantum memory scaling and Quality Factor

At the end of chapter 3 we discussed several technical challenges in making scalable quantum random access memories. The first of these challenges, the reduction of off-resonant interactions by minimizing cross-Kerr between non-target modes, may be addressed through the use of a flux tunable coupler that is comprised of a linear Josephson device. This is done by making $E_j \gg E_c$, pushing down the anharmonicity. The other challenge to scaling is the reduction of the idling error rate to allow for a higher number of modes to be multiplexed. To make this point more explicit, lets develop a naive model to understand the nature of this scaling based on the error rates of resonant gates versus the idling error of any individual mode.

The individual fidelity of any gate operation, assuming $T_{1,cav} \gg T_{1,q}$ is:

$$\mathcal{F}_{gate} \propto \exp\{(-t_g/T_{1,q})\} \tag{5.1}$$

Also, the error rate of the prepared state in a mode would subsequently be:

$$\mathcal{F}_{cav} \propto \exp\{(-t_g/T_{1,cav})\} \tag{5.2}$$

Where t_g is just the gate time. If we want to know how many modes can be multiplexed, assuming no loss in fidelity in the stored information, we can compare these two values. We can rewrite $\mathcal{F}_{cav}$ to reflect the loss in fidelity of the stored state over the time it takes to multiplex over n modes as $\mathcal{F}_{idle} \propto \exp\{(-nt_g/T_{1,cav})\}$, which has had the *idle* subscript

added to symbolize this is the fidelity after some idling time. To find $\text{MAX}(n)$ we can set the condition that the idling fidelity must be greater than or equal to the individual gate fidelity, after which the idling error rate will become the dominant source of infidelity. Solving for n gives us the expression:

$$n_{max} \propto \frac{T_{1,cav}}{T_{1,q}} \tag{5.3}$$

This criteria for the maximum number of multiplexed mode is, admittedly, a bit contrived and is only true for resonant gates as described in [124], however proposals for quantum-acoustic qRAM[1] that involve the use of virtual swaps that never populate the qubit, and are therefore only limited by the idling error and the inverse Purcell effect [84]. The downside with these virtual processes are speed, with virtual coupling rates of only $10-100$kHz versus the bare couplings of $60-100$MHz as seen in [39].

Besides the obvious benefits to the fidelity of multiplexed modes, higher quality factors also benefit systems where single mode operations and high cooperativities allow for the creation of non-trivial states which can be used for stabilization and autonomous error correction [130, 92, 34, 152]. With this in mind, and given the concepts previously discussed, what tools exist for the creation of a higher coherence cavity? One solution, of course, is to go to another bosonic system altogether. Phononic systems have been demonstrated with quality factors of $> 10^{10}$ with very small mode volumes [112], and with the advent of techniques for coupling superconducting circuits means that the cQED toolbox for state manipulation and encoding can be used [11, 160, 128, 44, 98, 113]. This does have several advantages, but pose it's own unique problems, with challenges in controlling dispersion and coupling an open research question.

The second option is to look at existing superconducting cavity technology for inspira-

1. qRAM differs from our proposed multimode RAM in that it takes a *quantum register*, or a superposition of addresses, instead of a classical indexing of a mode that contains a quantum state.

tion. A natural choice are superconducting radiofrequency cavities developed for particle accelerators. While these cavities are typically designed to maintain high quality factors at incredibly high electric field gradients, recent investigations into their single-photon performance have yielded promising results, with single-photon quality factors exceeding 10^{10} for 1.3GHz and 5GHz TM_{010} TESLA cavities [157, 156]. Unlike the cavities used in our multimode system, or previous 3D cQED experiments which were made of aluminum (both high-purity and alloy grades), these cavities are made of niobium and utilize a drastically different design philosophy than the TE_{101} rectangular or $\lambda/4$ coaxial cavities previously discussed.

In the following section we will discuss the material properties of these cavities and glean how we can adapt niobium to cavity designs that can be easily integrated with 3D transmon qubits. This will not be an exhaustive overview. For a far more comprehensive review of the accelerator applications and design philosophy, references [133, 134] are indispensable, and will be the primary source for much of the theoretical and historical background in this chapter.

5.1.2 Superconducting Radiofrequency (SRF) accelerator cavities background

To preface the following sections we may want some context for the logic behind the design of accelerator cavities. A charged particle inside of an electric field experiences a force and subsequently acceleration in the direction of the field (for positive particles). Naively this can be achieved by a set of parallel plates, but to accelerate a particle to relativistic velocities ($\beta = v/c \approx 1$), one would need a large field gradient over a very long distance. To negate this one could build a resonant cavity with a transverse electric field. A natural choice for this is the TM_{010} mode of a cylindrical cavity with the center of any two cavities spaced $\lambda/2$ apart. As the particle traverses a $\lambda/2$ distance in $1/2f_{cav}$ time, then the electric field will add

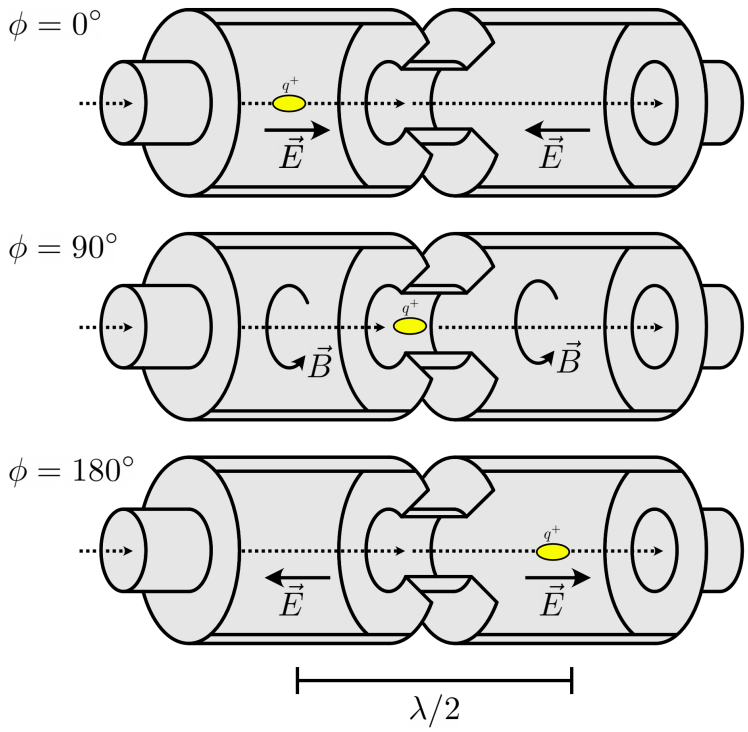

Figure 5.1: **SRF cavity operation and mode profile of TM_{010} π-mode**: At $t = 0$, a particle of charge q^+ is located at the center of the first cell, with the $\vec{E}$-field gradient at its maximum. At $t = 1/4f$, the energy is all stored in the magnetic field, but the particle is at a node in the field distribution. At $t = 1/2f$, the particle is located again in the center of a cell, this time the electric field vector has flipped, and is in the direction of motion. This cycle subsequently repeats as the particle moves through each cell.

constructively, accelerating the particle through the beamline. This process is schematically shown in fig. 5.1 for a simple pill-box geometry, depicting the relative phase of the $\vec{E}$ and $\vec{B}$ fields as the particle moves through a properly tuned pair of coupled cavities. The time-of-flight of the particle, the resonant characteristics, and field-flatness are all important tuning parameters to ensure maximum efficiency, or the ratio of input energy vs accelerating energy experienced by the particle, but how does the quality factor of this system effect the performance?

In Ch. 2 eq.2.3 we defined the internal, or intrinsic quality factor as the ratio of stored energy in the cavity (P_{tot}) over the energy lost (P_{loss}) over one resonant period (ω_0). For accelerators operating in continuous wave mode (CW), the amount of power dissipated per

unit length is:

$$\frac{P_{loss}}{L} = \frac{E_{acc}^2}{R_P} \tag{5.4}$$

Where E_{acc} is the accelerating field potential and R_P is the geometric shunt impedance, or the effective parallel resistance of fig. 2.1 per unit length, since this resistance is no longer discrete. For parallel resistance, $Q_0 \propto R_P$, and to reflect this the above equation is typically written in terms of the ratio of Q_0 and R_P. Nonetheless eq. 5.4 illustrates the idea that higher quality factor is equivalent to higher efficiency, which for accelerators that require field gradients of $E_{acc} > 1\text{MV/m}$ to accelerate charge particles to relativistic speeds, is quite important from an operational cost point of view. This is compounded by the fact that the superconductor, in this case niobium, must be cooled to liquid helium temperatures, requiring closed cycle cryogenics with typical Carnot efficiencies of $\eta_c < 2\%$. So far this description has presumed the cavity performance is agnostic to the accelerating potential, but in reality the quality factor's dependence on the accelerating potential makes up a large part of the active ongoing research in accelerator physics. In the case 3D cavities for cQED, neither this optimization, or the wall-power efficiency, is of much concern, since the operating power at the cavity is, in effect, on the order of a single-photon worth of energy. Instead, we want to focus on *how* we can adapt the optimization of accelerator cavity quality factors to 3D geometries that we can use to do interesting quantum physics.

In our simplified accelerator presented in fig. 5.1 we have drawn our TM_{010} cavity as a simple cylinder. In reality accelerator cavity designs are far more complicated and nuanced, with every one optimized for particular field gradients. All of these designs however belong to a class of TM_{010} known as *elliptical* cavities, named after the ellipses that define their inner and outer surface. For linear electron-positron colliders, where compactness and subsequently high accelerating potentials ($E_{acc} > 20MV/m$) the TESLA design [16], named after the international collaboration that designed it, has become the industry standard. Its

specific geometry, with its circular inner radius, has also been used to build some of the highest internal quality factor cavities ever made [154], reaching quality factors in excess of $Q_0 > 2 \times 10^{11}$ at 1.3GHz.

There are a few reasons why the internal quality factor of such cavities are so high. The first is the material properties of the Nb cavity, which will be discussed in more detail in the intervening sections of this chapter. The other is simply geometry; TESLA or elliptical cavities have far lower participation. A common metric used in the literature on SRF cavities is the geometric factor G, which is given as:

$$G = \frac{\omega_0 \mu_0}{S_m} \tag{5.5}$$

Note that the above is simply $Q_{mag} R_s$ from Eq. 2.52, and is a natural way separating the cavity's geometric contribution to Q and the intrinsic material loss R_s. For typical TESLA cavities, $G \sim 250 - 300\Omega$ [155, 157]. It should also be noted that the above only considers magnetic loss mechanisms, whereas dielectric loss can, and does, have a strong contribution. Confusingly, dielectric and magnetic losses are often both encapsulated into the value R_s, however, certain measurements can give us constrains on the contribution of both mechanisms.

This high geometric factor for SRF cavities however comes at a loss. To achieve this, S_m is made very small, which comes at the cost of making the field more diffused in the cavity volume, with less field near the surface. For an accelerator, this is fine; the charged particles live in the center of the cavity where the field gradient is highest. For qubits, this would mean a particularly long chip to achieve the coupling desired, and potential difficulty in incorporating a readout resonator. By making this chip longer, you also become more susceptible to vibration, since your chip is effectively a stiff cantilever, which can lead to pure-dephasing of the cavity (which will be discussed in a later chapter). Long chips also make qubit thermalization more difficult. The larger chip will also participate more strongly

Geometry	Mode	Material	f_0 (GHZ)	G (Ω)	R_s ($n\Omega$)	Q_0 ($\times 10^9$)
TESLA	TM_{010}	Nb	1.3	268	$1-10$	$10-100$ [157]
Coaxial	$\lambda/4$	Al ($\geq$5N)	6.5	$80-100$	1000	0.07-0.1 [38]
Rectangular	TE_{101}	Al ($\geq$5N)	5.4-7.0	$80-100$	1000	0.07-0.09 [38]

Table 5.1: **Comparison of cavity performance**

in the cavity field, making cavity susceptible to the chip's dielectric loss.

This provides an inherent trade-off in design; make the mode volume smaller, which makes the incorporation of a qubit easier, with the expense of a worst geometric factor. A natural candidate for a cavity design that fits the criteria is the already discussed $\lambda/4$ coaxial cavity, which has become a tried and tested workhorse of the cQED world [148, 75, 187, 158, 159, 106, 107]. The $\lambda/4$ coaxial cavity has typical geometric factors of between $80-100\Omega$ using the values determined from Sec. 3.3.2. This means, for the same R_s as the aforementioned ellitical/TESLA cavities, coaxial cavities will have a $3\times$ lower quality factor.

We can compare the performance of several other cavities that have been discussed, namely the TE_{101} multimode flute, and the coaxial cavities discussed in earlier chapters. We see that, for aluminum cavities, the residual resistivity is between $100-1000\times$ higher than for state-of-the-art niobium SRF cavities. This means, if we were to scale these material values for cavities of similar frequency to that of a coaxial cavity, we can expect internal quality factors between $1-10\times10^9$. As we will see in later sections, the performance of the cavity is dependent on more than the raw residual resistance. Nonetheless we will see that we can atleast *approach* the theoretical performance by switching to niobium.

Before we can discuss performance, we need to outline *how* this performance is actually achieved, and go over some of the main factors in surface treatment and chemistry that can limit or enhance cavity performance for niobium systems. This brings us to the second factor in why SRF cavities have such long lifetimes; the intrinsic properties of niobium itself.

5.1.3 Superconducting properties of Nb

Before we continue, it is informative to spend a moment on some of the unique properties of niobium, in particular its superconducting properties. In Sec. 2.2.4 we discussed BCS theory, which is a microscopic treatment of superconductivity in simple superconductors. We showed in Sec. 4.4.2 that, for a Type-I superconductor like aluminum, BCS theory, and Mattis-Bardeen theory, correctly predicted the response of our system in an AC-field. This microscopic treatment, which is predicated on the spatial homogeneity of the superconducting gap in space, becomes untenable for superconductors with non-trivial spatial dependence. For these systems, a more phenomenological approach is required, which reparameterizes the local density of superconducting electrons as a complex order parameter $|\psi(\vec{r})|^2$ based on a pseudowavefunction $\psi(\vec{r})$. This macroscopic theory, developed by V. L. Ginzburg and L. D. Landau, gives us the tools to describe the properties of Type-II superconductors, such as niobium, using empirically derived properties [77].

This section will *not* be a detailed discussion of how Ginzberg-Landau (GL) theory works, but instead as a pretense for explaining the key differences between niobium and the other superconductor previously discussed in this thesis, aluminum [2]. The main differentiator between the two is how they interact with magnetic fields. In addition to zero DC resistivity, another feature of superconductivity is the requirement of zero magnetic field in their bulk. This property, known as the Meissner effect, results in surface currents that counteract the external field inside the penetration length, λ_L.[3] To do this, work must be done, which in turn increases the thermodynamic free-energy of the Cooper pairs. Inversely, the condensation of cooper pairs leads to a decrease in entropy and subsequently the free-energy of the system. If the induced increase in free-energy from the external magnetic field is equal or greater than the change in free-energy required for condensation, the Cooper-pairs break

2. For more on GL theory see Ch.4 of Ref. [175]

3. Inside the London penetration length, the magnetic field *is not* zero, but decays exponentially as $B(x) = B_0 \exp(-x/\lambda_L)$ for values of $B_0 < H_c$

Material	T_c (K)	Coherence Length ξ_0 (nm)	London Length λ_L (nm)	H_c (mT)	H_{c1} (mT)	H_{c2} (mT)
Aluminum	1.2 [49]	1600 [49]	16 [49]	10 [49]	—	—
Niobium	9.26 [67]	38 [67]	39 [67]	200	$174 - 190$	$390 - 450$ [134]

Table 5.2: **Comparison of Nb and Al superconducting properties**

into normal-state electrons. This critical-field strength, H_c, is the thermodynamic critical field. For a material such as aluminum, where the coherence length ξ_0 is greater than the london length λ_L, the above condition is met, but what about for $\xi_0 < \lambda_L$, as is the case for niobium ($\xi = 38$ nm, $\lambda_L = 39$ nm)?

One result of the $\xi_0 < \lambda_L$ limit is that, below H_c, the superconductor can vary between a normal conductor and superconductor. If these normal regions are thinner than λ_L, then less energy is required for the neighboring superconducting sheets to expel the field, meaning that the field can not just penetrate into the bulk, but do so without breaking cooper pairs in the superconducting layer.

GL theory predicted that, for $\kappa = \lambda_L/\xi < 1/\sqrt{2}$, where κ is the ratio of London penetration length and coherence length, the superconductor exhibits Type I superconductivity, and the field is expelled. For $\kappa > 1/\sqrt{2}$, there exists an additional critical field $H_{c1} < H_c$, where the field is suppressed, and above which we see penetration due to the ordering described above. The nucleation of cooper pairs will continue in the bulk even above H_c, and until a third critical field is reached, H_{c2}, in which all superconductivity ceases in the bulk and at the surface, with $H_{c2} = \sqrt{2}\kappa H_c$, from GL theory.

The properties that distinguish aluminum and niobium are listed above in Table. 5.2. Other than the addition of its more complicated interaction with DC magnetic fields, the other notable feature of niobium is its high critical temperature T_c. For SRF cavities the high operating temperature means that the accelerator cavities can be operated at or near liquid helium temperatures without serious degradation in performance. As we saw in Sec. 4.4.2,

above $\sim 0.2T_c$, the number of equilibrium quasiparticle in the bulk contribute noticeably to the Q_{mag}, leading to an overall decrease in the internal quality factor.

While we have discussed the unique properties of niobium, as far as its superconductivity is concerned, the primary differentiatior, as we will see, is the temperature and power dependence of the quality factor due to two-level system defects in the oxide layer. These aspects will be explored in far more detail in the following sections.

5.2 Nb coaxial cavities manufacturing

We have outlined the background and some of the unique and desirable properties of niobium in the context of superconducting cavities. In this section we will shift from our discussion of Nb at high powers and in elliptical cavity geometries and shift to adapting the aforementioned processes to a coaxial cavity design. We have discussed this design in Sec. 2.3.1 and its impact on the cQED field as a whole. Moving forward we will focus on the design and manufacturing considerations made to properly characterize the loss mechanisms of the Nb surface.

5.2.1 Cavity manufacturing and etching

The cavity was machined out of a piece of pure niobium metal with an $RRR \geq 300$ which conformed to ASTM B393 R04220 standards and was purchased from Ningxia Nonferrous Metals. For more even material removal during etching, sharp edges and corners of the co-axial design were radiused for smooth transitions (see. Sfig 5.3). The cavity was initially degreased using an Alconox© and DI-water mixture at temperatures $> 50°$C under sonication. Following this, the cavity was thoroughly rinsed in DI-water and dried using an isopropyl alcohol (IPA) solvent exchange and dry nitrogen. A toluene-acetone-methanol-IPA cleaning was done using semiconductor (semi) grade solvents under sonication for $\geq$ 5min at each step, before being dried under high-purity nitrogen that had been passed through a 0.2μm filter and bagged in a ISO100 clean-bag. All cavities were stored in a dry box.

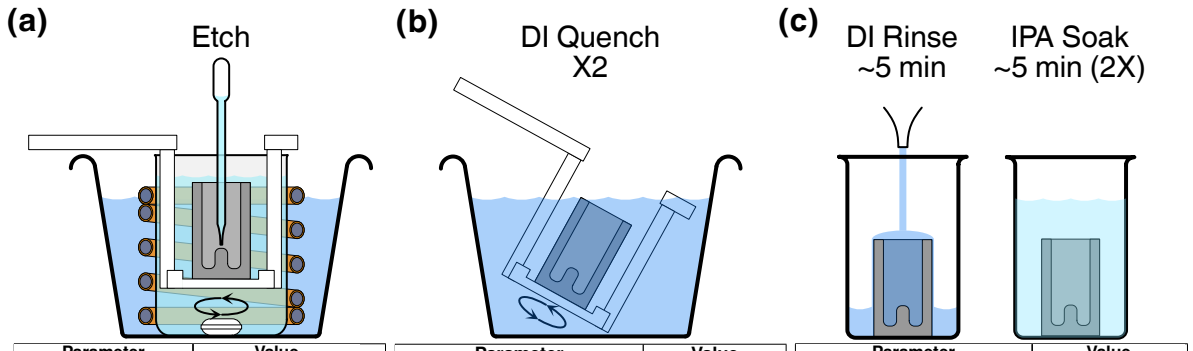

Parameter	Value
Etch	1:1:2 (HF:HNO_3:H_2O)†
Etch depth (μm)	$100 - 150$
Etched area (cm^2)	70.2
Time (min)	$60 - 120^{\ddagger}$
Temp ($^{\circ}C$)	< 10
Etch vol. (mL)	$\sim 500^{*}$
Stir-rate (min^{-1})	1250
Pipette Interval (min)	$5 - 15$

† Sold as RSE1:1:0:2H_2O from Transene Co. Inc.
‡ calibrated *in-situ* for specific etch-rate using Nb test-strip
* determined based on etch depth and surface-area (see text)

Parameter	Value
Water resistivity ($\Omega\ cm$)	$\geq 18.6 \times 10^{6\dagger}$
Water vol (L)	~ 2
Quench time, 1st step (min)	$\sim 1^{\ddagger}$
Quench time, 2nd step (min)	$\sim 5^{\ddagger}$
Temp ($^{\circ}C$)	22

† Water meets UHP Type-I standards using a Barnstead E-Pure 3-stage sieve and 0.2μm inline filter. TOC levels < 10ppb
‡ The cavity is constantly agitated and the water periodically evacuated from the cavity vol. while quenching

Parameter	Value
Water resistivity ($\Omega\ cm$)	$\geq 18.6 \times 10^{6\dagger}$
DI rinse time (min)	$\sim 5^{\ddagger}$
Drying solvent	IPA (semi-grade)*
Solvent exchange time (min)	~ 5
Temp ($^{\circ}C$)	22
Drying gas	N_2§

† Same UHP water specified in quenching step
‡ Rinsed under constant gravity fed flow
* Semi-grade acetone also used in Nb2-3[1] with no loss in performance
§ Oil-free, two-stage filter/drying, 0.2μm filter

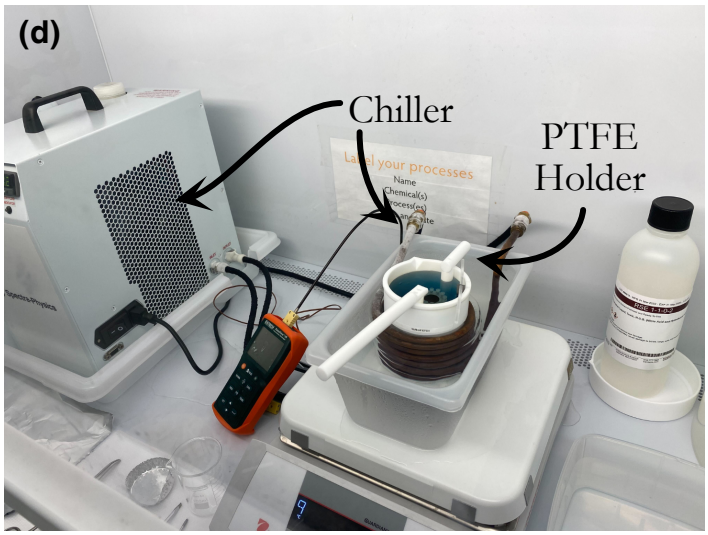

Figure 5.2: **Etching process steps**: A depiction of the primary etching steps and layout of components. (a) shows a cartoon of the etch setup in cutaway, depicting the copper cooling coil surrounding the primary Teflon etch containment. Below is a table of pertinent etching parameters used in this paper. (b) depicts the quenching process and information about the quench step in the table below. (c) depicts final DI rinsing and the solvent exchange process used for drying. Notes about the processing steps are provided below the tables for each step. (d) shows a picture of the complete etching setup during an etch of the cavity inside of the custom polypropylene hood and clean environment.

The top $100 - 200\mu$m of the surface is etched away to remove the damaged layer that is created from machining processes, and to remove any contaminants that may be on top or in

the surface. For niobium cavities, buffered chemical polishing (BCP), and electropolishing (EP) are typically employed [97]. For cavities used in particle acceleration, which require high electric field gradients, smoother surfaces are desired to reduce the risk of thermionic emission and multipacting at the surface by removing defects where the local field can exceed the work-function of Nb [134, 133]. While EP and BCP can both decrease surface loss and residual resistivity, EP has become more common for high-field applications [97]. In the low-field limit, however, BCP and EP have demonstrated similar performance and residual surface resistivity for high and medium-purity Nb stock [58, 45]. Electropolishing uses an applied voltage to electrochemically oxidize the niobium surface, before being dissolved into an electrolyte made of and HF/H_2SO_4, typically mixed in a 1:9 ratio. The EP process requires more infrastructure and has higher variability due to the larger number of parameters and sensitivity to acid aging, agitation, and temperature [8]. In contrast, BCP etching, which uses a chemical oxidizer in the form of nitric acid (HNO_3), is primarily dependent on temperature and relative concentrations to control the reaction kinetics. Because of our desire to develop an optimized process of easy implementation at the scale of a university research laboratory or clean-room, we will focus on BCP techniques that are specifically optimized for low-field (single-photon) conditions.

The addition of an H_3PO_4 or H_2SO_3 buffer has shown reduced etch rate and improved polishing action [100, 176, 6]. The limiting action of the buffer is believed to be due both to kinetics, but also the presence of a viscous product-rich layer at the Nb surface [14]. For lower relative concentrations of buffer, the polishing action is reduced, with selective etching occurring at the grain boundaries [176]. For the commonly used 1:1:2 mixture of $HF:HNO_3:H_3PO_4$ buffered chemical polished solution at an etching temperature of $\gtrsim 10°C$, $\sim 480-900$W of power is dissipated per $1m^2$ of Nb etched at a rate of $\sim 1\mu$m/min depending on the etch temperature [174, 97]. Another important factor in final etch quality is the amount of dissolved NbF_5 precipitate in the etch solution, with conventional wisdom being

< 20g/l of dissolved Nb^{5+} ions in solution, with our recipe designed to have < 10g/l for an $\sim 100\mu$m etch [174, 97].

The etch recipe used in this thesis diverges from the traditional BCP procedure, using a 1:1:2 HF:HNO_3:H_2O etch chemistry. The smaller surface area of our cavity, as listed in the Fig. 5.2 (a) table, means that our dissipated power is < 10W for even our highest etching rates. This, combined with the smaller volume of etchant required, means that an unbuffered etch can be used with minimal degradation in etch quality while keeping the reaction kinetics in a safe and manageable regime. In addition, the closed nature of the cavity geometry, with poor acid circulation, means that higher viscosity BCP mixtures can greatly effect the etch performance at the surface, leading to degradation in RF performance. This will be discussed in more detail in a following section.

The main etching steps of the cavity are depicted in Fig. 5.2. A custom Teflon PTFE holder was built to facilitate rapid and safe transfer of the cavity between etching steps (depicted in Fig. 5.2). PCTFE/PVDF screws were used to attach the cavity to the holder. The main etching was done in a 600mL PTFE beaker. The beaker was placed inside of a water bath, and surrounded by a copper coil attached to a 400W thermoelectric cooler which circulates a water/glycol mixture through the coil. The whole water bath was set on a magnetic stirring plate to circulate the primary acid volume. A pure niobium test strip is placed into the etch container alongside the cavity. The strip is periodically removed, rinsed, and its thickness measured using a micrometer to check etch-rate. Fig. 5.2 (a) provides a diagram of the aforementioned etch setup. To circulate the acid inside the cavity volume, the etchant is manually exchanged using a clean polypropylene pipette at $5 - 15$min intervals. Other pertinent etching parameters are listed in the Fig. 5.2 (a) table.

To prevent unwanted reaction products from adhering to the surface and leaving residue, the etchant is manually exchanged via a pipette for the final minute of etching, before being quickly evacuated of BCP solution and dunked into a (UHP/Type I) DI water bath and

vigorously agitated for $\sim$ 1min, quenching the etching reaction and reducing the chance of dissolved Nb salts and fluorinated compounds from depositing onto the surface. This step is done twice to further dilute any residual etching solution. Pertinent parameters of this step are presented diagramatically and in a table in Fig. 5.2 (b). Water purity is an important parameter in ensuring clean surfaces. It is standard practice in SRF accelerator cavity fabrication to use Type-I UHP high pressure rinsing to remove stubborn surface contamination.

Following this initial quenching step, the cavity is rinsed under a gravity fed stream of UHP Type-I water for $\sim$ 5min. Following this the cavity undergoes a two-step solvent exchange process. Here we use semi-grade isopropyl alcohol (IPA) to displace water, with each soak being $\sim$ 5min. Following this, the cavities are dried under high pressure dry nitrogen which has undergone two-stage filtration, with the final stage being a 0.2μm filter.

Following drying the cavity is sealed. An indium ring is sandwiched between the top of the cavity and the cap. A 3mm stainless steel tube is brazed into the copper cap. A Swagelok VCR© fitting is brazed at the other end. A second VCR fitting is attached to a copper pinch-off tube, allowing for easy replacement and reuse of the cap. A Corning-Gilbert GPO (SMP) hermetic connector (Part# 0119-783-1) is sealed into the top of the cap via a second indium o-ring, allowing for microwave feedthrough. An antenna is made on the vacuum-side of the feedthrough to couple RF power into the cavity. Because of the evanescent decay of the field energy, the coupling quality factor Q_{ext} is exponentially sensitive to this antenna distance from the pin. The cavity is left to pumpout for $\sim$ 24 hours before the copper pumping tube is sealed via a hydraulic pinch-off tool and placed into the fridge.

5.3 Nb cavity Measurements

5.3.1 Measurement setup

The cavity measurements were done in an Oxford Triton Dilution refrigerator with a base temperature at or below 50mK. The full measurement setup, including room-temperature hardware, is shown in Fig. 5.3

The cavity was measured in reflection (S_{11}). The input line uses discrete attenuation at 4K and the mixing chamber plate (MXC) of the dilution refrigerator, with an additional 20dB of attenuation being added by the insertion port of a Quantum Microwave cryogenic directional coupler, giving $\sim$ 60dB of total cryogenic attenuation. Reflected signal is passed through the directional coupler. A second directional coupler, which is used to couple the RF-pump power into a Josephson-parametric amplifier (JPA) is separated from the first coupler by a single-stage directional isolator. Finally a circulator couples the JPA into the output line, before entering a two-stage isolator. NbTi lines then pass the signal to 4K where they are amplified via a $4 - 8$GHz HEMT amplifier[4].

At room temperature, the cavity is measured via a Keysight PNA-X N5242A network analyzer. The time-domain setup is shown in Fig. 5.3. A sequence of upto four pulses are generated by the PNAX, with three combined into a single line via an quad XOR IC (DM74LS86) with $\leq$ 20ns switching delay. Two high-speed (20ns switching time) reflective SPDT RF switches [5] are used on the two available output sources, and combined using a splitter. One is used for the main drive tone of the cavity, while the second one is required for ringdown measurements, which will be discussed in a following section. Because the switches are operated slightly outside of their operational range, they add an additional $\sim$ 10dB of round-trip attenuation due to their increased insertion loss. An internal pulse is also used

4. Low Noise Factor model LNC4-8C

5. Mini-circuits model ZASW-2-50DRA+

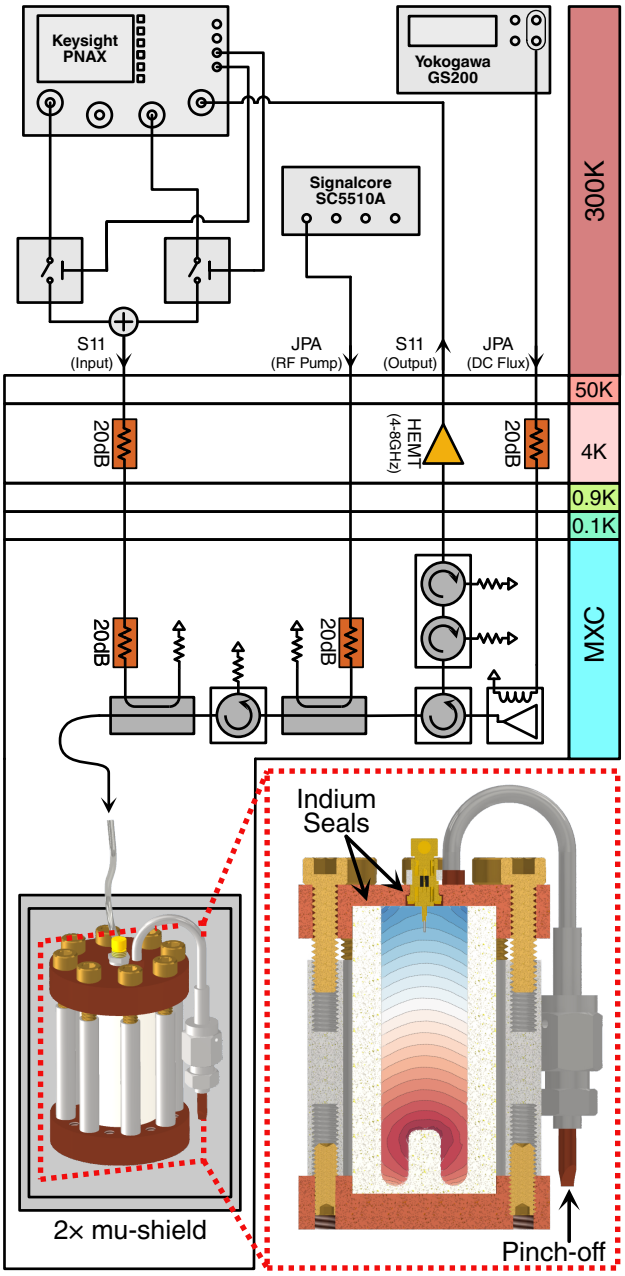

Figure 5.3: **Cryogenics and measurement Setup**: A diagram of the cryogenic and room temperature setup. An inset (framed in red) shows a detail of a cutaway of the cavity, highlighting the geometry and hermetic sealing components.

to turn on the PNA-X's ADC, shown on the rightmost port in Fig. 5.3.

The JPA RF pump-power is provided via a Signalcore SC5510A RF sythesizer, while DC biasing current, done via an external DC coil, is provided by a Yokogawa GS200 precision

current source.

5.3.2 SNR optimization and measurement considerations

In the next section we will look at the single-photon and low-power performance of our etched Nb cavities, but before we do that we must determine *what* constitutes a single-photon from the perspective of input power. In Ch.4 and later in Ch.6 we discuss ways of preparing single-photon Fock states using the qubit. These schemes have very little uncertainty in the prepared photon-number state that they create, however in our case, we only have access to a coherent drive of the system. This coherent state can be written in the Fock-basis as a superposition of Fock-states with associated amplitudes that follow a Poisson distribution [72, 163]:

$$|\alpha\rangle = e^{-\alpha^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle \tag{5.6}$$

To recover the Poisson-like probability distribution of finding n photons, we first note that the expectation value of particle number operator gives the mean photon-number, $\langle n\rangle = \langle\alpha| a^\dagger a |\alpha\rangle = |\alpha|^2$. Plugging this into Eq. 5.6:

$$P(n) = |\langle n|\alpha\rangle|^2 = e^n \frac{\bar{n}^{-\bar{n}}}{n!} \tag{5.7}$$

Where we will write $\langle n\rangle$ as $\bar{n}$ to express the mean photon-number. From this we see that, while we can't create a single-photon Fock-state with a coherent drive, we can create a coherent state that has a mean photon-number of $\bar{n} = 1$. If we want the steady state $\bar{n}$ based on a constant applied power, we can first write the input power P_{in} as [15]:

$$P_{in} = \hbar\omega \langle a_{in}^\dagger a_{in}\rangle \tag{5.8}$$

Where $\langle a_{in}^\dagger a_{in}\rangle = \langle\alpha| a_{in}^\dagger a_{in} |\alpha\rangle$, and where a_{in} and $a_{in}^\dagger$ are the associated creation and

annihilation operators of cavity photons under a constant input drive. We want to write the cavity number state in terms of these input operators. From Ref. [185] Ch.7 we can use the equation of motion from input-output formalism in steady-state conditions ($da/dt = 0$) to get an expression for a in terms of a_{in}:

$$a = \frac{2\sqrt{\kappa_{ext}}}{\kappa_{tot}} a_{in} \tag{5.9}$$

Where for generality we have written κ_{ext} to be the sum of external coupling rates (κ_1 and κ_2 in Eq. 2.15), and κ_{tot} to be the sum of *all* decay rates ($\kappa_i + \kappa_{ext}$). We can plug Eq. 5.9 into Eq. 5.8 to get:

$$P_{in} = \hbar\omega \frac{\kappa_{tot}^2}{4\kappa_{ext}} \langle a^\dagger a \rangle \tag{5.10}$$

Noting from above that $\langle a^\dagger a \rangle = \bar{n}$ we now have an expression of n_{bar} in terms of the input steady-state power:

$$\bar{n} = \frac{4\kappa_{ext} P_{in}}{\hbar\omega\kappa_{tot}^2} \tag{5.11}$$

With this in hand we now have a nice way of expressing the mean photon-number in terms of the coupling and internal loss rates (or internal and external quality factors), and the input power. Next it is insightful to look at what that means in terms of power *at* the cavity and what that means in terms of SNR and the time it will take to take a single measurement.

To do this we will use the Friis noise figure formula, which determines the signal-to-noise ratio for multistage amplification (or attenuation) based on the injected noise at each stage [172]. The noise factor of a component is given by the ratio of the input and output SNR:

$$F = \frac{SNR_i}{SNR_o} \tag{5.12}$$

For a single component, this values corresponds to the device's temperature. This is because electrons in a component at a finite temperature produce an output power density spectrum that follows Boltzmann-Maxwell statistics. This noise, known as Johnson-Nyquist noise, subsequently has a corresponding noise temperature T_n. It can be shown that the effective noise factor from Eq. 5.12 for a device is then:

$$F = 1 + \frac{T_n}{T_0} \tag{5.13}$$

Where T_0 is a reference temperature, which is typically set to room temperature, or more precisely $T_0 = 290K$. For an amplifier, or active components in general, T_n is based on numerous factors, and is usually an empirically derived value. For attenuators, from Johnson-Nyquist theory, the noise temperature is simply $T_n = (L-1)T$, where L is the attenuation ratio at the reference temperature, and T is simply the operating temperature.

For a single device, the total noise-figure is trivial, but what about a cascade of devices in series? If we take the SNR as S_i/N_i, where S_i and N_i are the signal and noise respectively at each stage of the chain, to derive the total noise-figure, we can then calculate the SNR at the output.

For the signal, the total output amplitude is simply the sum of amplification and attenuation. For noise, we must consider the fact that noise from the previous stage is either amplified (or attenuated) by a factor G_i, in addition to the noise that is injected at that stage. If we call the total output noise at each stage N_i, and the noise of the source N_s, we can show that the total noise figure is:

$$F_{tot} = 1 + \frac{N_1}{N_s G_1} + \frac{N_2}{N_s G_2} + ... \frac{N_n}{N_s G_1 G_2 ... G_n} \tag{5.14}$$

To write the above in terms of each component's noise factor, we can note that from the definition of F in Eq. 5.12 we can rewrite F as:

$$F = 1 + \frac{N_i}{N_s G} \tag{5.15}$$

Using the above relation we can then rewrite Eq. 5.14 as:

$$F_{tot} = F_1 + \frac{F_2 - 1}{G_1} + \frac{F_3 - 1}{G_1 G_2} + ... \frac{F_n - 1}{G_1 G_2 ... G_{n-1}} \tag{5.16}$$

To calculate the total signal-to-noise we need to take the ratio of total signal power and total noise power. The total signal power at the output of the chain is simply the power at the device and the total gain $P_{in} G_{tot}$. The Noise power is a bit trickier, but it can be shown from Johnson-Nyquist theory that it is simply:

$$P_n = F_{tot} k_B T_{dev} B G_{tot} \tag{5.17}$$

Where B is the bandwidth of the signal. In our case, the bandwidth is determined by the network analyzer's IF bandwidth. Here T_{dev} is the injected noise from the source, i.e. our cavity, and k_B is just Boltzmann's constant. In the real-world we also must consider the noise injected by the network analyzer itself, which can be added back into the above expression. Finally, because we typically express both noise factor in terms of a noise figure, in dB, and the total power in terms of dB, to calculate the SNR of the measurement we simply take the difference of $P_{sig} - P_n$.

Before we continue there are a few important things to note about the above expression. In particular, from Eq. 5.16 and 5.14, we see that the noise power is *dominated* by the first noise source in our system. This means that removing as much attenuation before the first amplifier, and reducing the first amplifier's noise temperature while increasing gain, is paramount for decent SNR. In our case, the HEMT amplifier has a noise figure corresponding to $T_n \sim 2$K at > 40dB of gain. By adding the potential 20dB and $T_n > 1$K of the JPA, we would see a further factor of $2 - 3$ improvement in the total SNR.

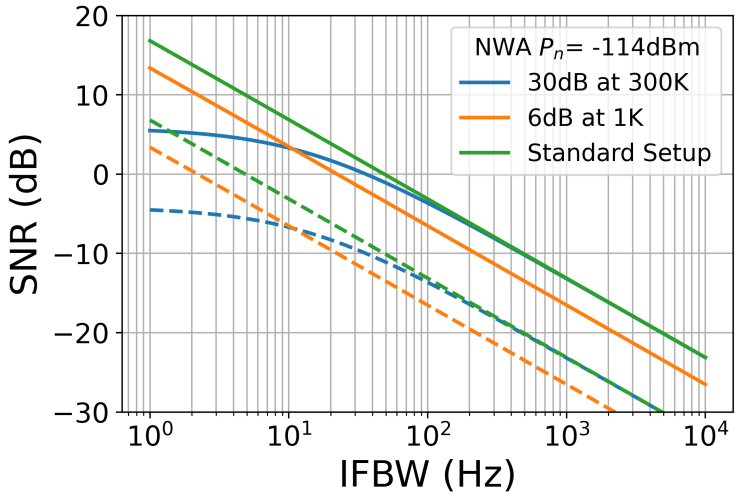

Figure 5.4: **SNR vs IFBW**: The SNR of various measurement chains, with added attenuation at high or low temperature leading to different types of degredation in signal integrity.

The above is presented in the form of a calculator in Appendix C. With it one can compute the estimated SNR versus IFBW for a given device. With the functions provided in the appendix, one can provide a list of components with their given gain or attenuation (given as negative gain) and their noise factor (or figure) or noise temperature. This can be expressed as follows:

```
generic_amplifier={'type':'amp',
                   'NF':1, #dB
                   'gain':1, #dB
                   'name':'generic'
                    }

generic_attenuator={'type':'atten',
                    'loss':1, #dB
                    'T_atten':1 #kelvin
                   }
```

With all of the individual components in the chain defined, one can then provide a dictionary with all the pertinent device components and details about the NWA's noise figure. Provided is an example of how this is entered for a hypothetical device:

```
test_dev={'meas_chain':[low_T_loss, HEMT,  high_T_loss, ZX60_LN,ZX60_LN, ZX60_LN,
↪  high_T_loss],
    'IFBW':np.logspace(.1,4,10000),
    'T_dev':.1, #temperature of device
    'PN_NWA':-117, #noise power of NWA
    'k_int':2.16, #internal kappa
    'kc_1':4.082, #kappa of first coupler
    'kc_2':0, #kappa of second coupler
    'f0':6.6271E9, #Hz
    'n_photon':10000,
    'meas_port':1, #which port you wish to measure from (1 or 2)
    'plot':True
    }

IFBW, sig, noise=SNR(test_dev)
```

The output of the SNR function is a plot of SNR (in dB) and the IFBW.

The above also estimates the amount of applied power *at the PNAX* from Eq. 5.11 needed for a given mean-photon number. For devices with $Q_{int} \geq 1 \times 10^9$, and similar Q_{ext} values, we anticipate PNAX output powers $\leq$ 90dBm, with an SNR of $<$ 10dB at an IFBW$\sim$ 1Hz from the above estimation using the measurement chain shown in Fig. 5.3. This is subsequently requires averaging $> 20\times$ for decent data, as is shown in the following section.

5.3.3 Single-photon quality factors

As we will see, and we have discussed in Ch.2 the internal quality factor can vary over power or temperature depending on the presence of two-level systems at the participating surfaces. Since the end-use of this cavity is for quantum information, where we encode the logical state on a single (or a couple) degree of freedom of the resonator, we need to directly interrogate the internal quality *at* the single-photon level.

To do this, cavity spectroscopy was done at a range of power from $\bar{n} \approx 1 - 2 \times 10^4$

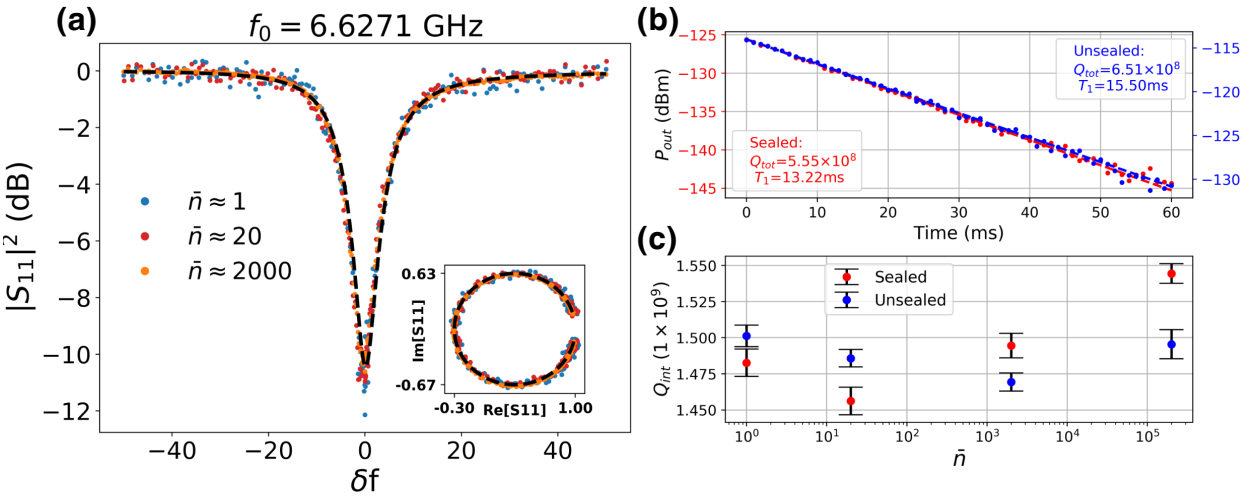

Figure 5.5: **Cavity spectroscopy and ringdown measurements**: (a) A detail of Nb2-1^3 S_{11} resonance and fit for various mean photon number, $\bar{n}$. (b) A comparison of cavity ringdown times for Nb2-1 sealed vs unsealed during cooldown, denoted by Nb2-1^3 and Nb2-1^4 respectively in Fig. 5.7 (b). The cavity was re-etched ∼120μm between subsequent cooldowns. (c) A comparison of Q_{int} vs $\bar{n}$ for both sealed and unsealed.

photons. Fig. 5.5 shows spectrocopic data and cavity ringdown for cavity Nb2-1 following the H_2O etch recipe outlined in Appendix 5.2.1. From $\bar{n} \approx 1 - 10^5$ there was no statistically significant increase in Q_{int}, suggesting two-level systems (TLS) were only weakly coupled to the field [74].

To determine if the cavity sealing and evacuation prior to cooldown played a significant role in preventing deleterious compounds from condensing on the surface and degrading the cavity quality factor, the cavity was etched an additional 150 μm using the same etch recipe and evacuated inside the dilution refrigerator within 90 minutes of etch quenching. Across the same range of mean photon-number there was no statistically significant change, with nearly identical performance as the evacuated cavity, indicating that, over the time scale of $1 - 2$ hours, and within the environment of the fridge, there was no degradation in single-photon quality factors. A second sample, Nb2-4, was also measured after initially being sealed, then re-etched and exposed to atmosphere for ∼8 days, showing degradation in single-photon quality factor, with maximum quality factor and residual resistant values

consistent with the cavities that had undergone the same etching process.

5.3.4 Temperature dependence

In the previous section we saw that at low to medium powers the quality factor was unaffected, suggesting that, at these powers, weakly coupled two-level systems (TLS) at the cavity surface play little or no role in cavity performance. Niobium is known to host TLS in its oxide layer, which is believed to be cause by the presence of NbO_x states within a disordered ∼5nm thick Nb_2O_5 surface oxide [157, 145]. The powers required to saturate the TLS to determine the conduction limited quality factor Q_0 saturated the amplification chain. So to study TLS the cavity was instead heated. The mixing chamber temperature was stabilized to ±5mK before spectroscopy was done.

The subsequent temperature dependent Q_{int} was fit to Eq. 2.35. The complete temperature dependent model that describes temperature saturation of weakly coupled TLS is given as[74, 110, 155].

$$\frac{1}{Q_{int}(T)} = \frac{1}{Q_0} + F_e \tan(\delta_{tls}) \tanh\left(\alpha \frac{\hbar\omega_0}{2k_bT}\right) \tag{5.18}$$

Where $F_e = t_{ox}S_e/\epsilon_r$, with S_e and t_{ox} being the surface to volume participation ratio of the electric field and surface dielectric thickness respectively. $Q_0 = \mu_0\omega_0/R_sS_m$, where S_m is the magnetic surface participation ratio, and R_s is the residual surface resistivity. From XPS measurements t_{ox} (see section 5.4), was determined to be $\sim 3.5 - 5.0$nm, depending on air exposure, while $\epsilon_r \sim 33$.

To study the difference of different etch buffer chemistry, cavity Nb2-1 was etched using a 1:1:2 H_3PO_4 BCP instead of H2O. Fig. 5.6 (a) shows that after two $\sim 100\mu$m etches only a mild improvement in Q_{int} was seen. Following the second etch, an aborted cooldown led to the cavity being held at ∼100K for ∼8 hours. Upon recooling, the cavity exhibited a distinct falloff in Q_{int} from $\sim 250mK$, suggesting a decrease in T_c, leading to a lower Q_0 in

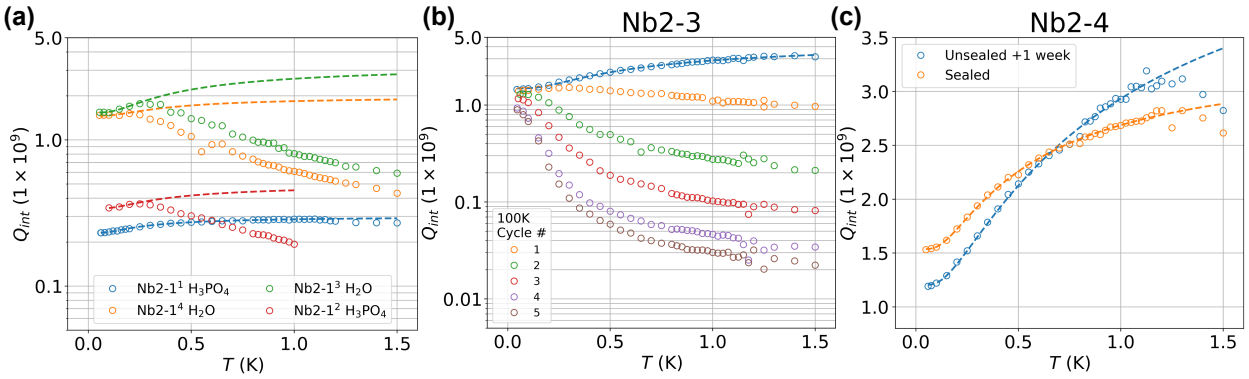

Figure 5.6: **Q_{int} vs Temperature**: (a) A plot comparing the effect of H_3PO_4 and H_2O etching recipes, and number of etches, on Nb2-1 cavity performance. (b) Nb2-3 Q_{int} versus temperature following repeated cycling of the cavity to 100K for ∼6 hours. (c) A comparison of Nb2-4 Q_{int} following rapid sealing and evacuation of the cavity volume, and exposure to atmosphere for ∼8 days.

Eq. 5.18. Re-etching Nb2-1 using the H_2O buffered etchant saw a significant improvement in the single-photon Q_{int}, with a nearly 5× increase, indicating an overall lower R_s. The T_c of the material, however, seemed unaffected, even after removing an additional $\sim 120\mu$m using the same etching recipe (Nb2-1^4).

The correlation between a decrease in T_c and being held at 100K seemed consistent with the formation of NbH_x from hydrogen interstitials at the surface between $50 - 150$K[103]. To study the evolution of NbH_x formation, Nb2-3 was subjected to repeated temperature cycling, being held at 100K for $\sim$ 6 hours, shown in Fig. 5.6 (b). After one 100 K cycle the cavity exhibited the same falloff in Q_{int} at $\sim$ 250mK as Nb2-1, with no statistically significant decrease in low-power ($\bar{n} < 1 \times 10^5$) Q. After 5 cycles the high-T Q_{int} had decreased by nearly 2 orders of magnitude, with the low-power Q_{int} only decreasing by $\sim$ 35%. All cavities, except Nb2-1^2, were initially cooled at a rate of $\sim 0.3 - 0.5$K/min between $150 - 75$K. This rate is slower than the recommended >1 K/min reported by Refs. [103, 9], however temperature dependent performance is consistent with performance of SRF Nb cavities cooled in dilution refrigerators.

The effects of oxide formation were also studied and is shown in Fig. 5.6 (c). With the

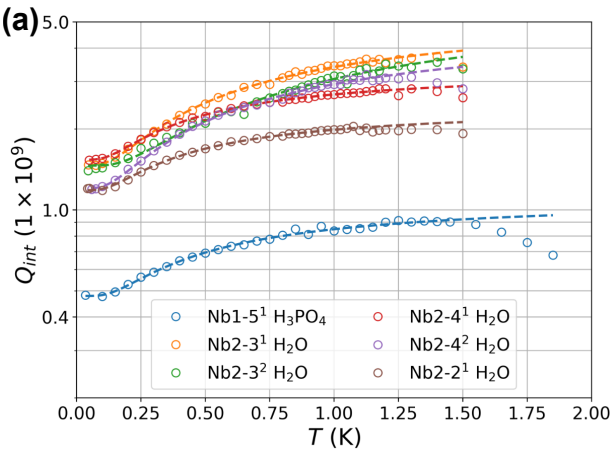

(b)

Device	f_0 (GHz)	Q_{int} (1×10^9)	$F_e \tan\delta_{tls}$ (1×10^{-10})	R_s $(1\times10^{-9}\Omega)$
Nb2-1^{3,†}	6.627	1.549	3.4	53
Nb2-1^{4,†,*}	6.682	1.474	1.8	63
Nb2-2^1	6.593	1.201	4.4	46
Nb2-3^1	6.618	1.469	5.1	25
Nb2-3^2	6.664	1.399	5.3	27
Nb2-4^1	6.578	1.532	3.6	34
Nb2-4^{2,*}	6.628	1.191	6.3	29
Nb1-5^1	6.692	0.481	12	105
Nb2-1^1	6.511	0.232	10	325
Nb2-1^{2,†}	6.575	0.343	9.4	258

$^\dagger NbH_x$ Q-degredation, *Unsealed during cooldown

Figure 5.7: **Cavity performance overview**: (a) Shows the temperature dependence of Nb1-5, Nb2-2, Nb2-3, and Nb2-4, showing typical performance following the described etching procedure. The superscript represents the etch number, while H_2O and H_3PO_4 represent the absence or presence of phosphoric acid buffer in the 1:1:2 BCP etching solution. (b) A table of all of the cavities measured, showing the loss tangent product $F_e \tan\delta_{tls}$ and residual resistance R_s. The yellow rows represent water (denoted as H_2O in (a)) buffer etchings, while aqua shows the performance of traditional H_3PO_4 BCP etching.

exception of Nb2-1^4, all cavities were hermetically sealed during cooldown, with $30-90$min between etching and evacuation. To study longer-term exposure, Nb2-4 was left open to atmosphere for 22 days following $\sim 150\mu$m etching. Nb2-4^1 was sealed and cooled following normal sealing procedures, and was exposed to air for ≤ 30min. Following extended exposure to air, Nb2-4 exhibited a $\sim 75\%$ higher loss tangent product ($F_e \tan\delta_{tls}$). The single-photon Q_{int} also decreased by $\sim$ 22%. Interestingly, the TLS saturated Q_{int} was also higher, suggesting a decrease in R_s.

With these results, and further etching, a more complete survey of the consistency of the processing was done. A survey of TLS characteristics is outlined in Fig. 5.7 (a) and (b). (b) presents a table with the finalized measured properties of all the cavities measured, showing consistent results over many individual cavities.

With these results in hand, and the results shown in Fig. 5.6 showing deviations between traditional BCP etch recipes and the one developed in this thesis, it is useful to ask where these deviations came from? In the following section we will take a deeper dive to connect

the surface chemistry to the dynamics that are shown above.

5.4 Material Analysis

In this section we will use a variety of techniques to explore how our etching procedure effects the surface chemistry of our cavities, and how these changes ultimately lead to better, or worst performance. We will first discuss the techniques used to carry out this investigation, before looking at a variety of data to then connect to the measured quality factors.

5.4.1 Material analysis techniques and processing

In later sections we will discuss data using two spectroscopic techniques, X-ray photoemission spectroscopy (XPS) and Time-of-flight secondary-ion mass spectroscopy (ToF-SIMS). Both of these techniques are capable of giving compositional information at various depths, however the two differ in both their analytic precision and the kind of information that can be extracted.

XPS determines the binding energy (E_b) of electrons via the photoelectric effect. The dynamics of this can be described by [178]:

$$E_b = E_p - (E_k + \phi) \tag{5.19}$$

A monochromatic X-ray source, with typical energies between $1000 - 1600$eV depending on the source target, is focused on the surface. Electrons with binding energies *less* than the X-ray energy are ejected. From the conservation of energy, the ejected electron's kinetic energy (E_k) is simply the difference between the binding energy and source energy. There are other systematics, such as the surface work-function of the sample and detector, ϕ, which can be calibrated for, and are typically small compared to the binding energy.

With E_k determined and the source energy known, E_b can be determined, which is

dependent both on the atomic species but also the exact hybridization and local bonding environment of compounds at the surface. While the X-ray penetration can be high, electrons that aren't in the top 10nm of the sample tend to scatter inside the material instead of being emitted [178]. This makes XPS great at studying the surface chemistry of a material. To study bulk composition, a milling process is introduced between measurements. This is done via argon cluster sputtering, which will be discussed in a later section.

While XPS is a versatile tool in looking at both composition *and* specific bonding between species, it is limited in its analytic capabilities, with a maximum sensitivity of $\sim 0.1\%$wt. In addition to this, XPS is unable to detect very light elements, in particular hydrogen. Because of these two factors, ToF-SIMS is used in conjunction with XPS to detect very low concentration, or light, contaminants. SIMS is a general class of surface analysis techniques that are all based on conducting mass spectroscopy on ionized particles emitted from surface. These secondary ions are produced by sputtering the surface with high-energy electrons, ions, neutral atoms or photons. The exact type of mass spectrometer used depends on the desired sensitivity. The simplest and most accurate is time-of-flight detection. Secondary ions are accelerated in an electric field, giving all particle the same kinetic energy. The velocity of the particle will depend on mass. The velocity is measured by determining the time it takes for the given ions to pass down a flight tube to the detector.

ToF-SIMS subsequently has comparable depth resolution as XPS, but has much higher sensitivity, with trace element detection on the order of a part-per million (0.0001%) or less [178]. Also, because ToF-SIMS is a mass spectroscopic technique, it is capable of measuring light atoms much more easily than XPS. The drawback of SIMS is that it is a destructive technique that is largely unable to determine bonding or stoichiometric details about the surface chemistry.

In Ch.2 we discussed a number of various loss-mechanisms that can degrade the quality factor, and in the previous section we showed the manifestation of that loss in our cavity

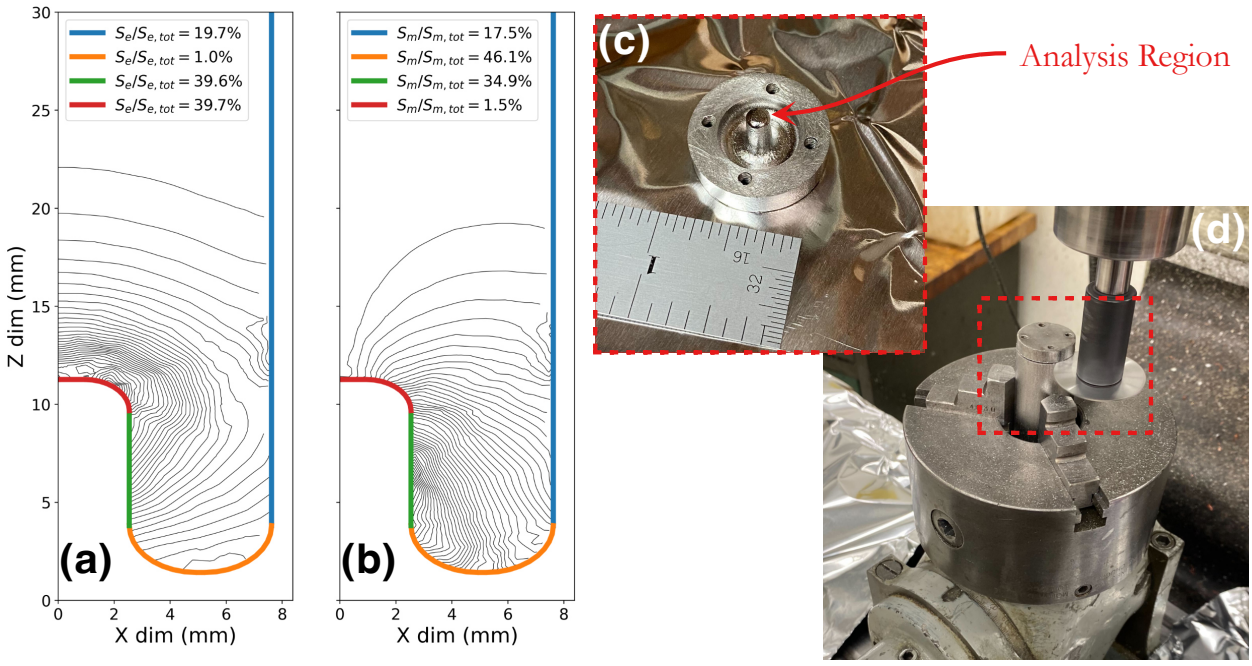

Figure 5.8: **Cavity participation and preparation**: (a) and (b) show the participation in the $\vec{E}$ and $\vec{B}$ fields of the cavity respectively, showing that the top of the pin has the highest participation in the $\vec{E}$ field, making it a good candidate for studying the presence of lossy TLS and dielectric compounds. (c) and (d) show a cavity that has been cut open to expose the participating surface, and the process of cutting it. Not shown is the use of IPA for cooling while cutting the cavity. A precision fine-tooth jewelers saw ($\sim$ 0.5mm thick) is used. All components used for cutting are thoroughly degreased and cleaned.

system. Dielectric and TLS loss mechanisms couple strongly to the electric field, which constitutes the first $5-10$nm of the surface (see Eq. 2.49), whereas magnetic defects, which primarily take the form of magnetic impurities and lattice defects that interact with surface currents, live in the London penetration length λ_L, or $\leq$ 50nm depending on the exact bulk quality of the Nb (see Eq. 2.54).

Fig. 5.8 shows the participation of the E (a) and B (b) field at various points on the surface. The top of the pin has the highest participation, and given its elevated location makes it an easy target for ToF-SIMS and XPS ion/X-ray sources/detectors to study. Fig. 5.8 (c) and (d) show the process in which the cavity is cut to get access to the interior surfaces. Given the sensitivity of both these analysis techniques, it is imperative that the process of

cutting the bottom of the cavity is done in a clean manner. The precision slitting saw and arbor is Alconox+TAMI cleaned in a sonicator to thoroughly degrease the surface. All other components, such as the chuck are also thoroughly cleaned with acetone and IPA. During cutting, semi-grade IPA is dripped onto the blade at low RPM to keep chips from filling the blade and cool off the part during cutting. As shown in (c), the whole process is done with the cavity mounted upside down, to ensure no chips contact the surface. Samples are carefully handled to reduce dust or cross contamination, with cut samples placed in ISO100 cleanroom bags filled with dry-nitrogen.

5.4.2 XPS experimental procedure

With the samples prepped, it was now possible to conduct XPS on the top of the pin. This analysis was performed on a Kratos Axis Nova spectrometer using monochromatic Al K source (1486.6 eV). Nb 3d, C1s, and O1s high-resolution spectra were collected using an analysis area of 0.3×0.7 mm^2 and a 20 eV pass energy with the step size of 100 meV. The survey spectra were collected using 100 eV pass energy. Charge neutralization was performed using a co-axial, low energy (≈ 0.1 eV) electron flood source to avoid shifts in the recovered binding energy. The C 1s peak of adventitious carbon was set at 284.8 eV to compensate for any remaining charge-induced shifts. Deconvolution of the high-resolution XPS spectra was performed in CasaXPS software using Lorentzian asymmetric curves for Nb metal phase, symmetric Gaussian-Lorentzian curves for all other elements' fitting, and a Shirley fitting for the background. The Nb 3d region consists of the two $3d_{5/2}$ and $3d_{3/2}$ spin-orbit split components. The peak area ratio of $3d_{5/2}$ to $3d_{3/2}$ was fixed to 3:2. The Nb 3d region was fit using six doublet components (Nb metal, Nb_2O_5, NbO_2, NbO, NbO_x, NbH_x) of $3d_{5/2}$ and $3d_{3/2}$ components for each sample. The energy split of one doublet component is 2.7 eV. The depth profiling XPS analysis was applied by 5 keV monatomic argon ion sputtering to investigate the oxide thickness and chemical states with depth.

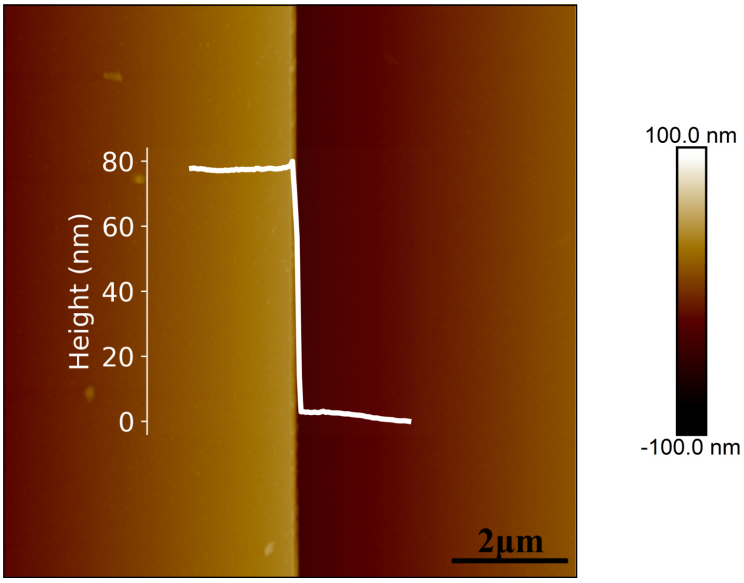

Figure 5.9: **AFM image of Nb thin-film step edge**: The above shows the edge profile of the test Nb thin-film. This film was patterned via flourine reactive ion etching, and differs from the un-patterned and etched sample analyzed in Fig. 5.11.

Depth profiling analysis via argon ion sputtering was performed to resolve the oxide thickness. To estimate the etch rate of the argon ion sputter, a thin film of known thickness was sputtered. The thickness of the Nb thin film was confirmed by Atomic Force Microscopy (AFM) on Bruker Dimension Icon AFM. The thickness of the Nb thin film is 73 nm, shown in Fig 5.9. A 5 KeV monoatomic argon ion sputter was utilized for all the depth profiling analysis. For the Nb thin film sample, the sputter was 3 min per cycle, and after 11 cycles the Al peaks from the sapphire substrate began to appear in the spectra. The etch rate for the 5 KeV monoatomic argon ion sputterer was determined to be $\sim$ 1.909 nm/min.

5.4.3 XPS oxide thickness analysis

In Sec. 5.3.4 we showed that the filling fraction from the loss-tangent product, F_e, is partially dependent on the oxide thickness. It has also been postulated that the TLS loss-tangent is dependent on the presence of sub-oxide states in the native Nb2O5 surface oxide.

A fine sputter etch (15 s per cycle) was was used on the Nb 2-4 cavity. Following an

Depth	Sample	Parameters	Nb	Nb_2O_5 $3d_{5/2}$	NbO_2 $3d_{5/2}$	NbO $3d_{5/2}$	NbO_x $3d_{5/2}$	NbH_x $3d_{5/2}$
Surface	Nb2-4 +22 days	Binding Energy (eV)	202.3	207.95	206.2	204.2	202.95	203.89
		FWHM (eV)	0.68	1.35	1.15	1.2	1.2	1.3
		Atomic Comp. (%)	37.41	51.52	3.65	0.57	1.1	5.75
	Nb2-4 +30 min	Binding Energy (eV)	202.31	207.96	206.1	204.25	202.9	203.45
		FWHM (eV)	0.56	1.39	1.50	1.20	1.20	1.51
		Atomic Comp. (%)	47.2	23.63	6.90	8.25	3.46	10.55
	Nb 73nm film	Binding Energy (eV)	202.26	207.86	206.3	204.2	202.9	—
		FWHM (eV)	0.60	1.31	1.80	1.70	0.40	—
		Atomic Comp. (%)	32.32	38.45	23.17	5.96	0.02	—
+35nm Depth Ar+ etch	Nb2-4 +22 days	Binding Energy (eV)	202.38	—	—	—	—	203.90
		FWHM	0.84	—	—	—	—	2.0
		Atomic Comp. (%)	76.82	—	—	—	—	23.18
	Nb2-4 +30 min	Binding Energy (eV)	202.38	—	—	—	—	203.83
		FWHM (eV)	0.82	—	—	—	—	2.0
		Atomic Comp. (%)	74.77	—	—	—	—	25.23
	Nb 73nm film	Binding Energy (eV)	202.3	—	—	—	—	—
		FWHM (eV)	0.48	—	—	—	—	—
		Atomic Comp. (%)	100	—	—	—	—	—

Figure 5.10: **XPS extracted parameter table for Nb2-4 oxide**

initial scan, Nb 2-4 was re-etched following the H2O buffered recipe outlined in Sec. 5.2.1, and placed in the XPS load-lock within 30 mins of etching. For the Nb 2-4 cavity that had been exposed to air for 22 days, the relative concentrations of Nb and O were consistent after 11 cycles. However, Nb 2-4 after etching/30 mins exposed to air only required 8 cycles to reach the Nb bulk.

Two analysis methods were used on the XPS spectra to investigate the Nb oxide thickness. First the XPS spectra are fit with six doublet components, revealing that there are four coexisting oxidation states and one hydride state besides the pure Nb metal phase (202.4 eV). These include NbOx (x<1), NbH_x (x<1), NbO, NbO_2, and Nb_2O_5, with respective shifts in binding energy compared to the Nb metallic phase peak[145, 168].

The first method is to use the ratio of the measured Nb 3d oxide to metal peak intensities (I_0/I_m) to calculate the oxide thickness at the surface of Nb cavities. The surface oxide thickness,d_{xps}(nm), was estimated using the equation[1, 83, 6]

$$d_{xps}(nm) = \lambda \sin\theta \ln\left(\frac{N_m \lambda_m I_o}{N_o \lambda_o I_m} + 1\right) \tag{5.20}$$

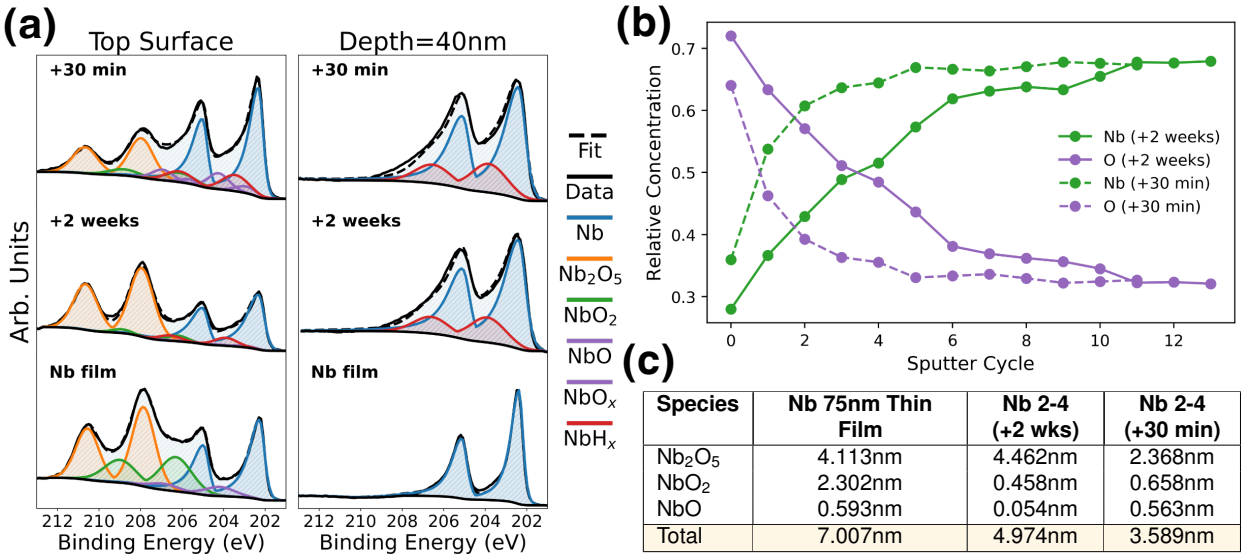

Species	Nb 75nm Thin Film	Nb 2-4 (+2 wks)	Nb 2-4 (+30 min)
Nb_2O_5	4.113nm	4.462nm	2.368nm
NbO_2	2.302nm	0.458nm	0.658nm
NbO	0.593nm	0.054nm	0.563nm
Total	7.007nm	4.974nm	3.589nm

Figure 5.11: **XPS and ToF-SIMS Surface Analysis** (a) shows a comparison between XPS spectroscopy of the Nb 3d peak for the top surface and the 35 nm depth of Nb2-4[2] after being exposed to atmosphere and after an additional H_2O etch followed by an $\sim$ 30min exposure. This is compared to a thin-film Nb that has been e-beam evaporated onto sapphire. (b) Shows a comparison in the relative Nb/O concentrations after various sputtering cycles for fresh etched (+30 min) and extended atmospheric exposure (+2 weeks). The sputter etch rate is 0.47 nm per cycle. One cycle is 15 s of 5 keV monoatomic Ar iron sputtering. (d) provides a table of calculated oxide thicknesses based on relative peak heights of various constituent oxides.

Where the ratio of the volume densities of Nb atoms in metal to different oxide phases (N_m/N_o) are listed in Fig. 5.10. The inelastic mean free path (IMFP) values, λ, for Nb_2O_5, NbO_2, and NbO can be found at Ref. [143]. These values are specified as being emitted normal to the surface (θ= 90°) of the Kratos Axis Nova XPS instrument, which has an angle of 45° between the X-ray source and the axis of the electron lens. The estimated oxide thickness of Nb thin film, and Nb cavities with different air exposure duration are shown in Fig. 5.11 (d).

The data indicates that the exposed cavity does see an increase in oxide thickness. The relative composition of the NbO_2, NbO, and NbO_x oxidation states, however, seemingly decrease in that time period, with lower relative concentrations, with the native oxide becoming dominant. From the increase in the loss tangent seen in Fig. 5.6, which seems to follow pro-

portionally with the increased oxide thickness. This seems to suggest that increase TLS loss due to cavity aging in air is almost solely due to continued Nb oxide growth, whereas the minimum loss-tangent is likely limited by the exact chemistry of the surface oxide. Recently, additional processing steps have been developed that can decrease the presence of deleterious oxidations states, such as NbO_2, NbO in electropolished cavities [156]

5.4.4 Loss due to bulk impurities

In the previous section we connected the makeup and thickness of the oxide layer with the increase in TLS loss after extended exposure to air. If we want to determine the limiting factor in either the single-photon quality factor, residual resistance R_s, or the degradation in T_c, we must look at the presence of impurities in the bulk.

In Fig. 5.6 (b) we see a decrease in T_c after a number of temperature cycles to 100K. This implies the evolution of NbH_x species in the London-penetration length. In Fig. 5.11, we also see the presence of NbHx, with relative concentrations of $\sim 23 - 25$at% present at 40nm. The precipitation of interstitial hydrogen into niobium hydrides has been studied both theoretically and observed experimentally [68, 19]. Empirical evidence shows that these hydrides lead to degradation in high-power quality factor in SRF cavities, and can be suppressed if cooled faster than $\sim$ 1K/min through the $50 - 150$K transition [103, 9]. Furthermore ToF-SIMS analysis has shown that for BCP processing, hydrogen interstitials may be present at depths $\geq 50\mu$m into the surface, due in part to the extremely high diffusivity of hydrogen into niobium [7]. The reason that we saw NbHx species in the freshly re-etched Nb 2-4 cavity, shown in Fig. 5.11, even though it hadn't been cooled, was likely due to the presence of existing NbHx species that had already precipitated in the bulk.

To confirm the presence of hydrogen in the bulk, at depths far greater than λ_L, we conducted ToF-SIMS on Nb 2-1[6], which had shown extensive Q-degredation and the char-

6. ToF-SIMS analysis done by Eurofins EAG Laboratories

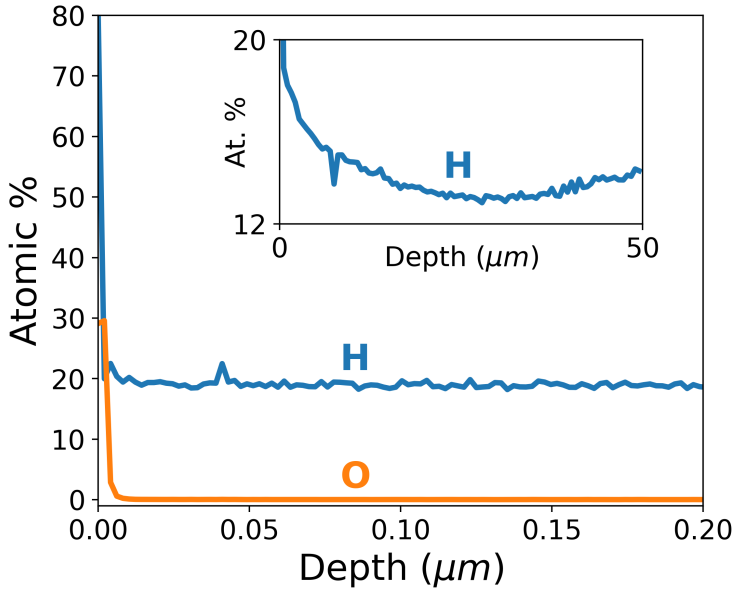

Figure 5.12: **Nb 2-1 ToF-SIMS Surface Analysis of Nb2-1**: Dynamic time-of-flight secondary ion mass spectroscopy for depths to 200nm for hydrogen and oxygen analysis, with the inset showing hydrogen down to 50. The relative atomic percent of hydrogen versus other species is much higher, indicating high diffusion from the etching process and the production of hydrogen interstitials. The presence of these interstitial, particularly in the London penetration length and deeper bulk, corroborate XPS data and the detection of NbHx species, which likely formed during the cooldown process. Oxygen concentrations drop precipitously after only a few nm, which is also consistent with oxide thickness analysis of Nb2-4.

acteristic decrease in T_c associated with hydride formation in Fig. 5.6 (a). Fig. 5.12 shows both depth analysis down to 1μm, showing oxygen and hydrogen relative concentrations, along with hydrogen concentrations at depth of $\leq 50\mu$m in the inset. The falloff in oxygen concentration at $\sim$ 5nm is consistent with XPS analysis described in Sec. 5.4.3, while H concentrations remain very high (and constant) throughout, with atomic percentages that are inline with the compositional data shown in Fig. 5.10 from XPS analysis. While there is a small falloff, we see that hydrogen concentration remains high down to depths of 50μm. While not groundbreaking, these results closely match the concentrations associated with H_3PO_4 based BCP etching, showing that the difference in chemistry neither improves (or makes worst) the relative concentration of interstitial hydrogen [7].

While NbH_x formation is associated with a characteristic drop in T_c, an overall drop in both single-photon quality factors and the maximum quality factor could suggest the presence of mechanisms that increase the overall residual resistance. This is most present in cavities etched using H_3PO_4 buffer, namely Nb1-5, and the first two etches of Nb2-1. As mentioned in Sec. 5.2.1, a pure Nb test-strip is placed in the etching bath to determine the etch rate. These test-strips are saved in case further analysis needs to be done. For studying the differences between the two buffered etch recipes, both strips from Nb2-1's first and fourth etch were analyzed using the same XPS protocol used in the cavity samples.

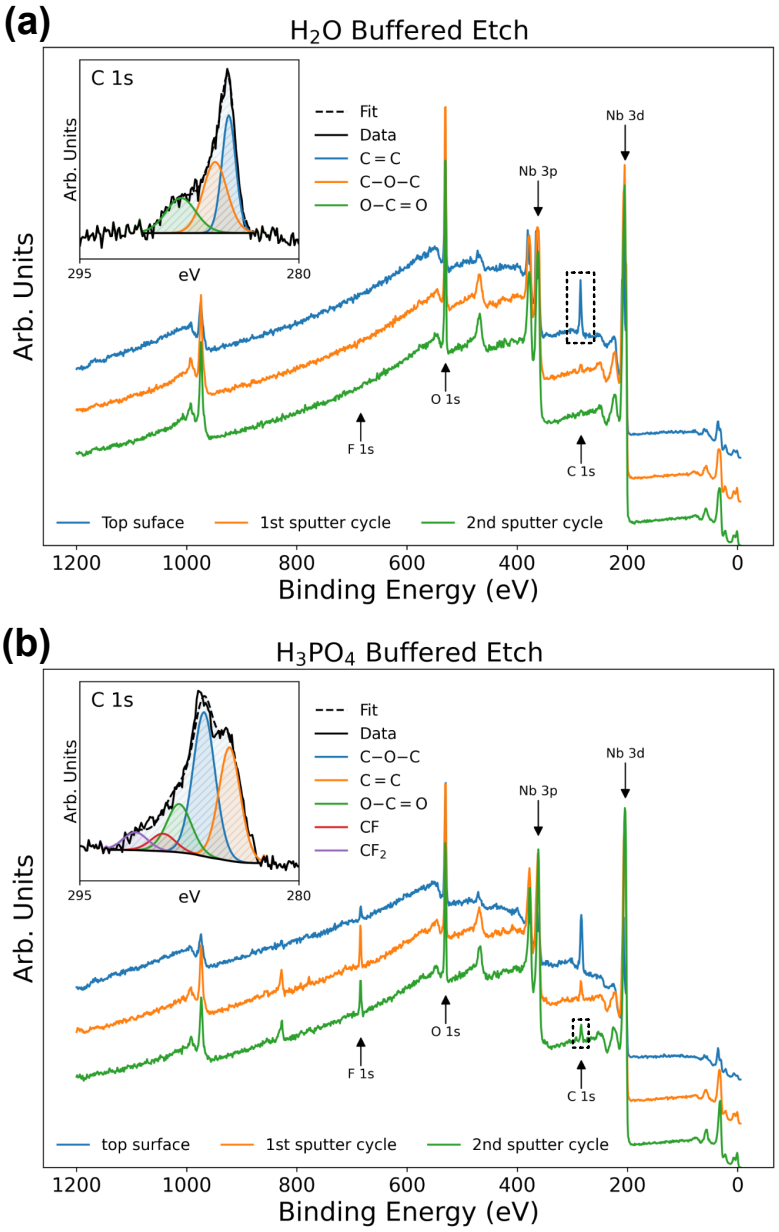

Figure 5.13: **XPS survey scan comparison of H_3PO_4 and H_2O etches** (a) shows survey scans for the test strip associated with Nb2-1[4]. The data shows that, for the H2O based buffered etch, there is no fluorine present at the surface or in the bulk, with the carbon peak disappearing after a single sputtering cycle, indicating surface-level carbon contamination. (b) shows the Nb test strip associated with Nb2-1[2], which was the second H_3PO_4 etch. The test strip exposed to phosphoric acid buffer shows a pronounced F 1s peak that initially *increases* with depth. In addition to this, the C 1s peak shows a characteristic broadening. Fitting this peak with C−F compounds give close agreement with the data.

Fig. 5.13 (a) shows survey scans for the test strip associated with Nb2-1[4]. Highlighted are the locations of the F 1s, O 1s and C 1s peaks for three sputtering cycles or $\sim$ 40nm. Also shown are the peaks associated with the metallic Nb phases. The inset shows a detail of the C 1s peak between $280 - 295eV$, fit to the C−O phases associated with surface level organic carbon contamination. The data shows that there is no fluorine present at the surface or in the bulk, with the carbon peak all but disappearing after a single sputtering cycle, suggesting carbon contamination is confined to the first few nanometers of the surface. (b) shows the Nb test strip associated with Nb2-1[2], which was the second H_3PO_4 etch shown in Fig 5.3.4 (a). In contrast to the water based etch, the test strip exposed to phosphoric acid buffer shows a pronounced F 1s peak (and F KLL line at $\sim$850eV) that initially *increases* with depth. In addition to this, the C 1s peak shows a characteristic broadening. Fitting this peak with C−F compounds give close agreement with the data, suggesting the presence of lossy fluorinated hydrocarbons on and in the surface.

These two survey scans showed that the water based etchant had no flourine at the surface or in the London-penetration length, with carbon confined only to the surface, as shown in the inset detail of the carbon (C 1s) peak for the top most surface. H_3PO_4 based etchant however showed flourine present in the top surface and the metal bulk. Fitting the broadened C 1s line with C−F species showed the presence of flourinated species in the bulk. The presence of these two features may explain the overall increase in R_s for Nb1-5, Nb2-1[1], and Nb2-1[2] without a notable change in the apparent T_c of the cavity. It should also be noted that Ref. [89], which used a nearly identical coaxial design and actively cooled H_3PO_4 BCP etching recorded near identical results in TLS and single-photon performance to Nb1-5.

A detailed analysis of electropolished Nb elliptical cavities, using an H_2SO_4 electrolyte, found the presence of similar CF_x compounds at the surface, in addition to hydrolyzed niobium flouride species [43]. A correlation was found between the agitation and removal of a viscous product layer at the Nb surface, and thorough post-etch rinsing, with the presence

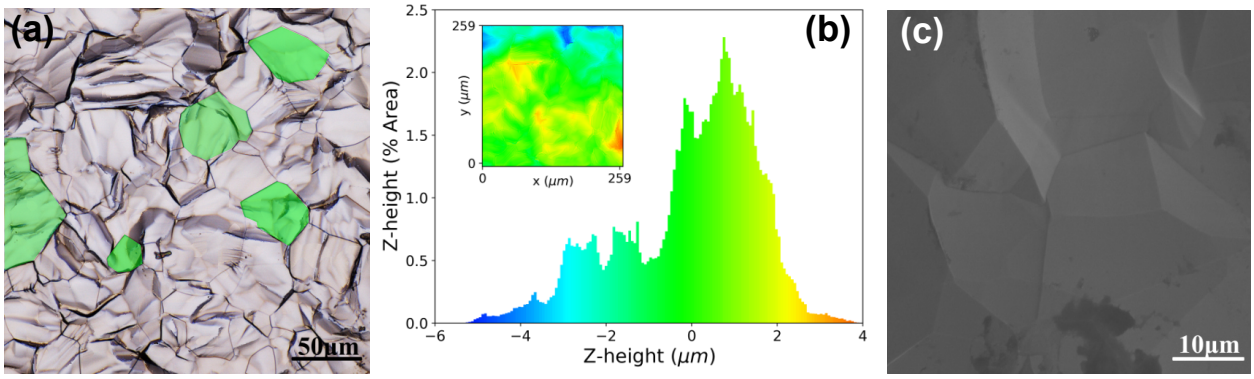

Figure 5.14: **Surface roughness and crystallinity of etched cavity**: (a) A high contrast image using a white-light confocal microscope of a post-etched surface, showing clearly delineated crystal domain edges. (b) surface roughness mapping of the patch shown in (a) showing height an inset height map and corresponding distribution of the surface roughness. (c) shows an SEM backscattered electron image showing a faceted surface structure.

of these features. The closed nature of the coaxial design makes constant agitation and steady-state removal of products from the participating surfaces difficult. By comparison, BCP etching is typically done under constant flow for elliptical cavities [174, 28]. During regular pipetting, the viscous layer can be seen as a dark blue for H_2O BCP) or green (for H_3PO_4 BCP) film that forms at the surface. The difference in color, and the presence of these deleterious species in XPS suggests that the presence of H_3PO_4 changes the reaction kinetics in a meaningful way, however further investigation into the exact mechanism would need to be undertaken.

5.4.5 Surface roughness and crystal domain size

In addition to chemical characterization, the roughness of the participating surfaces were also measure. The surface roughness was measured using an OLS5000 laser confocal microscope.Since the confocal design allows light that is reflecting off of a single 2D plane, by precisely sweeping the $z-$height of the sample one can create a 3D reconstruction of the surface. With post-processing one can accurately determine the height of features at the surface and do statistical analysis of the variation of these features. This is identical to the

technique used in Sec. 4.3.2 and shown in Fig. 4.5. In addition to this, SEM backscattering microscopy was also done to get information of surface morphology at the single grain level.

The results of this analysis is shown in Fig. 5.14, with the backscattered image of a single grain domain shown on the left and confocal microscopy of the surface with statistics of roughness on the right. From confocal micrsocopy the surface roughness was found to be $\sim 1-2\mu$m over the sites measured at the top of the cavity post. Confocal images show clear crystal domains with sizes ranging from $\sim 10-50\mu$m. While these grains may seem large, they would be considered polycrystalline or "fine-grain" when compared to "large-grain" SRF cavities [167]. While fine-grain cavities can have comparable performance to large-grain cavities, it has been shown that the fine-grain cavities are more susceptible to loss due in part to the precipitation of impurities to the grain boundaries, with performance being recovered through additional high temperature annealing and purification techniques [81, 57]. This could be one of the reasons why our H_3PO_4 buffered BCP etched cavities are more sensitive to the creation of $C-F$ compounds, which may coalesce in the grain boundaries, and may explain the elevated signal in the bulk, however further analysis of the exact source of $C-F$ in the cavity surface.

5.5 Discussion

Surface analysis using XPS and ToF-SIMS shows both the significant introduction of hydrogen to depths of $\geq 50\mu$m and the strong dependence of exposure time to atmosphere, increased oxide thickness, and a subsequent increase in the TLS loss of the resonator. Rapid sealing and evacuation ($<$ 90min) seems to keep the TLS loss product at values comparable to electropolished Nb elliptical cavities before any additional mid-T annealing to diffuse the oxide [156, 157]. This is notable since the geometric participation is nearly $\sim 3\times$ ($G = \mu_0\omega_0/S_m \sim 96\Omega$) greater in the coaxial design than the TM_{010} mode of an elliptical cavity, and a filling fraction F_e that is also $\sim 3\times$ greater than that reported for 5GHz

TESLA designs [156]. This means that, when scaled, the R_s are within a factor of 3 of the state-of-art Nb elliptical cavities of similar frequencies, and effective TLS loss-tangents that are between $2-5\times$ lower than normal EP treated cavities with no additional oxide annealing. XPS analysis in Fig. 5.11 (a) seems to hint at this discrepancy, with the H_2O BCP etched cavities showing a thin oxide that is largely composed of Nb_2O_5. The use of HNO_3, which is a strong oxidant, may preferentially produce pentoxide over other suboxides, producing a lower-loss dielectric surface.

Another conclusion is that, while reducing hydrogen concentration in the surface via high-temperature ($> 650°$C) vacuum annealing would likely decrease R_s, normal cooling of the cavities, and the presence of NbH_x species at the surface do not lead to significant decrease in single-photon Q_{int} for normal cooldown rates between $0.3 - 0.5$K/min without annealing.

When compared to equivalent aluminum coaxial cavities, the Nb cavities processed using the water based BCP etch recipe described in Sec 5.2.1 show a >15 times improvement in single-photon quality factor [149, 106]. This single-photon quality factor is largely limited by the presence of TLS loss mechanisms at the surface, with the intrinsic, conduction limited quality factors of all the cavities shown in Fig. 5.6 being in excess of 2×10^9, with the highest values being $> 3.5 \times 10^9$. If TLS was removed, this would correspond to single-photon lifetimes approaching 100ms for our $f_0 \sim 6.5$GHz cavities. Techniques for removing this TLS, in the form of mid-temperature ($\sim 450°C$) oxide bake, have shown a marked reduction in the TLS loss product [156].

In addition to TLS reduction, there have also been the development of techniques which infuse nitrogen into the surface, which has shown the ability to reduce the residual resistivity by a factor of three, leading to potential single-photon lifetimes in excess of 250ms [82, 56]. This, along with the oxide diffusion to reduce TLS, and hydrogen removal from the surface, all require high-temperature vacuum annealing at pressures below 1×10^{-6}mbar. Work has

been done to develop a small scale furnace capable of these processing steps, and will be discussed later in the thesis, however this work will likely be the task of future graduate students within the lab.

In summary, the work outlined in this chapter represents an important advancement in making state-of-the-art quantum memories that can be manufactured at the scale of a university research lab, with the processes all taking place inside of the lab itself [7]. While this process is capable of producing high-performance cavities with high levels of consistency, it is only really worth the effort *if* a qubit can be incorporated into the system. By choosing the coaxial cavity—which has been a staple in the cQED community for nearly a decade—we have avoided the potential headaches associated with incorporating qubits into more traditional SRF cavity geometries. With that said, because of the extremely high coherence, and thus high sensitivity to loss, there are a few additional considerations that must be made before incorporating a qubit into the mix. In the next chapter we will discuss those challenges and incorporate those ideas into a design that can preserve the performance we have worked so hard to achieve.

7. All of the data shown did not take place in the Pritzker Nanofabrication Facility given the lack of infrastructure for etching large scale devices

CHAPTER 6

ULTRA-HIGH Q CAVITY-QUBIT SYSTEMS

6.1 Introduction

In the previous chapter we have shown how to build a high-Q bulk niobium cavity, but as we saw in Ch. 4, to turn a cavity into a quantum memory requires the introduction of a non-linear element. This non-linearity comes in the form of the transmon qubit, which was discussed in detail in Ch. 1 and Ch. 3. Before we can incorporate a qubit into our pristine high-Q cavity, however, we need to first consider the *what* the qubit does to the coherence of our cavity mode.

In the previous chapter we went to great lengths to remove as many sources of loss as possible, so it may seem a bit counter-productive to now place a two-level system that has a coherence that is 3-orders of magnitude *lower* than the mode. This is, however, a necessary evil, since the inherited anharmonicity of the qubit allows us to conduct quantum operations on the system. To make our jobs easier, and build upon the knowledge gained in the previous chapter, we will use the tested coaxial cavity for our system, and begin by studying single-mode dynamics to constrain the complexity of our system, and the parameter space in which we must optimize.

In this chapter we will discuss how we can mitigate some of the introduced loss to protect our mode. Then we will look at the design parameters of this qubit-cavity system.

6.2 Modelling qubit induced losses

As previously discussed, the qubit is an anharmonic oscillator, in the form of an LC-circuit in which the inductor, thanks to the Josephson-effect, gives the system non-linearity. We once again use the transmon qubit for this purpose, and use the techniques outlined in Ch.3. In sections. 3.4.2, 3.5.1, and 3.5.3, we used simulation tools to understand how the

qubit participates with the cavity mode, and as a result, determine the pertinent dispersive Hamiltonian parameters that fully describe the dynamics. In this section we will be using these same tools and techniques in order to instead determine, based on the interaction strength, how much the qubit will decohere our cavity. First we will discuss the ways *mechanical* coupling of the qubit chip to the mode can potentially lead to dephasing. Next we will study how qubit location can be optimized to minimize the reverse-Purcell effect, where the qubit decoheres the cavity, before discussing how the participation of the electric field in the dielectric of the chip can limit mode coherence. Finally we will investigate how the geometry of the qubit mount can effect mode coherence. These insights will then inform our design considerations in the following sections.

6.2.1 Qubit vibration and induced dephasing

The frequency of our cavity is dependent on both the relative permittivity and permeability of the material in our mode volume. The presence of a dielectric in the mode volume will lead to a *decrease* in the the mode frequency. We have already seen this *explicitly* in Eq. 2.41, which describes the frequency of the TE_{nml} mode, showing that $f_{nml} \propto 1/\sqrt{\mu\varepsilon}$, where ε and μ are the relative permittivity and permeability respectively for the dielectric medium. In most cases, we can take these values to be ~ 1, since the mode volume is just vacuum. With the addition of the qubit, however, we have *also* introduced a chip made of dielectric material, which will inevitably lead to a shift in the cavity frequency.

In the static case, this shift is easily accounted for. In Ref.[149] the frequency shift was used as a way to determine the insertion of the qubit into the cavity volume, with shifts of $\sim$ 100MHz for insertion lengths on the order of a few hundred μm when inserted into the coaxial cavity from the side. The next inevitable question is what would happen if the chip—which is, in effect, a cantilever fixed at one end—were to vibrate due to a constant drive? How *much* of a shift can we expect for vibrations on the order of nanometers? A

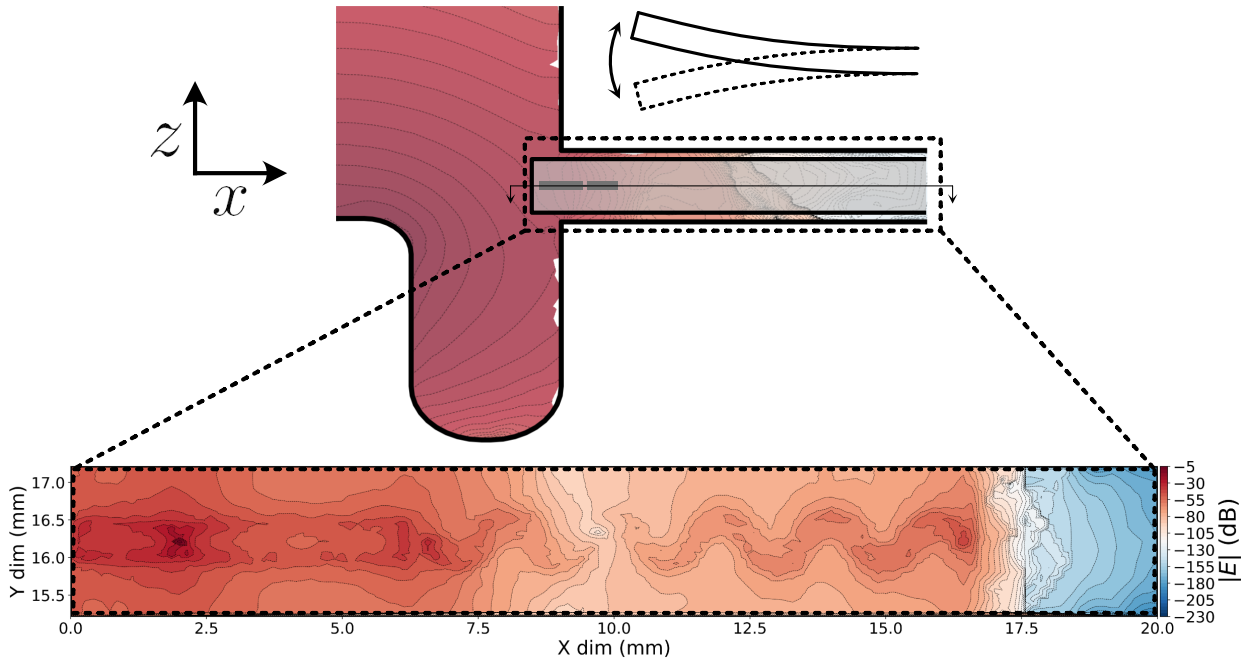

Figure 6.1: **Chip orientation and participation in cavity $\vec{E}$-field**: A diagram showing the orientation of the chip in the cavity field. Above is a cross-section the chip undergoing deflection due to cantilever vibration. When compared to the inset below, the chip with highest deflection also has the highes $|E|$-field participation,

time-dependent detuning of the cavity frequency would lead to an effective dephasing of the state. If the frequency of these vibrations are on the order of $1/T_{1,c}$, where $T_{1,c}$ is the cavity frequency, then the effective intrinsic linewidth of the cavity will be broadened. Because of this, it is important to quantify the scale of shift for a given vibrational amplitude, and the vibrational frequency one can expect for a rigid cantilever.

To determine the expected frequency shift for a given displacement, we can begin with Maxwell's curl equations which gives the identities:

$$\nabla \times E = j\omega\varepsilon\vec{E} \tag{6.1a}$$

$$\nabla \times H = -j\omega\mu\vec{H} \tag{6.1b}$$

A perturbation in the field caused by a volumetric change in a media with permittivity and permeability ε and μ will lead to a change in frequency from $\omega_0 \rightarrow \omega$. We can solve for $\Delta\omega = \omega - \omega_0$ from Eqs.6.1 (a) and (b), with the approximate form becoming[1]:

$$\frac{\Delta\omega}{\omega_0} \simeq \frac{\int_{\Delta V}(\mu|\vec{H}|^2 - \varepsilon|\vec{E}|^2)dv}{\int_{V_0}(\mu|\vec{H}|^2 + \varepsilon|\vec{E}|^2)dv} \tag{6.2}$$

We can see that the integrand in both the numerator and denominator are simply the magnetic and electric stored energy in the perturbing material, and can be rewritten using Eqs.2.47 (a) and (b) to give a simplified expression:

$$\frac{\Delta\omega}{\omega_0} = \frac{\Delta W_m - \Delta W_e}{W_m + W_e} \tag{6.3}$$

The above indicates that the fractional change in stored energy inside the chip is directly related to the expected perturbation in frequency. The above also assumes that the field profile due to the perturbation remains largely unchanged. In our approximation for the chip this is a safe assumption given the small amount of the mode volume we expect the chip to occupy.

Before we compute the above, we can note that, for sapphire, which make up our qubit chip, we expect the permeability to be equal to that of vacuum, so any change in the position of the chip should not change the stored energy of the $\vec{H}-$field. This simplifies Eq. 6.3 to simply $\Delta W_e/W_e$. We can get an idea of how the stored electric energy is distributed in chip. In Fig. 6.1 we see that the $\vec{E}-$field participation is quite high, especially at the end of the chip where we expect the largest changes in location due to vibration. As the chip displaces in a given direction, we have a small volumetric change associated with it, ΔV. By computing the integral in *only* this volume, we can calculate the frequency difference.

1. For a more detailed derivation see Sec. 6.7 in Ref. [144]

To do this calculation, we can employ techniques discussed in Ch. 3, namely in Sec. 3.3.2. Before we continue, you may be wondering, if we plan on using simulation techniques such as those discussed in Ch. 3, why not just displace the chip and use eigenmodal simulations to extract the new frequency? The main issue with eigenmodal simulations is the general lack of numerical precision; for the shifts we are interested, which are $\sim$ 1part per billion, we would need to set our convergence criteria to be *atleast* on the order of that level of precision. In practice, the numerical floating point precision of the eigenmode simulation means that successive simulations would have eigenmode frequencies that deviate by orders of magnitude more than our required precision. Instead, we want to create a *single* eigenmode simulation, with meshing that has a characteristic length *smaller* than our chip perturbations, and *then* compute the above integrals for various displaced volumes. Ideally we could reuse the mesh and simply move the chip location by our desired displacement to then compute the stored energy, however doing so nullifies the previous mesh setup in HFSS. Instead, we will create a number of *non-model* objects that represent the displaced chip volume, and simply integrate over the object with sapphire's dielectric constant.

The *way* the chip can be expected to vibrate is determined by geometry. With a length of $\sim$ 25mm, width of 2mm and a thickness of $< 500\mu$m, the fundamental eigenmode of the chip will be due to a displacement in the $\hat{y}-$direction as shown in Fig. 6.1. Fig. 6.2 shows the HFSS setup described above. The setup in pyHFSS for the depicted setup is:

We additionally set our maximum mesh lengths over the volume of the perturbed chip locations to ensure that the small perturbation to stored energy can be properly determined. The assignment of a non-model object to each box in line 9 ensures that HFSS does not attempt to treat the box boundaries as an electrodynamically interacting surface. This once again shows the approximate nature of this simulation, in that the field profile should also be perturbed. This however is already accounted for in Eq. 6.2.

Once the above simulation has run we can use the techniques discussed in Sec. 3.3.2,

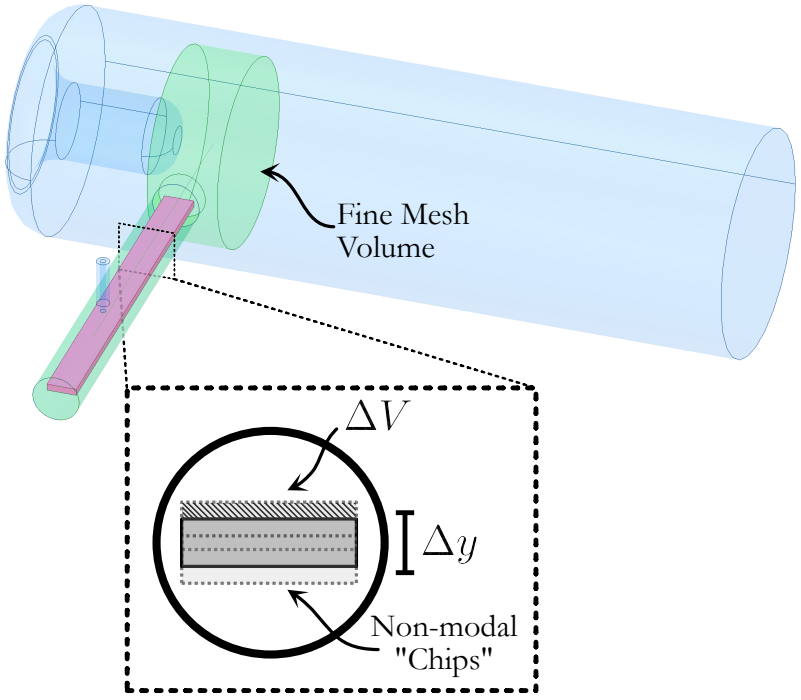

Figure 6.2: **Perturbative chip simulation setup in HFSS**: A depiction of the simulation setup, with the sapphire chip highlighted in pink, the mode volume simulated in blue, and the mesh volume in green. A diagramatic cutaway of the chip showing the displacement and the location of overlapping "non-modal" mesh boxes used to determine changes in the stored electric energy. It is assumed that small displacements will only weakly change the overall field profile of the resonator.

where we instead only compute the stored electric energy over the perturbed volume and original chip volume. The calculation to do only this computation is:

Where we have written a compact function that takes the volume and relative dielectric constant to compute W_e. First we need to solve for the numerator, which represents the *total* stored electric energy over the entire mode volume, and can be split into the stored energy of the chip and vacuum. Using the above `vol_energy_calc` we can do this computation as follows:

Where `E_energy_total` on line 3 is W_e for the entire system. From there we can calculate the perturbed stored energy `stored_E_perturbed` in the same manner for each displaced chip box. To get just the energy in ΔV, however we need to subtract the unperturbed chip stored

```
#create new chip box at new chip position that represents the perturbed location for later integration
new_z_locs=[]
new_chips=[]
perturb_vars=np.linspace(-chip_perturb, chip_perturb, variations+1)
for I, vars in enumerate(perturb_vars):
    new_z_locs.append(EM_design.set_variable('Chip_var_%i_loc_z'%I, value='%.3fum'%vars))

    new_chip_loc_box=model.draw_box_corner([chip_x_origin, chip_y_origin, chip_z_origin+new_z_locs[-1]], size=[chip_x_dim, chip_y_dim, -chip_thickness], name='new_chip_loc_box_'+str(I))
    model.assign_non_model(new_chip_loc_box)
    new_chips.append(new_chip_loc_box)

#do some meshing to better sample the E-field in the chip volume(s)
model.mesh_length(name_mesh='substrate', objects=[chip_mesh_box], max_length='%.3fum'%100)
model.mesh_length(name_mesh='vac_mesh', objects=[vac_mesh_vol], max_length='%.3fum'%200)
```

```
def vol_energy_calc(self, volume, variation, eps_r):
    calcobject=HFSS.CalcObject([], self.setup)
    smooth=False
    vecE = calcobject.getQty("E")
    if smooth:
        vecE = vecE.smooth()
    A = vecE
    B = vecE.conj()
    A = A.dot(B)
    A = A.real()
    A = A.integrate_vol(name=volume)

    lv = self._get_lv(variation)
    return eps_r*const.epsilon_0*A.evaluate(lv=lv)
```

energy. The final result for the perturbed frequency of the system is thus given by:

One design aspect we can vary in the design of our qubit-cavity system is the orientation

```
chip_E_energy_0=energy_calc(eprh, volume=str(chip), variation='0', eps_r=11.2)
vac_E_energy=energy_calc(eprh, volume=str(vac), variation='0', eps_r=1)
E_energy_total=chip_E_energy_0+vac_E_energy
```

```
solns=EM_setup.get_solutions()
freq_0=solns.eigenmodes()[0][0]
delta_omega=freq_0*1E9*(np.array(E_perturb)-chip_E_energy_0)/(E_energy_total)
```

of the qubit chip. We can mount the chip horizontally, i.e. perpendicular to the axis of our coaxial design, or vertically, where the chip is perpendicular to the coaxial axis. We will see in a subsequent section that we can achieve the necessary coupling to the mode regardless of this orientation, so either are valid solutions.

Fig. 6.3 (a) and (b) shows the sensitivity of the cavity mode to chip vibration. One interesting takeaway is how the chip orientation and the natural symmetry of the coaxial cavity's fundamental mode effect the amount of dephasing that *can* occur for a given displacement. Because the field gradient decays exponentially along the $\hat{z}-$axis, small displacements in the chip lead to dramatic changes in the amount of stored energy in the chip volume, leading to large fluctuations in frequency that follow an exponential relationship. Over the small displacements we are interested in, $\Delta\omega$ looks, locally, linear. Inversely, the vertically mounted chip, with its fundamental mode leading to displacements in the $\hat{y}-$direction, the field, which extends radially from the pin, leads to only a quadratic relation between frequency detuning and displacement. This means that, by centering a chip along the $\hat{y}-$axis, one could expect a quadratic suppression of dephasing amplitude.

Before we continue, we should make a few important notes about the above approximation. For one, this is a *worst-case* scenario; this simulation is treating these cantilever vibrations as a displacement in the entire chip. In reality, the displacement of a cantilever beam, $\Delta d \propto x^3$, where x is some distance along the length of the cantilever such that

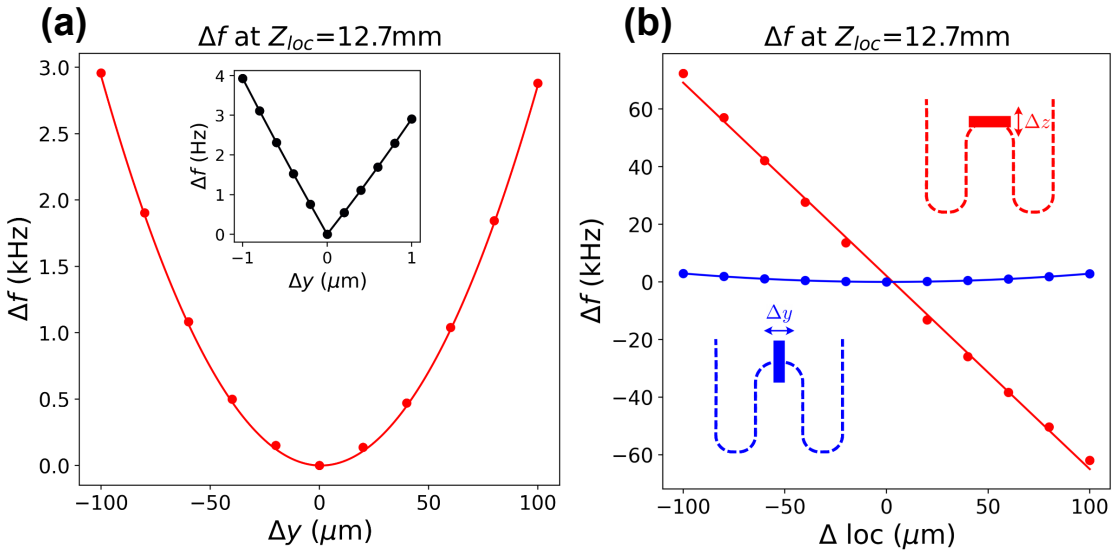

Figure 6.3: **Chip displacement result**: (a) shows the change in frequency for a given displacement for a chip in the vertical orientation, showing a quadratic suppression in vibrational sensitivity. The inset shows a zoom in about zero, with the discrimination of the simulation mesh leading to the step-like sensitivity and asymmetry of frequency to position. (b) compares the vertical line in (a) to a chip at the z location mounted horizontally. The line, which appears linear, but is actually weakly exponential over this range, shows a much higher frequency dependence on location, due in part to the location of the chip in an evanescent region of the cavity near the top of the pin.

$0 < x < L$ where L is the cantilever length. We can see if Fig. 6.1 (b), however, that most of the field's stored energy is in the end of the chip where the displacement is the highest.

Two additional points of consideration is the displacement and natural frequency of our free-vibrating cantilever chip. To calculate the displacement from the equations of motion, we would need to know the driving amplitude, which would require a detailed understanding of the way vibrational energy is coupled into the chip. This is difficult to fully quantify, however we *can* deduce the *maximum* displacement of the chip, which is simply the point in which the chip breaks. The point of breaking occurs when the stress on the chip is greater than or equal to the maximum tensile strength of the sapphire. The stress on the chip is given by:

$$\sigma = \frac{My_n}{I} \tag{6.4}$$

Where σ is the stress on the beam at a given location, M is the moment at the fixed point of the cantilever, and I is the moment of inertia for the beam [76]. y_n is the distance from the neutral axis of the beam, which can be taken as $y_n = t/2$ where t is the thickness. The moment at the fixed end of the beam is simply given as $F_a x$, where F_a is the force applied and x is the distance from the fixed point. We can also write the static force applied to the end of the beam ($x = L$) as:

$$F_A = \frac{Ewt^3}{4L^3}d \tag{6.5}$$

Where E is the modulus of elasticity (Young's modulus) and d is the displacement of the beam at the free end. The moment of inertia for a rectangular member is simply $I = wt^3/12$ where w and t are width and thickness of the chip respectively. Plugging equation 6.5 into $F_a x$ and substituting it into M and Eq. 6.4 we can get an expression that relates the stress to the displacement:

$$d = \frac{2\sigma L^2}{3Et} \tag{6.6}$$

The maximum stress, σ_{max} is equal to the tensile strength of the material. For sapphire, the tensile strength ranges in literature from $190-400$MPa, whereas the modulus of elasticity is given to be $\sim 375-400$GPa [108, 162]. For the chip depicted in Figs. 6.1 and 6.2, $L = 16.5$, while $t = 460\mu$m. This gives estimates of maximum deflection at the end of the chip being between $\sim 170 - 400\mu$m. Given the high modulus of elasticity, however, for our 2mm wide chip, a force of > 1N would need to be applied to the end of the chip. Given the small mass of the chip, this means that the driving amplitude would need to be very high to even approach this level of displacement. The inset in Fig. 6.3 (a) however shows what one may expect from small displacements. The shape deviates from the parabola that we see over larger displacements due to the finite precision of our simulation and minimum size of our mesh,

however we can see that to get a $\sim$ 1Hz broadening of our resonance, we would still need displacements of $\sim$ 100nm in the vertical orientation, whereas in the horizontal orientation we estimate an $>$ 50Hz shift for the same displacement.

Finally, we would like to estimate the natural resonant frequency of our cantilever. We can calculate the natural frequency by treating the cantilever as a 1D beam of uniformly distributed mass. By then splitting the beam into elements and solving the equation of motion for each point, taking into account the boundary conditions, we can determine the natural frequencies f_i such that $i = n - 1$ where n is the number of nodes. We can also use an approximate form for f_i which is given by [192]:

$$f_i \approx \frac{K_i}{2\pi}\sqrt{\frac{EIg}{m_L L^4}} \tag{6.7}$$

Where m_L is the load per unit length, which for an unloaded beam, is simply $m_L = \rho wt$, where ρ is the specific density, which is 3.98g/cm^3 for sapphire. The finite element method (see Appendix B) and approximate form in Eq. 6.7 give a natural resonance of between $\sim 90 - 100$kHz for the chip in question. At these frequencies the period of oscillation is shorter than the ringdown time of the cavities outlined in Ch. 5, meaning that we would expect this dephasing to lead to homogeneous broadening of the cavity resonance. Although we have seen that the chip *can* be displaced enough to broaden the cavity linewidth, given the chip stiffness, and by sufficiently fixing the chip, we can hope to ensure that all chip vibrations remain common mode with the cavity.

6.2.2 χ and the reverse-Purcell effect

In Sec. 1.1.2 we discussed the process in which a lossy cavity can increase the decoherence of a qubit (or atom). We also discussed how, inversely, a qubit with a decay rate $\gamma \ll \kappa$, where κ is the intrinsic linewidth of our cavity, can increase the effective linewidth of the cavity. These enhanced decay rates, known as the Purcell and reverse-Purcell effects, are described

by Eqs. 1.7 and 6.2.2 respectively. For our purpose, we are interested in the qubit's effect on the storage mode of our high-Q cavity. In Eq. 6.2.2, we relate the effective rate in terms of the coupling rate, g. Since we are in the dispersive regime, a more natural way of expressing this effective enhanced decay rate is in terms of the dispersive coupling χ. For a transmon qubit, we have computed χ in terms of g in Eq. 1.24. Solving in terms of g and plugging it into Eq. 6.2.2 we get an effective Purcell rate of:

$$\kappa_p = \frac{\chi(\Delta + \alpha)}{\alpha\Delta}\gamma \tag{6.8}$$

The rate above is $\kappa_p = 1/T_p$ where T_p is the Purcell limited lifetime, which will give us the effective cavity lifetime from Eq. 1.9. We can rewrite this more naturally in terms of the linewidth κ as:

$$\kappa_{eff} = \kappa_{int} + \kappa_p \tag{6.9}$$

From the above we see that we want $\kappa_p \ll \kappa_{int}$. To illustrate this point, lets estimate a value for κ_p from values comparable to the multimode cavity discussed in Ch.4. With an $\omega_a = 5$GHz, $\omega_c = 6.5$GHz, an $\alpha = 180$MHz, and a qubit decay time of $T^1 = 80\mu$s. From this we get a $\kappa_p \sim 80$Hz. This value is over $2\times$ higher than the internal linewidth for the cavities described in Ch.5, which have $\kappa_{int} \sim 30$Hz. This means that the effective lifetime of the mode would be decreased from > 30ms to ~ 12ms in this example. To combat this, we can decrease the coupling of the qubit by lowering the χ. We remember from Sec. 3.5.3 that the coupling strength is directly proportional to the strength of the electric field at the qubit location. We can use the simulation techniques to determine, based on the $\hat{z}-$location and the insertion length, as indicated from Fig. 6.2 what we can expect the effective χ to be.

If we want to ensure that our mode is not being limited by the qubit, we can aim for

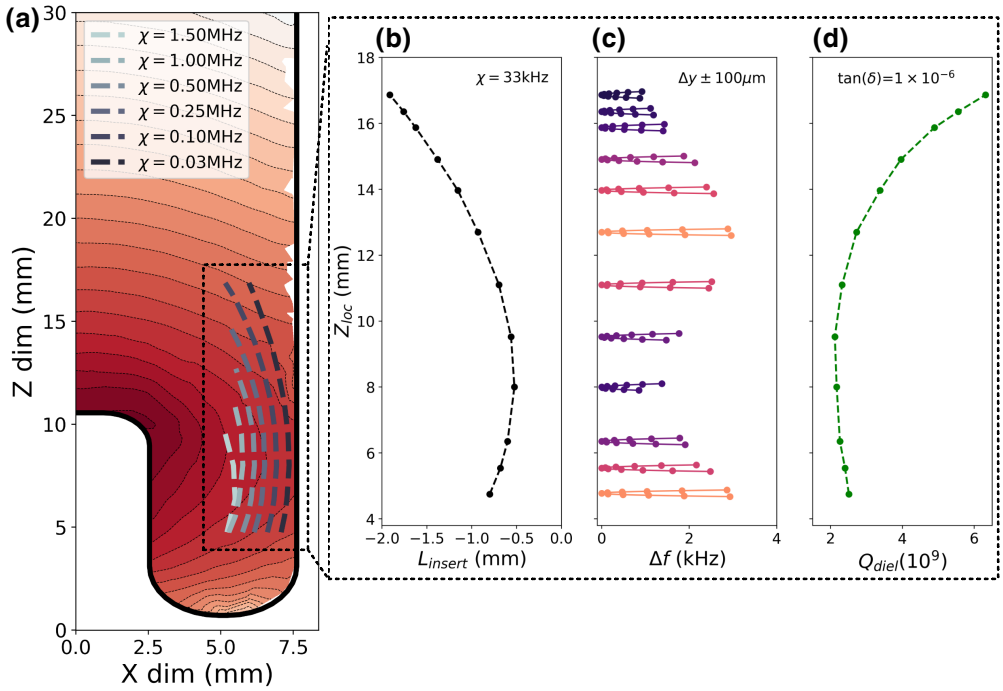

Figure 6.4: **Chip location vs** χ: (a) shows an overlay of lines of constant χ between the qubit and storage resonator, with the magnitude of the $\vec{E}$ field also shown. The inset is highlighted in (b) for the target, $\chi = 33$kHz line. (c) shows the sensitivity to vibration of the chip at various locations. (d) shows the limited Q from dielectric with a loss tangent of $\tan\delta \sim 1 \times 10^{-6}$, which represents a worst-case scenario for measured loss in sapphire, with typical cryogenic loss tangents being $< 1 \times 10^{-7}$ [146].

$\kappa_p \leq \kappa_{int}/10$. With $\kappa_p \propto \chi$, that means we need to reduce our dispersive coupling by a factor of ten from the above example, which gives an effective dispersive rate of $\chi \sim 100$kHz. Using `pyEPR`, we can look at location vs χ for our sample. Fig. 6.4 (b) shows lines representing $\hat{z}-$location and insertion length for values of constant χ. Fig. 6.4 (b) gives the shift in cavity frequency vs displacement for the chip at certain locations for a single ($\chi = 33$kHz) line, showing how the location and insertion length affect sensitivity to dephasing as demonstrated in sec. 6.2.1.

Finally, we also compare the location, based on the participation of the field in the chip, the introduced dielectric loss given that our chip has a dielectric loss tangent $\tan\delta \sim 1\times10^{-6}$ in (d). We compute the value of Q_{diel} as an extension of the concepts outlined in Sec. 2.3.1,

by using the the expression:

$$Q_{diel} = \frac{1}{p_{e,chip} \tan \delta} \tag{6.10}$$

Where $p_{e,chip}$ is the ratio of electric field participating in the volume of the chip versus the cavity volume, and is taken from Eq. 2.48 (b). Using the techniques outlined in Sec. 3.3.2, we see that, while there is a minima in the sensitivity to dephasing near $z \sim 8$mm, the static field participation in the chip's dielectric is still very high. The combination of minimizing the reverse Purcell limited lifetime, the sensitivity to chip vibration, and the dielectric loss from the chip, we want to locate the chip further into the evanescent region of the cavity above the pin. With these three things considered, we can estimate the estimated loaded quality factor from:

$$Q_L = \left(\frac{1}{Q_{int}} + \frac{1}{Q_p} + \frac{1}{Q_{diel}} \right)^{-1} \tag{6.11}$$

From this, and a choice of $\chi = 33$kHz we get a $Q_p \sim 16 \times 10^9$, and similarly, for $z \sim 16$mm, a comparable $Q_{diel} > 50 \times 10^9$, meaning that for an internal single photon quality factors, with $Q_{ext} \gg Q_{int}$, Q_L *should* be $\sim 1.3 \times 10^9$ if we are only looking at energy relaxation of the storage cavity.

6.2.3 Other geometric considerations

While the participation of the qubit and the chip add a considerable source of loss to the system, the addition of a qubit also necessitates additional holes and seams which can also add to the loss of the system. One example of this is the qubit hole, which allows for the insertion of the qubit into the cavity volume. We have shown in Sec. 2.3.3 that a hole with a cutoff frequency higher than the cavity mode will lead to exponential decay of the participating field along the length of the hole. In that section, we used a hole aligned along

the $\hat{z}-$direction of a TE_{101} cavity to illustrate this scaling. In the case of the coaxial cavity, the dipole coupling between the mode and hole is more complicated. In the optimization in the last section, to achieve the desired coupling while simultaneously minimizing the effects of dielectric loss and sensitivity to chip vibration, we located the chip in the evanescent region of the cavity, where the mode is described by a TE_{11}/TM_{01} modes of the circular waveguide. We could try and analytically determine the scaling of this coupling, which is complicated. Instead, we can see *how* the scaling effects loss for our specific geometry by parametrically sweeping the diameter and length of the coupling hole, and see if we can glean some basic intuition.

In the next section we will discuss the qubit and readout geometries in more detail, but for now we will just note that the qubit chip has a width of 2mm (which was also used in the estimates in Sec. 6.2.1). A starting point for a hole diameter is then $r = 1.5$mm. From the relation derived from Eq. 2.64 we can get the waveguide cutoff frequency $f_{c_{nm}} = p'_{nm}/2\pi r\sqrt{\mu\epsilon}$. For the 1.5mm radius hole we get ~ 58GHz, nearly a factor of 10 higher than the target mode frequency. With $\omega_{hole} \gg \omega_r$, where ω_r is the cavity mode, we have ensured that the mode's field profile should evanescently decay exponentially with hole length. The second factor, which we had originally derived using Bethe aperture coupling, is how much of the mode's field is actually coupled into the hole. In addition to this, we also have to consider the amount of field participating at interfaces. We saw that, while the Poynting vector, or the net flow of radiative energy out of the evanescent hole was imaginary, i.e. the hole only added reactance to the system, there *was* non-zero field at the end of the hole which could participate with lossy surfaces. We also saw that lossy surfaces could act as parallel resistive shunts, meaning that higher impedance at the hole end led to lower loss as a whole.

Fig. 6.5 (a) shows the boundary conditions of the qubit hole. The design shown has the hole located at $z = 16$mm, as we determined in the previous section. For practical reasons,

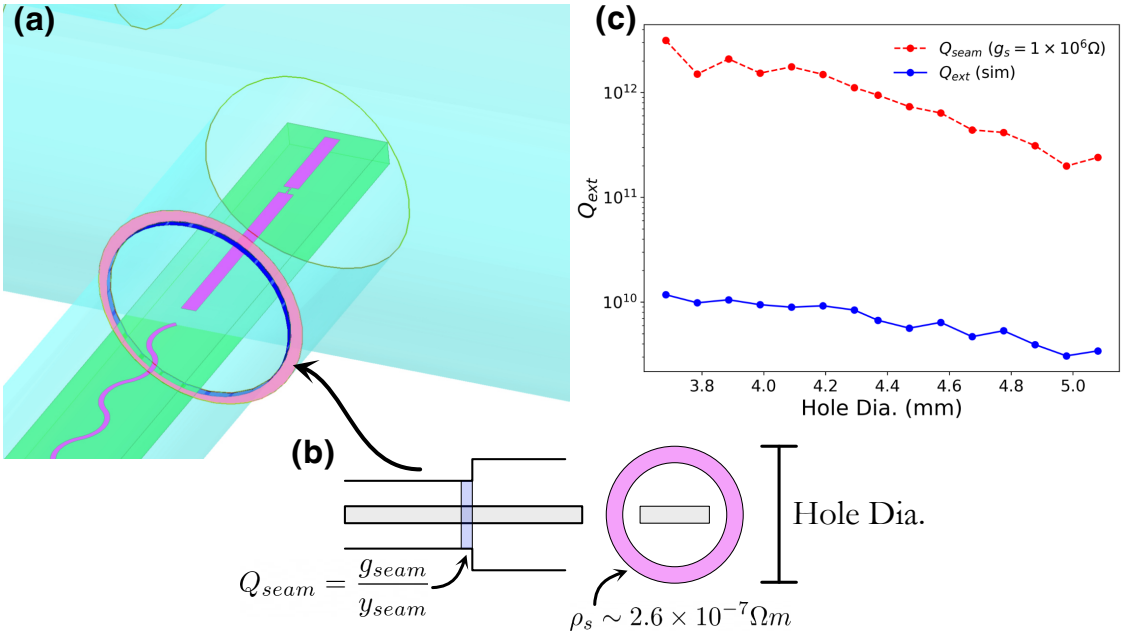

Figure 6.5: **Chip and readout section boundary conditions**: (a) shows a detail of the interface between the readout section and cavity. Highlighted in blue is the cavity-readout boundary, with a g_{seam} set to $1 \times 10^{-6}\Omega$, which has been experimentally found to be typical for indium joints in aluminum cavities [31]. In pink is the interfacing aluminum surface, with boundary resistivity set to that of aluminum, representing a worst-case scenario. (b) shows a side and front view for clarity. (c) shows simulation results for making the hole at the cavity wall larger in diameter. We see a near exponential decrease of both Q_{seam} (red) and Q_{ext} (blue), as a whole. Q_{seam} is taken from the computed admittance of the boundary from simulation data, whereas Q_{ext} is the simulated Q value extracted. This shows that the admittance at the seam is likely not dominant in the loss, but rather field participation in the qubit, or lossy readout, or field interacting with the aluminum boundary (pink) are dominant.

which will be discussed in the next section, the 2D qubit readout is contained in a secondary aluminum section that is attached to the side of the cavity via a hermetic flange. The qubit hole can then be split into two sections with an interface between the integral section (the hole that is machined directly into the Nb cavity wall), and the primary readout section which is made up of the aluminum readout attachment. This interface between the two parts can, in effect, act as a lossy seam which acts as a parallel current shunt.

We can simulate the effect of the seam loss and the effect that changing the hole diameter at the wall of the cavity. The reason we want to understand the effect of this hole diameter

is to understand how successive etchings of this cavity can effect performance. In the qubit-less coaxial cavity described in Ch.5, successive etching only changed the frequency in a meaningful way. The diameter and depth of the cavity was also changed, but these changes in geometry had seemingly little effect on performance between successive etchings. In this case the evanescent coupling to this seam *can* lead to noticeable changes in quality factor if the seam loss and material properties of the readout section, combined with the geometry, are not properly accounted for. To conduct the simulation we incorporated the qubit chip with readout into the cavity. We also added a coupling port with a 50Ω termination, setting the coupling such that the readout mode, located at $\sim$ 7.5GHz has an effective $Q \approx 1.7\times10^4$. The qubit is treated as a lumped-element inductor with a value that gives the qubit a linear frequency component of $\omega_{01} \sim$ 4.99GHz. We then set the boundary as an 25μm thick gap with a resistivity equal to the native aluminum ($2.58 \times 10^{-8}\Omega$m), and the wall is also set to that value. In the actual device the seam is made via an indium ring. The Nb-In-Al interface should all be superconducting at these temperatures, meaning that the native oxide at the interface should dominate. To simulate this we used values derived from Ref. [31] for Al-In interfaces which gave seam conductances ($g_{\oslash}$ from Eq. 2.61) of $\sim 10^6/\Omega$m, while leaving the wall to have intrinsic resistivity of aluminum. This may represent a *worst-case* scenario for our design, as the wall of the aluminum readout section at these temperatures is superconducting, albeit it has a much higher residual resistivity to the bulk Nb cavity.

In Fig. 6.5 (b) we parametrically sweep the hole diameter at the cavity wall under these different situations showing the relationship it has with hole diameter. Also listed are the evanescent cutoff frequencies for the TE_{11} and TM_{01}. Over this range of hole radius we see that the scaling is roughly linear with diameter. The seam admittance is calculated in simulation and plotted on the right axis of the plot in Fig. 6.5 (b), showing an inverse relationship indicating the current density across the seam increases as more $\vec{H}-$field participates in the hole. This calculation was done using Eq. 2.63. The dotted line represents the theoretical

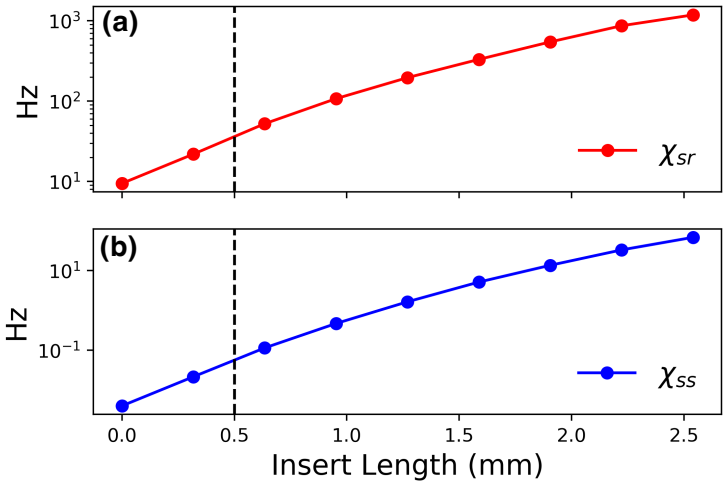

Figure 6.6: **Kerr and cross-Kerr vs insertion length**: (a) and (b) show the storage-readout cross-Kerr, χ_{sr} and the storage self-Kerr χ_{ss} respectively. The black line shows the target insertion length for a $\chi \sim 33$kHz from above.

Q from the seam alone if $g_{\oslash} \sim 1 \times 10^6$ is fixed, showing that most of the contribution in Q_{int} degradation is actually the evanescent participation with the lossy aluminum readout section. We compare this with and without the addition of the lossy exposed surface of the readout section to determine the contribution of participation from the aluminum readout attachment and the seam itself assuming the readout coupling is included. We see that the contribution of the seam is negligible when compared to the participation of the evanescent field interacting with the higher-loss aluminum waveguide section of the readout.

6.2.4 Readout-storage cross-Kerr and induced readout loss

In our simulations of the qubit-storage cavity coupling, we can also extract the effective cross-kerr interaction strength between the readout and storage cavities. These terms give rise to a dispersive shift, mediated by t he qubit, which leads in a shift of the storage cavity for a given occupation of the readout. We have seen these terms previously when describing the multimode Jaynes-Cummings Hamiltonian in Eq. 4.7, which manifested in a state-dependent shift with strength k_{sr}, where we have replaced indices m and n for s and r respectively.

Fig. 6.6 shows the predicted cross-Kerr interaction strength as a function of insert position at the optimized $\sim$ 16mm $\hat{z}$-position.From Ch.3 we showed that the cross-kerr can be expressed as $\sqrt{\chi_{ss}\chi_{rr}}$ where χ_{ss} and χ_{rr} are the storage and readout self-Kerrs respectively. We see that the storage-readout Cross-Kerr value is on the order of $10 - 100$Hz. This means that a direct readout of the storage cavity state via the state-dependent shift of the readout is all but impossible, since the shift is much less than the readout cavity linewidth, which for the design shown has a linewidth of $>$ 100kHz. We will discuss how we actually determine the storage-cavity occupation in a later section. This also highlights another way of seeing the shot-noise dephasing problem discussed in some detail in Ch.4. In this case the storage dephasing rate follows the form of Eq. 6.2.1, where the thermal occupation of the readout leads to enhanced dephasing of the qubit. This, of course is inherited by the storage cavity. In this limit, however, the approximation used in Ch.4 does not hold, since we have purposely designed χ to be very small, meaning that for a 100μs qubit, $\chi \sim \gamma$.

6.3 Design Considerations

In the previous section we explored and made an attempt to optimize the design parameters to minimize unwanted additional qubit induced loss in our storage cavity. With these considerations in mind, combined with the insights gained in Ch.5 concerning the creation of ultra-high Q niobium cavities, we can now create our quantum memory. In this section we will discuss the realized design and construction of this system, in particular the 2D planar readout, the additional considerations for hermetic sealing, and the unique qubit mounting design used to facilitate these two design constraints.

6.3.1 Qubit and readout design

The 3D transmon qubit used in the Nb cavity shares the same basic geometry as that discussed in Sec. 4.5.1. In contrast to the multimode cavity, which used a 3D readout design,

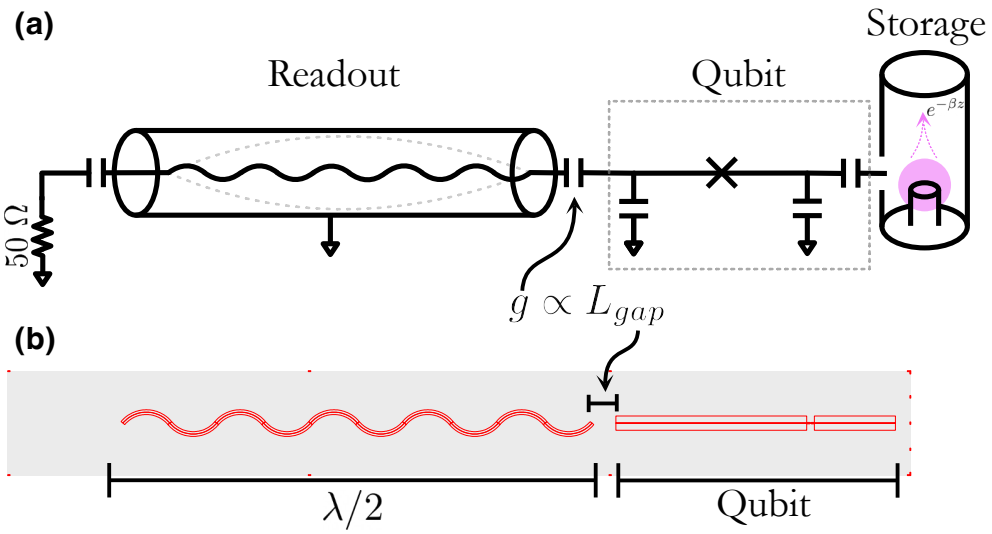

Figure 6.7: **2D readout, orientation of chip and coupler** (a) layout of qubit and readout (b) pin geometry and coupling vs distance:

with the readout made using the flute technique, the niobium quantum memory makes use of a 2D readout design, following the design discussed in Ref. [17], which utilizes a coaxial stripline resonator that is capacitively coupled to the qubit. This was done, as with *most* aspects of the cavity's design, to accommodate the ability to be hermetically sealed for the preservation of the Nb cavity surface, with the 2D design requiring fewer indium seals.

Fig. 6.7 (a) shows the 2D design of the chip. The mode of the readout is a $\lambda/2$ stripline mode, with the length setting the frequency. By surrounding this stripline with groundplane via the mounting hole, the effective electric field profile and resonant characteristics approach that of a shunted section of coaxial waveguide, allowing for predictable control of frequency. To make the stripline more compact a gentle serpentine structure was added to the design. Optimetric simulations were used to produce a readout resonator with a center frequency of ~ 7.5GHz.

Coupling between the qubit and the readout was controlled by adjusting the spacing of qubit pad and the stripline. As discussed in Ch.3, $g^2 \propto \chi \propto |\vec{E}|^2$ at the qubit. The hole in which the chip was placed was designed with the aforementioned cutoff frequency of $f_c > 50$GHz. Because of this, the readout mode is evanescent in the waveguide section

separating the qubit pad and the readout, meaning that $|\vec{E}| \propto e^{-\beta z}$, as shown by Eq. 2.67. This has been shown to be in good agreement with measured systems in Ref. [17], and provides accurate control and the creation of strong readout-qubit dispersive couplin χ_{rq}.

To couple power into and out of the readout, and provide a charge-drive for the qubit, a port is added that allows for the evanescent coupling of the coaxial stripline mode to the TEM mode of a coaxial stub-coupler. In literature, separate strongly coupled readout, and weakly coupled drive ports are used to make the direct implementation of qubit operations more straightforward. Due to the hermetic sealing, and the additional contraints that requirement adds, only a single port was added to the readout in which both drive and readout operations would be fed. In the later measurement section we will see how this port was utilized. The hermetic connectors utilized a coaxial waveguide with a $\oslash$ 0.89mm provided the correct impedance for the $\sim 400\mu$m diameter coupling pin. The extremely high cutoff of this port made coupling highly sensitive to the coupling length. Simulation was used to estimate the correct length for the desired readout coupling, as shown in Fig. 6.7 (b)

The qubit pad and readout was patterned from evaporated niobium. The recipe is similar to that discussed in Sec. 4.5.1, with the same thickness and quality sapphire used, and the same junction fabrication parameters, however the use of the Nb base layer necessitates the use of fluorine reactive ion etching instead of chlorine plasma. For the purposes of this design the choice of niobium over aluminum baselayer was an arbitrary one. Qubit junction parameters were computed using both q3d and pyEPR simulations to determine optimal E_c and E_j parameters for a qubit $|g\rangle - |e\rangle$ frequency of between $4 - 5$GHz. The details of this simulation are discussed in Secs. 3.4.2 and 3.5.3.

6.3.2 Hermetic design and optimized qubit clamping

In Sec. 5.2.1 we demonstrated a simple technique for hermetically sealing the cavities, and later showed that sealing, while not strictly necessary, can help ensure that the deleterious

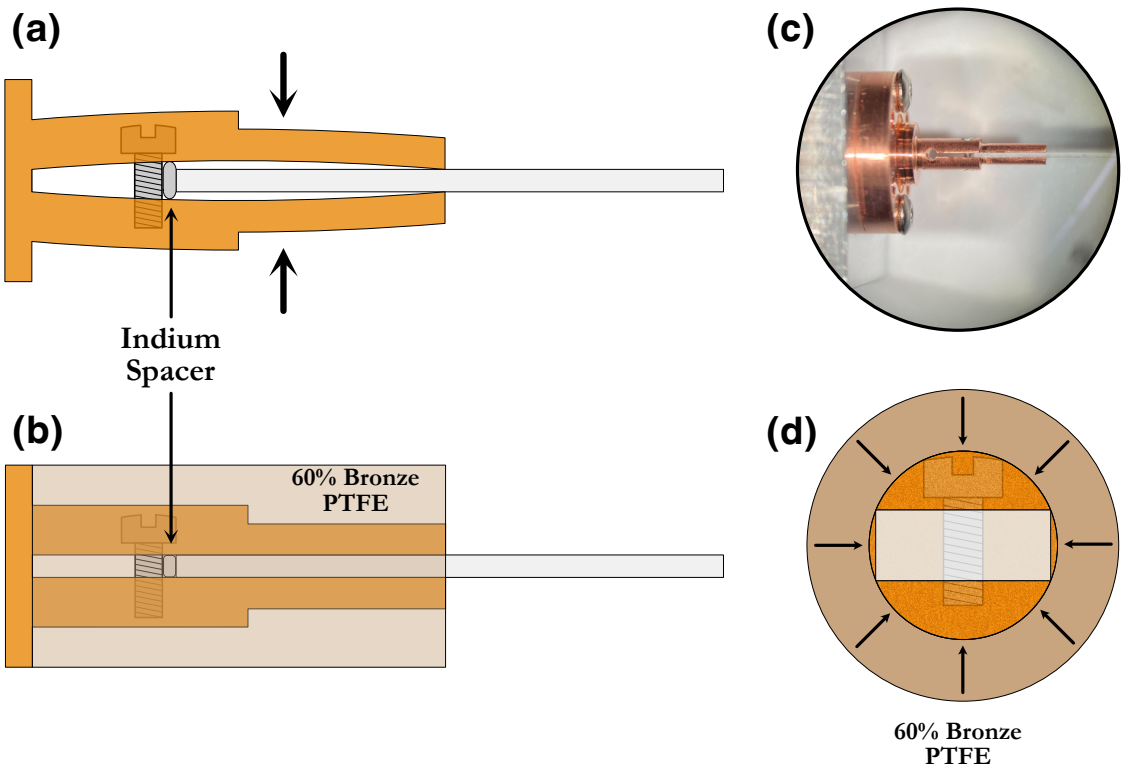

Figure 6.8: **Schematic of 1 piece design and clamping action** (a) action of clamping (b) increased hold force from low-temperature contraction:

oxide growth is kept minimal. In the cavity case, this sealing was simple, with only the top of the cavity, and the coupling port, being sealed via an indium wire. In the qubit-cavity case, the addition of the readout section, which itself requires a coupling port and a way of holding the qubit makes the task of hermetic sealing more difficult. In addition to creating a gas-tight seal, the hermetic sealing also ensures a light-tight environment, which we saw in Sec. 4.5.2 is an important requisite for reducing qubit dephasing.

The most challenging part of this system is the qubit mount. Previous two-piece mount designs would be difficult to seal, given the perpendicular sealing surfaces. To solve this problem a new single piece qubit mount was created. Fig. 6.8 (a) shows a schematic of this design. A circular flange with an integral indium o-ring groove is attached to a post with a slot cut to the exact thickness of the sapphire chip using a precision slitting saw. The tines of the clamp are bent to have a slight convex geometry. A small screw (#000 – 120

thread) flattens the tines around the chip, with the convexity and natural elasticity of the copper providing a pre-loading force to the clamp. To combat differential thermal expansion and the loosening of this clamp upon cooldown, a 60% bronze-filled PTFE compression ring, machined to a slip fit, is placed over the clamping section. The integrated thermal contraction of copper between 298K and 4K is $\Delta L/L \sim 326 \times 10^{-5}$ [52], whereas polytetraftuorethylene (PTFE) in its virgin form has an integrated thermal contraction nearly 5× greater ($\Delta L/L \sim 180 \times 10^{-4}$) [22]. The addition of $40-60\%$ bronze adds dimensional stability and toughness while reducing cold-flowing of the PTFE at low temperatures. This does however reduce the effective thermal expansion by $\sim 20\%$. Because of the precision fit of the parts, no additional indium or compliant material was required for adequate clamping. The cylindrical design also ensure even pressure application on the bearing surfaces, reducing the chances of fracturing the sapphire chip. A lip at the rear of the mount was made to be a precision slip fit with the rear of the aluminum readout section as to ensure precise indication of the chip inside of the section of cylindrical waveguide while reducing constraints.

The readout coupling port was also made hermetic. A copper flange was made with an hermetic male SMP connector, with a 3.8mm×400μm pin[2], was soldered into the flange. A cylindrical boss cut into the flange was machined to a precision fit with a matching hole on the readout side to precisely indicate the pin location, while an indium o-ring groove was used to consistently seal the connector using 0.5mm diameter indium wire. This arrangement is shown in Fig. 6.9 (a), with a detail showing the sealing surfaces of the chip holder and readout port flange. A mounting flange was also added to the readout section to seal the readout to the cavity side.

The entire cavity and readout assembly was bolted to an OFHC copper (alloy C101) mounting bracket for adequate thermalization, with a sub-bracket thermally linking the qubit clamp to the remainder of the mounting assembly. This detail is shown in Fig. 6.9 (b).

2. Pasternack part PE44921

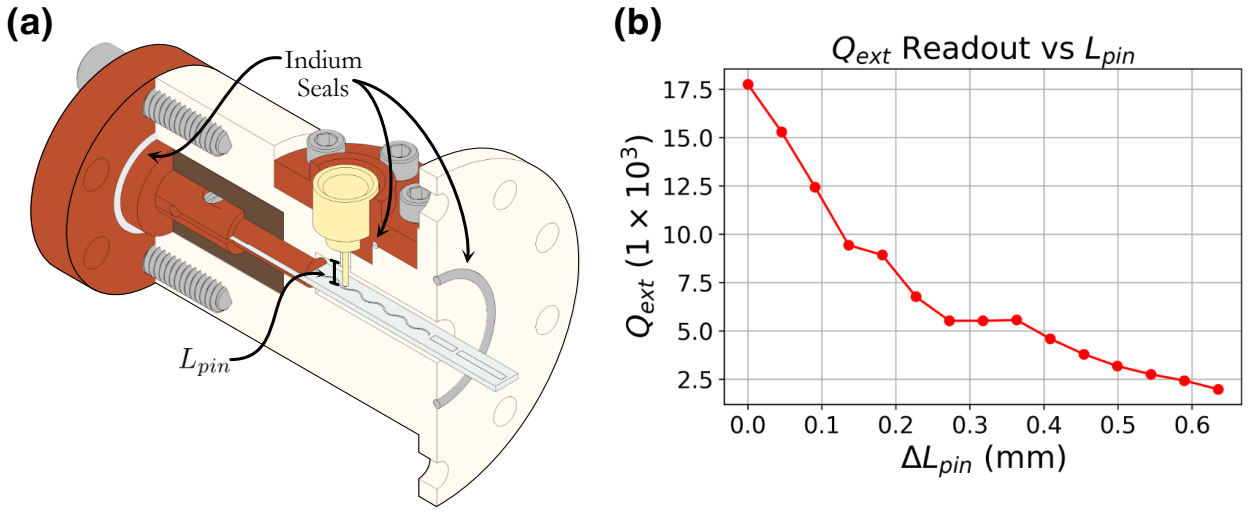

Figure 6.9: **Cutaway for 3D cavity and qubit** detail showing the hermetic readout coupler:

The bracket also supported Eccosorb low-pass IR filters. The entire assemly was mounted inside of a 2-layer μ-metal shield with an internal copper liner to thermalizer any errant IR photons that still may be incident on the cavity.

6.4 Measurement

The qubit-cavity system measurement differs substantially from the cavity measurements discussed in Ch.5. The basic setup for the readout of the qubit, as we will see in the next section, shares the same basic setup as the cavity measurement, however we cannot interrogate the storage cavity directly. Instead, we can use the unique properties of cavity QED (and by extension, 3D cQED), to measure the cavity *via* the qubit. In this section we will discuss how this is done, some of the unique challenges involved with such long-lived cavity modes (and weakly coupled qubits), and discuss the first data from the system, showing its potential as a store of quantum information.

6.4.1 Measurement setup

The cryogenic side of the measurement setup for the coax and qubit, shown in Fig. 6.10, shares the same basic topology as the setup previously used in Sec. 5.3.1. Instead of measuring the cavity in reflection, however, the output line is coupled into the readout resonator port. All qubit drive signals are also passed into this port via the directional coupler, with the qubit being driven off-resonantly through the readout. A separate drive line is also added directly to the cavity to allow for independent displacement drives on the storage mode. This differs from the setup used in Sec. 4.5.2, where storage cavity drives were similarly driven either off-resonantly through the readout, as was the case for the storage Rabi drive, or the cavity was manipulated via the qubit using $|f0\rangle - |g1\rangle$ sidebands, as discussed in more detail in [39, 37].

The largest changes to the measurement chain was done at room temperature. For time-domain and pulse measurements, instead of sequencing pulses and conducting readout via the PNAX, as was done in 5.3.1, am FPGA based Xlinx RFSoC (model ZCU216) was utilized. The RFSoC hardware was used in conjunction with the Quantum Instrumentation Control Kit (QICK) to generate time-domain pulse-sequences [173]. The ZCU216 features 16×2.5 GS/s DDC and 16×10 GS/s DUC channels which are configured into 16 ADC and 16 DAC channels respectively. The added XM655 mezzanine add-on board breaks out the double-ended I/O into an 8DAC/8ADC single-ended channels via on-board baluns, with each block having frequency ranges of < 1GHz, $1-4$GHz, $4-5$GHz, and $5-6$GHz [3].

For the readout, an external local-osscilator (LO), taken from a SignalCore SC5511A, is mixed into the output DAC intermediate frequency (IF) using double balanced $2300-8000$MHz mixers [3]. The readout signal is split and the input signal is mixed down using a conventional homodyne topology. One point of note in this scheme is the absence of analog IQ-mixers, reducing the need for mixer calibration. The $I-Q$ quadratures are obtained

3. Mini-circuits part ZX05-83-S+

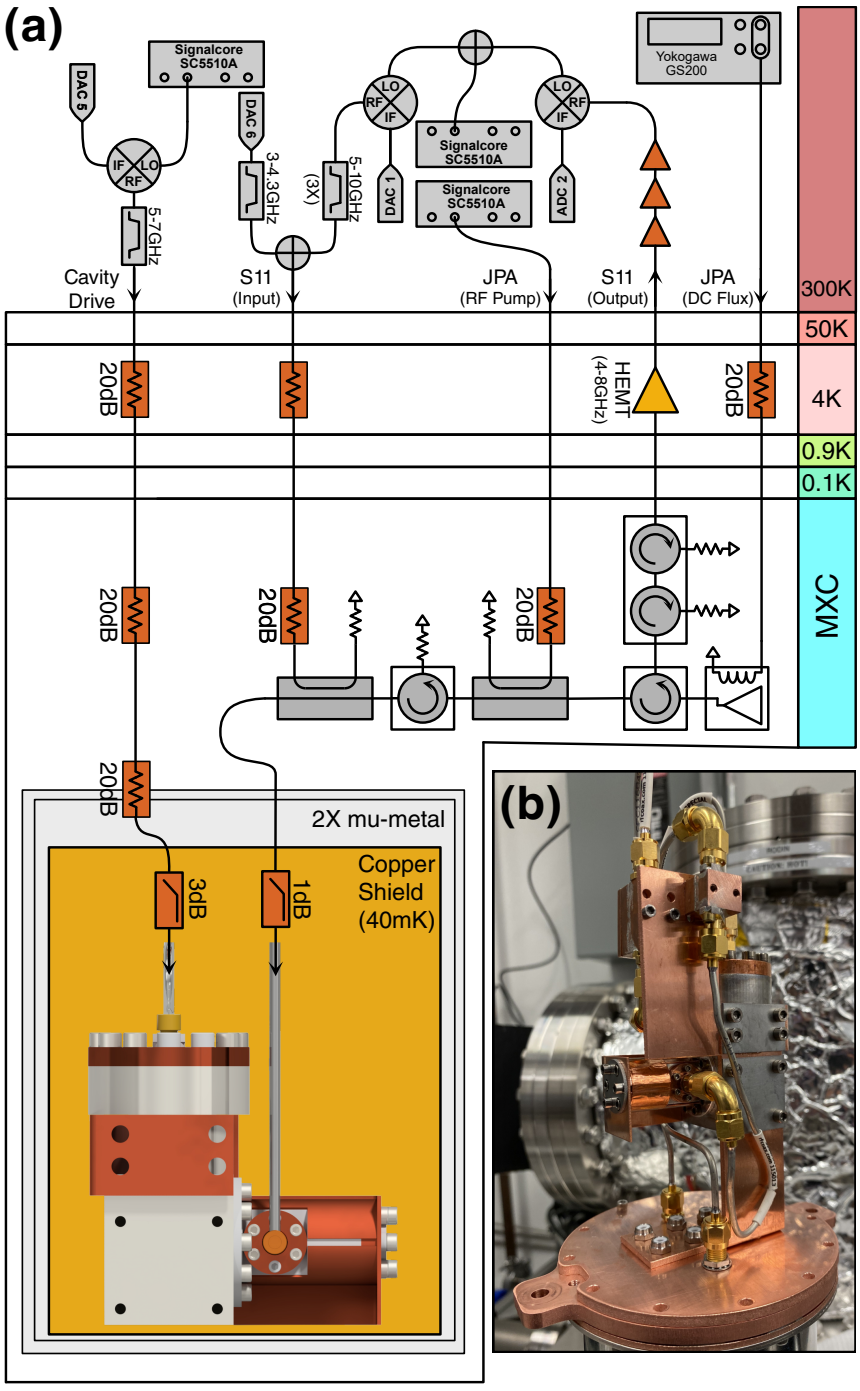

Figure 6.10: **Measurement setup for cavity and qubit**

in the digital down-conversion process block of the ADC, which incorporates digital IQ downconversion as part of the FPGA architecture [173]. To reduce unwanted sidebands $3 \times 5 - 10$GHz bandpass filters are placed on the readout output. A separate qubit drive is

directly synthesized from the $1-4$GHz output of the ZCU216 (second nyquist zone), and passed through a narrow bandpass filter of $3-4.3$GHz. Finally, a separate cavity storage drive is sythesized via mixing a second SC5511A LO into another double-balanced mixed and filtered through a $5-7$GHz bandpass filter.

6.4.2 Readout and Qubit Characterization

Initial measurements were done under continuous power to determine the frequency of both the readout resonator and, using two-tone spectroscopy, the qubit frequency. After this the readout spectroscopy was done using the RFSoC measurement setup shown in Fig. 6.10 using discrete pulses in the time-domain. The result of this is shown in Fig. 6.11. Following readout spectroscopy, the second qubit drive tone is added. The spectrscopy, shown in Fig. 6.11 (b) for the $|g\rangle - |e\rangle$ transition, can then be used to optimize readout signal-to-noise. By adding a $\pi-$pulse of the $|g\rangle - |e\rangle$ qubit transition, before doing readout spectroscopy we are able to determine the dispersive coupling of the readout and qubit from the splitting of the peak, which is equivalent to $2\chi_{qr}$[4]. We can also park the readout tone at the $\omega_r + \chi_{rq}$ to maximize the SNR of the qubit readout.

Next we can conduct time-domain experiments to study the time-dependent evolution of the qubit. By coherently driving the qubit we can induce a Rabi oscillation, in which the qubit rotates between $\pm\hat{z}$ [126]. This rate, shown in Fig. 6.12 is power dependent, and by changing the qubit drive strength it is possible to induce oscillations that allow for swapping the qubit from $|g\rangle - |e\rangle$ in only $100-200$ns to allow for fast resolved qubit readout. Once the π-pulse has been calibrated we can conduct experiemnts tha study the evolution of the qubit due to decoherence mechanisms. By initializing the qubit in $|e\rangle$ and waiting some time before doing a resolved π-pulse onto the qubit, we can determine the decrease in excited-state population due to energy relaxation of the qubit, or the T^1. Finally, we want to study the

4. In Ch.3 and in the above sections we have denoted χ as equivalent to two times the dispersive shift as defined in Eq. 1.24, however here we are using that definition for χ_{qr}

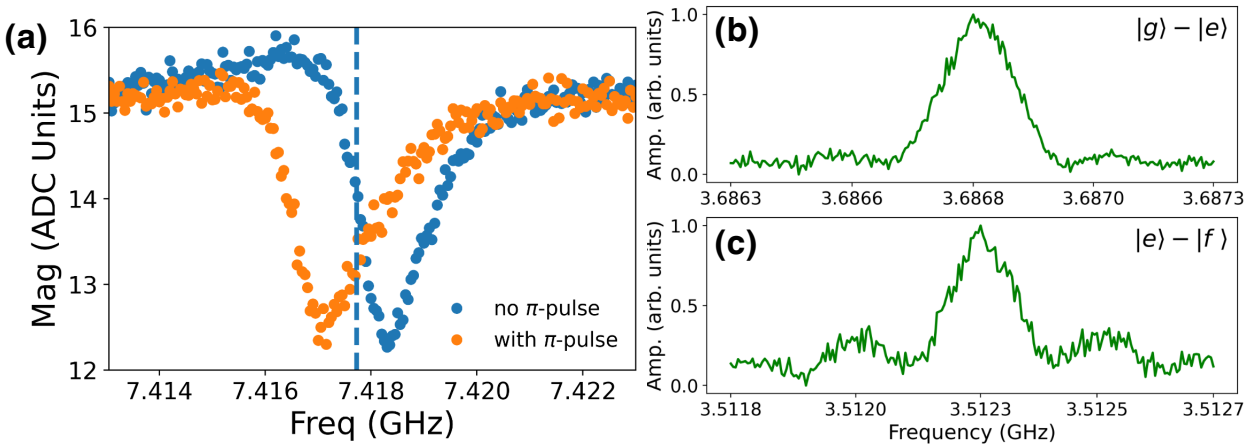

Figure 6.11: **Readout and qubit spectroscopy** (a) shows the splitting of the readout resonator line, with a splitting of χ. The dashed line represents the point of maximum readout fidelity. (b) and (c) are the $|g\rangle - |e\rangle$ and $|e\rangle - |f\rangle$ wubit peaks respectively. The anharmonicity from spectrscopy for this sample is $\alpha \approx 175$MHz

dephasing of the qubit, or T^{2*}. To do this, as is canonical for qubit measurements, we apply a $\pi/2-$pulse to the qubit, which produces a mixed $(|g\rangle + |e\rangle)/\sqrt{2}$ state. After some delay we can use a second $\pi/2$-pulse to map the state back onto the $|e\rangle$-state of the qubit. By detuning the drive strength from the qubit's $|g\rangle - |e\rangle$ transition frequency we can induce osscilations, known as Ramsey fringes, into our measurement. This serves teh purpose of reducing the effects small-detuning from the target qubit frequency, and inversely allows us to accurately determine the qubit frequency by fitting the Ramsey frequency, which corresponds to the detuning. This difference from the input Ramsey frequency and the measured oscillations can be used to correct the qubit ω_{ge} transition. The exponential decay of these fringes also give the T^{2*} of the qubit, or dephasing time. Rabi, T^1, and Ramsey measurements are shown in Fig. 6.12 (a), (b), and (c) respectively.

With the qubit frequency optimized, and π-pulse calibration done, we can also optimize the high-power dispersive readout to obtain high single-shot fidelity by tuning the readout power and integration time of the ADC [151]. Doing so gives us single-shot fidelity, shown in Fig. 6.12 (d) of $> 75\%$ without the need for the JPA which is shown in Fig. 6.10 (a).

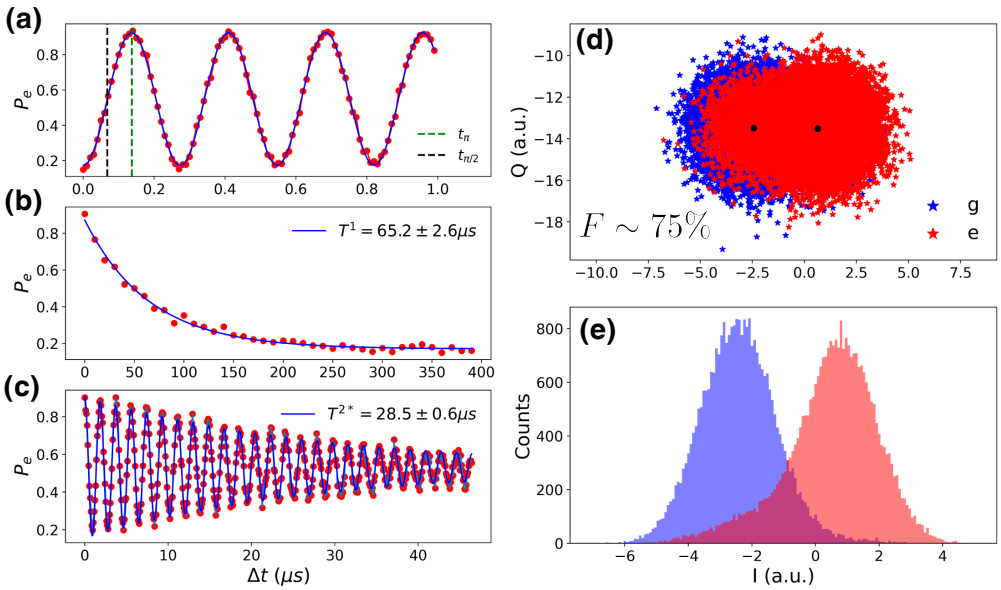

Figure 6.12: **Qubit Characterization and π-pulse calibration** (a), (b), and (c) are Rabi, T^1, and Ramsey/T^{2*} respectively. (d) and (e) are single-shot qubit spectroscopy data in $I-Q$ and histogram form respectively, showing fidelity of $F \sim 75\%$

6.4.3 Storage characterization with qubit

With basic readout and qubit characterization done, we can begin the process of characterizing the storage resonator. Instead of directly interrogating the storage resonator, as we did in Ch.5, we will now use the dispersive interaction of the qubit to do both basic resonator spectroscopy and measure decoherence. The first part, the resonator spectroscopy, is done via a second direct storage drive line. The signal is synthesized by upmixing the RFSoC IF to the frequency of the resonator. The RF storage is swept, with a resolved readout on the qubit $|g\rangle$ state. The dispersive shift of the due to photon occupation of the storage leads to a characteristic resonance at the cavity frequency. To better resolve this peak and accurately determine the frequency, the storage cavity strength and drive time are reduced until the peak is no longer homogeneously broadened. This process, and the cavity spectroscopy is shown in Fig. 6.13 in (a) and (b) respectively.

Once the cavity frequency is determined, we must calibrate the coherent drive strength of the direct cavity drive. To do this we can use the dispersive interaction to calibrate

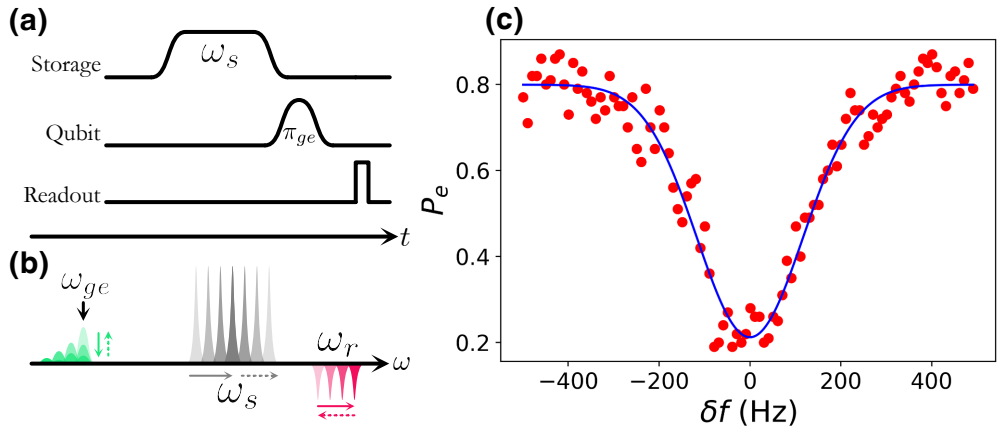

Figure 6.13: **Storage cavity spectroscopy pulse sequence and spec** : (a) shows the generalized time-domain pulse sequence for the storage cavity spectroscopy experiment. A coherent tone applied to the storage cavity is swept in frequency, followed by a π-pulse on the qubit $|g\rangle - |e\rangle$ peak, followed by readout. The presence of photons in the storage will lead to a displacement in the qubit peak, leading to a drop in signal. This process, however, only happens when the drive is on resonance with the cavity. (b) shows this process in the frequency domain. (c) is the measured spectroscopic data, with a cavity $\omega_0 = 2\pi \times 6.655$GHz.

the drive strength. Normally, the interaction is designed such that qubit linewidth γ is much greater than χ, however, as demonstrated, we have intentionally made χ between the qubit and storage cavity as small as possible to minimize the reverse Purcell effect. Because of this, $\chi \approx \gamma$, which means we need to resolve the intrinsic qubit linewidth as closely as possible. The width of the measured qubit peak is a convolution of the pulse (in the frequency domain) and the intrinsic qubit lineshape. For a Gaussian pulse, the width of the frequency domain lineshape is simple $1/\sigma$, where σ is the time-domain width of the pulse envelope. Because of limitations in the memory buffer of the RFSoC, the max Gaussian width, at these frequencies, is only 2.5μs, which is much less than the extracted T^{2*} from Fig. 6.12, which sets the qubit's intrinsic linewidth. To combat this, a flat-top pulse with a gaussian ramp is used. This new pulse, in frequency space, takes the form of a function, with the width of the function being $1/2T_p$, where T_p is the flat-top pulse width. This means that, with a flat-top pulse $T_p \geq T^{2*}$, the linewidth should approach the intrinsic lineshape.

This, however, is only half the problem, since the power of the qubit pulse also changes

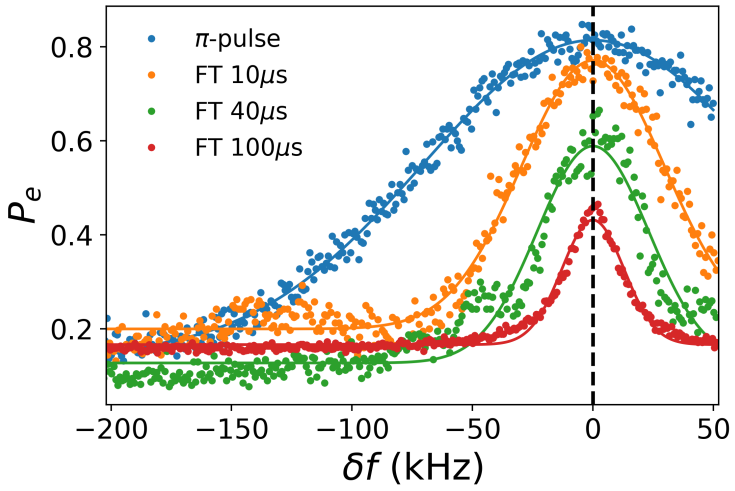

Figure 6.14: **Calibration of resolved $|g\rangle - |e\rangle$ peak** : The plot shows the effect of the pulse bandwidth on the peak width of the qubit $|g\rangle - |e\rangle$ transition. The blue line represents a typical gaussian resolved π-pulse, with a $\sigma = 2.3\mu$s, while the subsequent lines show flat-top pulses with varying lengths and gaussian rise-times of 2.3μs. We see that for total pulse lengths less than MAX(T^1, T^{2*}), the contrast remains the same (orange), with no change in peak height. For pulses $\sim$MAX(T^1, T^{2*}), there is a clear loss of SNR since the qubit is beginning to decay into a mixed $|g\rangle - |e\rangle$ state (green). For times much longer than the coherence time we see that the fidelity has dropped to roughly half (red), showing the qubit is in a fully mixed state, although the overall peak width is now comparable to χ, allowing for resolved number state spectroscopy.

the intrinsic linewidth. To see why, we must remember that the coherent drive produces a cohernet state in the resonator with a poisson distribution, as shown in Eq. 5.6, that has a photon variation of $\sqrt{n}$. Refs. [72] and [164] show that for situations where the qubit probe time is longer than the readout cavity decay rate, the width of the qubit peak scales linearly with mean photon number, while also approximating a lorentzian line shape. With this in mind we must reduce the power under the induced linewidth, $\Gamma_m \sim 2\bar{n}\kappa_r\theta_0^2$, is less than the intrinsic qubit linewidth, where $\theta_0 = \chi_{qr}/\kappa_r$, and κ_r is the readout decay rate. Fig. 6.14 (a) shows the effects of power reduction and pulse length on the intrinsic linewidth, while (b) shows the effect that these changes have on the number resolved peaks of the storage resonator under a coherent displacement.

Now that we can resolve the individual number-split peaks, we can conduct calibrations

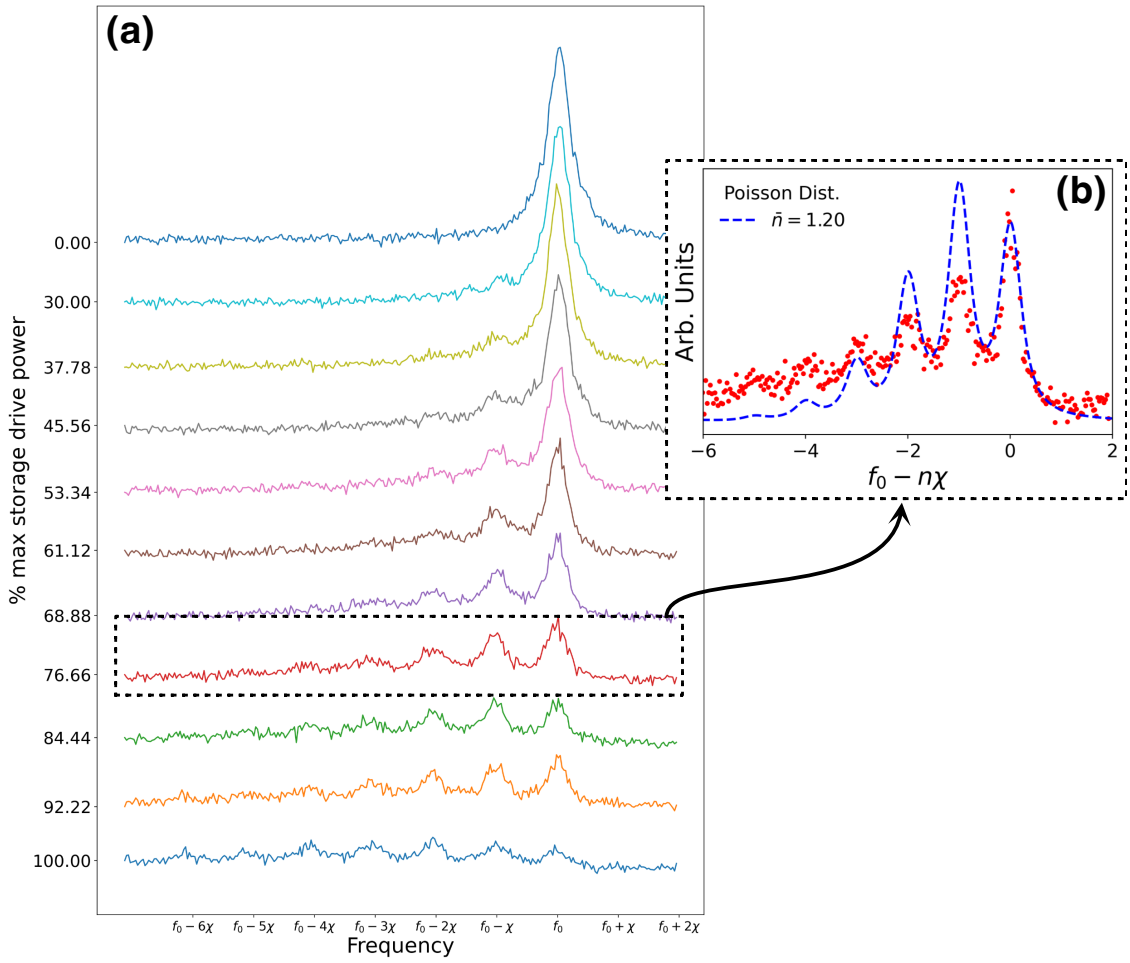

Figure 6.15: **Number splitting vs DAC output.** : (a) shows qubit spectroscopy under varying drive power on the storage resonator, with the power scaled to the highest power listed along the x-axis. (b) shows a detail of one trace, overlaid with the predicted Poisson distribution for a mean photon number of $\bar{n} = 1.2$, highlighting the deviation from the expected coherent distribution. The falloff in amplitude with successively higher number-states indicating a thermal instead of coherent state. The measured $\chi \approx 29$kHz for the above data, showing the ability to resolve weakly dispersive coupled systems.

of the coherent drive strength. Ref. [163] showed the expected distribution of resolved number-states of a cavity under a coherent drive. From this, in theory, we should be able to calibrate the drive strength to achieve a known mean photon number and, for our purposes, to extract the storage cavity-qubit χ and to conduct cavity ringdown using the displacement of the qubit peak. We show in Fig. 6.15 the evolution of the resolved peaks vs DAC output power. The first thing to note is that the splitting of the peaks is $\chi \sim 35$kHz, which indicates

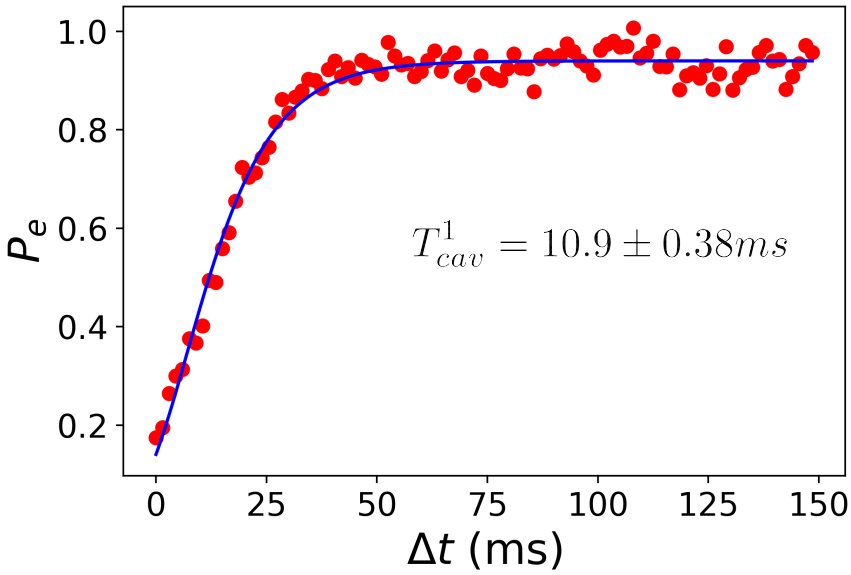

Figure 6.16: **Storage cavity T^1:** Using a displacement drive with a strength that fully displaced the qubit $|g\rangle - |e\rangle$, ringdown was done. This process is identical to the steps described in Fig. 6.13 (a), except the storage drive frequency ω_s was kept constant and on resonance with the cavity, while the delay between drive and resolved qubit readout was delayed. The increase in P_e indicates a drop in mean photon number in the cavity. The fit, done using Eq. 6.12, shown in blue, gives a $T^1_c av = 10.9 \pm 0.38$ms, which is a record for coaxial geometries, although it does underperform the theoretical internal quality factor from Ch.5 by a factor of 3.

the success of the above simulation in the calibration of the qubit-storage dispersive coupling in Sec. 6.2.2. Another observation is that the distribution does not follow the poisson-like distribution expect of a coherent-state, but instead appears as an almost thermal state [163], which is shown in Fig. 6.15 (b).

We can use the above calibration to ensure that the cavity is sufficiently displaced such that the qubit, when pulsed at the $|g\rangle - |e\rangle$ transition, will have zero amplitude. Following this procedure, we can do a time domain measurement on the displaced storage cavity, where we do two resolved measurements of the qubit $|g\rangle - |e\rangle$ transition separated by some time interval. In doing so, we can watch the evolution of the cavity state in time. The decay of the cavity coherent state can be modelled using the following function, as derived from Ref. [150].

$$P_0(t) = e^{-|\beta_0|^2 \exp(-\kappa_s t)} \tag{6.12}$$

Where κ_s is the storage cavity linewidth, and subsequently $T_s^1 = 1/\kappa_s$. The decay of the resonator under this classical displacement is shown in Fig. 6.16. When fit, the data gives a $T_s^1 = 10.9 \pm 0.38$ms. This value, which is less more than half of the loaded cavity lifetimes found in Ch.5, is still nearly 10× greater than nearly identical aluminum cavities after a qubit has been introduced [150]. This seemingly indicates that there are additional aspects of the above system that must be considered for a full optimization of this system for use as a quantum memory

6.5 Discussion and outlook

The above results are only preliminary, and represent only a few cooldowns. A more careful accounting of loss in the system would be necessary to understand why the performance above is less than what is theoretically possible based on the results of chapter 5. The first, and most glaring omission from the above data is the measurement of cavity T^{2*}. Traditionally T^{2*} is done by preparing a $(|0\rangle + |1\rangle)/\sqrt{2}$ cavity state using SNAP gates, blockade, as detailed in Ch.4 or using the $|f0\rangle - |g1\rangle$ sideband transition [124, 38]. With the extremely weak dispersive coupling, and the necessary hierarchy of scales, requiring the rabi drive rate to be much less than χ, the preparation time for such a state would mean that preparation fidelity, even with these long coherence times, would be poor. SNAP gates [87] would be faster, however they are limited by timescales of $T_g \leq 1/\chi$. The final option, $|f0\rangle - |g1\rangle$, similarly has a limitation to the Rabi drive rate, however for simple state-preparation may be the easiest to implement. Another solution which has recently been demonstrated for weak-dispersive systems, is the use of Echo Conditional Displacement pulses to gain an enhancement in the interaction rate that scales like $|\alpha|$, where α is the

coherent drive strength on the cavity [66]. A challenge with this technique is the creation of large displacement drives in a short period or time. To gain a real advantage it is necessary to displace the storage resonator to $|\alpha| \geq 100$, which corresponds with $\bar{n} \sim 10000$ photons. Doing so with the coupling and internal quality factors of our cavity may be difficult, since the drive strength required to displaced the cavity in timescales on the order of the gate time would require large amounts of incident power at the cavity. A final and very rudimentary option is to simply drive the cavity coherently to a number state that has a non-equal mixture of $|0\rangle$ and $|1\rangle$. If the state is coherent, it should be possible to do so at mean photon numbers of $\bar{n} < 1$, however as shown in Fig. 6.15 (a), there is no drive power that seems to produce a coherent state, but instead the number splitting seems to be thermal in nature. If, however, one were to make a state with only a mixture of $|0\rangle$ and $|1\rangle$, it would be possible to do a Ramsey style experiment, with the only loss being that, because the state isn't an equal mixture, the Ramsey fringes would have subsequently less contrast.

The thermal state of the number splitting, and the decrease in coherence, may indicate a more pervasive issue, namely that the system is rather hot. During the above experiments the qubit temperature fluctuated between $\sim 50 - 65$mK, or photon occupations of $> 3\%$. With an $E_c \sim 175$MHz, and a $\omega_{ge} \sim 2\pi \times 3.6$GHz, the qubit has a ratio of $E_J/E_C \sim 60$, which makes the qubit's $|e\rangle - |f\rangle$ transition far more susceptible to dephasing. The main way to improve these is to simply switch the qubit for one with higher E_J, and to improve the thermal environment of the qubit. The fridge this sample was cooled with, an Oxford Triton 200, has a base temperature of only 35mK on a *good* day, and no MXC shield, making a change in fridge an easy solution to this problem. With that change, a better characterization of the novel qubit clamp can also be done to determine if the one-piece design is on-par with existing two-piece solutions. Once this issue is sorted—and a robust protocol for doing cavity Ramsey experiments is developed—a more complete picture of the loss in this system can be made.

Finally, developing an active reset protocol may be necessary as experiments become more complicated. The long decay times, even on the order of 10ms, means that experiment times have become *excruciatingly* long. In fact, one of the main limitations in completing these experiments in the last few months has been the time it takes to simply conduct an experiment, with cavity T^1 experiments taking upwards of an hour for the date in Fig. 6.16. With an active reset protocol, likely that swaps the cavity state into the qubit, would greatly improve usability of the system.

In total, the above data is just a glimpse at the possibilities. With the allure of cavity T^1s in excess of 25ms, based on Ch.5s resonator measurements, and the additional possibilities for improving those lifetimes to $T^1 > 100$ms, means that this is only the first step in a series of future experiments and research that may see a $100\times$ improvement in 3D quantum memories. With the addition of SRF technology, with T^1s of 300ms or greater, there is a lot of headroom to see even more improvement. One question we can ask is, will there be a point that the drawback in weak dispersive coupling to save cavity coherence comes at too high a cost? Only time will tell, but either way, it is an exciting moment for superconducting quantum memories as a whole.

CHAPTER 7

FUTURE OUTLOOK

The work outlined in the previous chapters represent a next step in the evolution of 3D cQED quantum systems. With improvements in coherence that correspond with over an order of magnitude increase in coherence times, these techniques will usher in a new and exciting chapter in superconducting quantum systems and quantum optics. In addition to this, the advent of flexible multimodal quantum cavities, with high cooperatives across many modes, represent a completely new class of scalable quantum memory architecture. While these advancements in superconducting cavities, and the techniques, hardware, and software development they have spurred, are exciting, there are still a number of things that can be done to truly see them reach their full potential. In this chapter we will briefly go over some of these future (and concurrent) research objectives, and discuss the exciting future for both multimode and long-lived 3D quantum systems.

7.1 Improvements to multimode architecture

7.1.1 Reducing off-resonant effects

In Ch.4 we discussed two challenges associated with scaling multimode quantum memories. One was the reduction of accumulated idling errors while doing operations on the register of modes by increasing the mode coherences (for a given qubit coherence), which was addressed in the intervening chapters. The second was dealing with the off-resonant effects caused via the simultaneous dispersive coupling of the qubit to all of the modes in the register. As the size of the register becomes larger, than the state-dependent shift on the qubit, and subsequently on the other modes of the cavity, become more numerous. This is made particularly worst if the Fock-state being prepared itself contain multiple photons. While in theory one can calibrate for this, the number of kerr and cross-Kerr terms to keep track of

scale quadratically with the number of modes.

There are a few options to combat this issue. One idea is to implement protocol that dynamically decouples the qubit from the non-target modes. Dynamical decoupling (DD) protocol have been popular in spin-systems for decoupling the quantum spin from a bath of lossy degrees of freedom, suppressing decoherence [180, 179], with more advanced DD protocol allowing for universal control and reduction of operational errors [181, 99]. Instead of decoupling from a both of spins, however, the qubit would need to selectively decouple from non-target modes, while preserving the robust state-preparation toolbox, and speed, that have made these systems so appealing. Efforts currently exist to study the viability of using these techniques, however the research is still in its early stages.

Another option to atleast alleviate the problem off unwanted off-resonant interaction is to simply reduce the coupling between the qubit and the modes, effectively reducing the shift. This only has limited utility, however, as the timescale over which gates can take place is set by the dispersive coupling rate, $T_g \sim 2\pi/\chi$ [186, 26]. By reducing this rate, one would lose a lot of the advantages that having such a high-cooperativity system may bring, and exacerbate both gate infidelity and idling errors by extending the time of operations without extending the coherence times of either the cavity modes or the qubit. One recently demonstrated solution to this problem involves the addition of a coherent cavity drive, which leads to a phase-state displacement [66]. In this new displaced frame, the dispersive dynamics are enhanced, with the displacement acting as a lever-arm, leading to large separations in phase space depending on the qubit state. The net result is an enhancement in $T_g \propto 1/|\alpha|$, or the sixe of the coherent state imparted onto the resonator. This is particularly exciting in the context of ultra-high coherence systems, such as our Nb cavities, where weak dispersive coupling is required to reduce reverse-Purcell decay of the storage mode. As of yet, no multimode implementation has been demonstrated, but is actively being pursued by our team in the existing multimode (MM2) system.

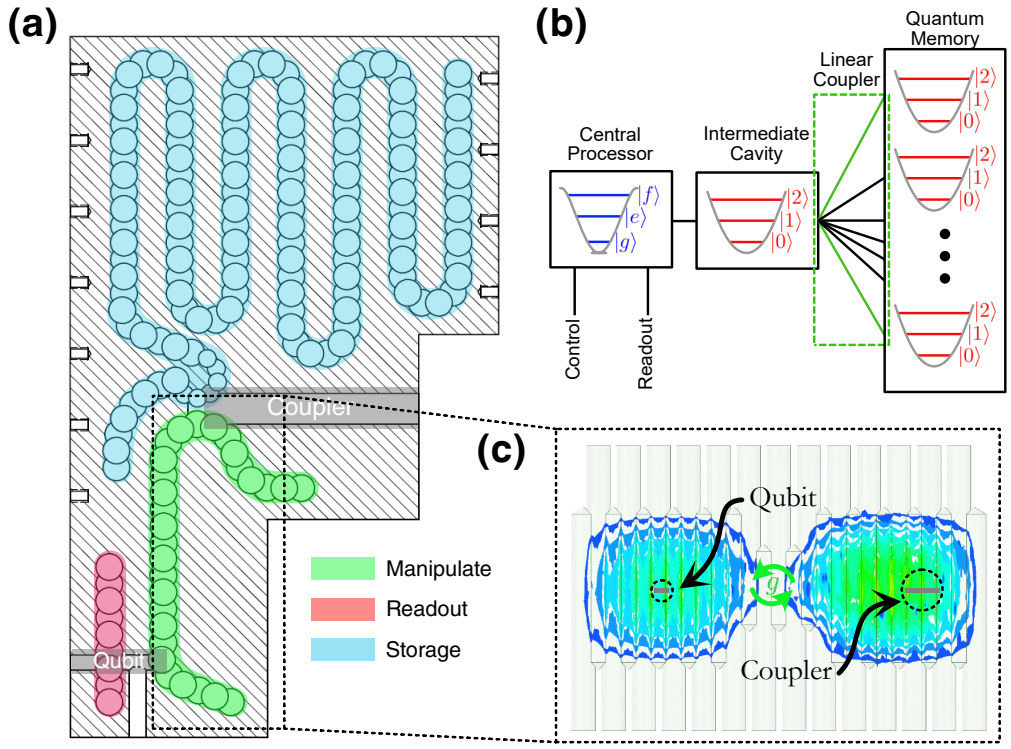

Figure 7.1: **Manipulate based multimodal memory** : (a) shows the realized multimodal architecture, with the addition of an intermediate "manipulate" cavity, which acts as a high Q intermediate to limit the cross-talk amongst storage modes. (b) shows a more generalized schematic of the system's interactions, while (c) shows a detail of the split manipulate cavity geometry, with the delocalized manipulate mode split to be strongly coupled to the coupler and qubit simultaneously.

7.1.2 Linear coupling and next-gen MM architectures

The previous section proposed techniques for reducing the effects of off-resonant interactions in the existing multimode architecture, however these techniques don't *completely* remove the underlying problem. Only a complete decoupling of the non-linearity from the modes would alleviate this issue, however doing so would remove the ability to have universal control over the modes of our quantum memory. One solution, which is currently being developed, is to have the qubit interact with an intermediate high-coherence cavity with one or two modes. To transfer a prepared state into a storage mode, a SQUID based coupler, can be parametrically driven to swap the state into the target mode.

Fig. 7.1 (a) shows the modified multimode memory with the intermediate manipulate

cavity. State preparation in the manipulate cavity is done via traditional universal control techniques, as outlined in [38]. To swap information from the intermediate cavity to the storage mode, a linear coupler is employed.

Recently, couplers that utilize a 3-wave mixing element, or SNAIL, have been developed, removing the 4^{th}-order non-linearity and subsequently cross-Kerr interactions all together [69]. Another option is the use of SQUID based couplers, which have been used in a number of cQED where tunable interactions are desired, including qubit-qubit interactions [118, 40], qubit-resonator [2, 170] and resonator-resonator [193, 190]. It was demonstrated in Ref. [193] that a SQUID loop driven via an RF-flux with detuning of $\omega_i - \omega_j$, where i and j are the adjacent cavity's mode frequencies, a beam-splitter swapping operation is induced between modes. By reducing the anharmonicity of the squid-loop, such that $E_j \gg E_c$, the anharmonicity can be effectively minimized to a point where cross-Kerr terms become negligeable.

Flux-modulation of the DC-SQUID loop requires the introduction of wires that must have the required bandwidth to flux-pump the coupler while also reducing undesired loss at the storage mode frequencies. Recent work, shown in Fig. 7.1 (b), has been done to design a 3D linear coupler which utilizes a 5-pole bandpass filter, made of alternating capacitive and inductive sections of stripline, to act as a Purcell filter, with $\sim 100dB$ of suppression of frequencies in-band with the desired storage modes. This filter is, in essence, analogous to stepped-impedance Purcell filters used to protect qubits from decoherence when dispersively coupled to a strongly coupled readout line [32].

One last challenge addressed in this unique geometry is the ability to strongly couple the transmon qubit and linear coupler to the manipulate cavity while reducing spurrious cross-coupling. Using the flute technique, one can create a split TE_{101} cavity made of two identical rectangular waveguide sections. The aperture separating the two is evanescent to the cavity modes, however, making the section short allows for a controlled coupling across the qubit

and coupler halves of the cavity. The net result is a two-mode cavity with modes split by some g_{man}. This style of coupling is analogous to the chain of CPW waveguide resonators used in creating the 2D multimode in Ref. [124], where the capacitive coupling has been replaced by a waveguide section. The dynamics, however, still follow a tight-binding mode dispersion as described in Ref. [119]. By setting the splitting by $g_{man} \sim 200$MHz, it ensures that the effective tunneling rate is higher than the effective swap time of the design qubit and coupler with their respective readout and storage modes.

7.2 Improved Nb cavity based memories

In CH.6 we briefly discussed several immediate avenues for improving Nb based quantum memories, and experiments to better characterize the loss mechanisms that can be effecting the cavity T^1. In addition to this there are aspects of the Nb cavity manufacturing process that can lead to further improvements in cavity performance. In addition to this, the techniques outlined in this thesis may also be useful in improving 2D niobium resonators. This last line of inquiry is currently being undertaken by F. Zhao in collaboration with our lab, and will hopefully lead to a tickle-down effect in the performance of 2D systems. The following sections will explore these directions more thoroughly.

7.2.1 Cavity annealing and Nitrogen doping

The optimization discussed in Ch.5 was primarily concerned with improving the initial etch process. It was shown that the limiting factor in the performance of the cavity was two-level systems inside of the surface oxide. By minimizing the time exposed to air, we were able to minimize the effect. Recent work has shown that vacuum annealing the cavities at ~ 450°C can diffuse the oxide layer into the surface, minimizing loss [156]. In addition to this, the addition of nitrogen into the surface, via N_2 backfilling of the vacuum chamber at T$\geq$ 750°C, has also shown a $\geq 3\times$ decrease in the residual resistivity [82, 57, 58, 56].

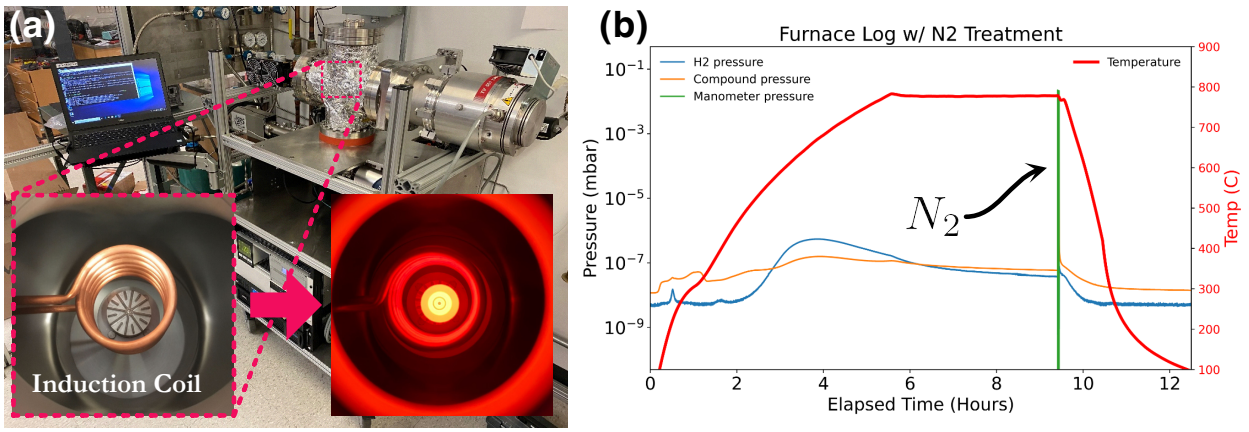

Figure 7.2: **Vacuum annealing preliminary results**: (a) shows a picture of the furnace exterior, with the large turbo-molecular pump located at right, with an inset showing the interior water-cooled induction coil at room temperature and with a cavity under heating at $\sim 750°$C. (b) shows a plot of a typical annealing run in the furnace, with a peak in H_2 partial pressure at 450°, and decreasing by $40-80\times$ over the course of 6 hours. The green line is a 6 min long N_2 pulse for nitrogen surface doping. The furnace was constructed by the author.

This style of vacuum annealing system is also used to remove interstitial hydrogen from the surface via high-temperature diffusion [19].

Fig. 7.2 (a) shows a picture of a recently built custom vacuum furnace that, in theory, meets the specifications for conducting the above surface treatments. Based on the work in Ref. [120]. The furnace uses a 75kHz fixed frequency induction coil to heat the cavity. A residual gas-analyzer (quadrapole mass spectrometer), gives real-time information on the composition of outgassing species. For nitrogen doping, the furnace also has a ultra-high purity N2 gas line hooked up through an impedance and valve for precise control of N2 back-filling. The entire furnace is controlled via a programmable logic controller and automation is initiated via a computer interface. The furnace also has automated, PID controlled bake-out heaters (8 channels in total). A $\sim$ 700L/s @ $< 1\times10^{-4}$ mBar turbo-molecular pump allows for sustained pressures of $\leq 1\times10^{-6}$mBar at temperatures in excess of 750°C.

Preliminary results, shown in Fig. 7.2 (b), show a $> 40\times$ reduction in the hydrogen partial pressure, showing the efficacy of the system in removing trapped interstitial hydrogen

at the surface, however measured cavities have shown a marked *decrease* in performance over their un-annealed counterparts. One theory is that, because the hot-zone is confined to the cavity and does not include the surrounding furnace, radiant heat from the cavity leads to the outgassing of contaminants at the vacuum chamber walls, which can recondense on the hot surface of the cavity. Ref. [81] showed the sensitivity of cavities to these types of contaminants. The introduction of a susceptor that surrounds the cavity [59], and the use of a getter, such as titanium, may act to reduce the effect of outgassing contaminants. Further analysis of post-treated cavities, using the techniques outlined in Sec. 5.4, would be necessary to make any conclusions, however. In addition to this, very little research has been done with BCP etched cavities undergoing N2 treatment, with previous results using electropolishing techniques.With the advent of these techniques, however, it is perhaps not unreasonable to believe that single-photon cavity coherence could be further improved by a factor of $3-9\times$, leading to lifetimes of $>$ 100ms for the $\sim$ 6.5GHz coaxial cavity discussed in this thesis.

7.2.2 Improving of 2D superconducting materials

The work in this thesis may also provide insights into making 2D niobium resonators better. Recent results with tantalum showed marked improvements in qubit coherence times, which have been demonstrated to be as high as 0.5ms [140, 188]. Tantalum, like niobium, is a group V transition metal, and such has very similar surface chemistry and structural properties, along with a high superconducting transition temperature ($T_c \sim 4.38$ K). When compared to niobium thin films, the properties appear better, however a longitudinal study that compares thin films using similar growth techniques, etching, and intrinsic qualities such as residual resistivity ratio values (RRR), has not been done. For example, recent results with Nb films grown to 500nm using a nearly identical high-temperature sputtering process to that used in Ref. [140] gave RRR values comparable to that produced in the high-performance Ta devices (RRR$>$ 40), and record coherences for 2D microwave resonators of $\sim 40 \times 10^6$

at 3.9GHz [129]. It was found that the thicker film, sputtered at high temperature, has much higher crystallinitity. Also, with a thickness of 500nm, which is much longer than λ_L, even in a dirty limit, the base material begins to take on bulk like characteristics. This is in contrast with the evaporated devices previously used, which utilize a 75nm film thickness and RIE etching, which produces a thin film with typical RRR$<$ 10. Also, given the similarity of the wet-etch chemistry to that used in 2D Ta device fabrication (Ta etch uses a 1:1:1 HF:HNO_3:H_2O mixture) as that used in the 3D Nb cavities, there is also some question if the recipe developed can be extended to 2D devices as well. This, combined with better information as to the effect of sapphire substrate orientation, and loss tangent [146], could produce 2D resonators with quality factors approaching 1×10^8.

CHAPTER 8

CONCLUSIONS

This thesis has been a detailing of the natural evolution of 3D cQED systems for use as quantum memories. The development of multimodal cavities with record breaking cavity-qubit cooperativities, and processes for reliably making seamless TE_{101} cavities and niobium coaxial cavities will give the field as a whole more options for platforms to build quantum hardware or investigate interesting quantum phenomena. In that sense this thesis is a small stepping stone in a larger arch of quantum research. While it has not elucidated any fundamentally new quantum phenomena, it will hopefully be the motivation for future researches to explore regimes that did not exist before the creation of these techniques.

By analogy, this work is much like the "tock" in the Intel tick-tock model of development. The work outlined by Matt Reagor, Eric Holland, and Theresa Brecht in 2013-2018 represented the first major move forward in coherence times in 3D systems. Since then, qubit fabrication, and protocol for controlling these systems have improved dramatically, an optimization stage much like the "tick" cycle of CPU microarchitecture. But now, these platforms are feeling dated, and the whims of theoreticians are beginning to butt up against the reality of these systems. Now, with a factor of 10 improvement in coherence, we need to rethink and re-investigate a lot of the processes and methods we take for granted. In that sense we are entering a new cycle of development, with active research in niobium based cavity systems now being undertaken by national and university research labs becoming more prevalent. The trickle-down effect of these innovations will also work synergetically with ongoing research in improved materials for qubit and 2D resonator fabrication, leading to an uplift of the entire field. With this work, and other recent investigations into qubit improvements, we will also see more mindfulness by quantum physicists into the underlying material properties, bringing in a wider range of material scientists, physical chemists, and condensed matter physicists into the quantum field.

Stepping back and putting all of this into context, we are seeing the trajectory of this field from nascent fundamental physics research, to a technical discipline in its own right. The level of sophistication that has been introduced, at all levels, is requiring physicists to think in more applied ways than ever before. My only hope is that, as we move forward, the field embraces these more technical aspects of the job, and outlines standards and practices that turn a small arcane field of study into a legitimate enterprise. We have seen in this thesis that better accounting of the cryogenic environment as a whole, with some nuts-and-bolts cryogenics engineering, can have drastic effects on the performance of these quantum systems. This is not to mention my *other* work in developing new kinds of cryogenic platforms, which is a story for another time, and has only increased my appreciation for perhaps the less glamorous aspects of our job. Needless to say, as these platforms also scale, an adoption of these ideas, and stepping outside of the traditional rolls we play as physicists, will become necessary to progress as a whole.

APPENDIX A

NUMERICAL SOLVING OF CPB HAMILTONIAN IN PYTHON

The following is adapted from the amazing Appendix C. example listed in Ref. [164] C, which was originaly written in Mathematica, aka the lesser programming "language". For posterity and thoroughness, I thus provide my example below.

To begin we must create cosine and ϕ matricies:

```
import numpy as np
import scipy
import scipy.linalg as la
import matplotlib
from matplotlib import pyplot as plt
from scipy import constants
from scipy.optimize import brentq
import math

def cosmat(phi):
    n=len(phi)
    cos_vals=np.cos(phi)
    return np.diag(cos_vals, 0)

def ThetaMat(phi):
    #create diagonal theta matrix using nd.diag
    return np.diag(phi,0)
```

Then we can create the finite difference matrices with the requisite boundary conditions:

With the helper functions out of the way, we can construct the hamitonian and solve it:

Now that the functions are setup, we can run the code to solve the CPB hamiltonian for various offset charge states, then plotting as a function of n_g

```
def periodic_ddn(phi):
  dims=len(phi)
  ones=np.ones(dims)
  delta=(max(phi)-min(phi))/dims
  #using 4th order central difference with periodic BCs
  dndn=-np.diag(ones[0:dims-2], -2)+16*np.diag(ones[0:dims-1],
  ↪  -1)+16*np.diag(ones[0:dims-1],1)-30*np.diag(ones,0)-np.diag(ones[0:dims-2], 2)
  #set BCs
  dndn[0,-1]=16.
  dndn[-1,0]=16.
  dndn[0,-2]=-1.
  dndn[-2,0]=-1.
  dndn[1,-1]=-1.
  dndn[-1,1]=-1.
  return (1./(12*delta**2))*(dndn)

def periodic_dn(phi):
  dims=len(phi)
  ones=np.ones(dims)
  delta=(max(phi)-min(phi))/dims
  #using 4th order central difference with periodic BCs
  dn=-np.diag(ones[0:dims-2], 2)+8*np.diag(ones[0:dims-1],
  ↪  1)-8*np.diag(ones[0:dims-1],-1)+np.diag(ones[0:dims-2],-2)
  #set BCs
  dn[0,-1]=-8.
  dn[-1,0]=8.
  dn[0,-2]=1.
  dn[-2,0]=-1.
  dn[1,-1]=1.
  dn[-1,1]=-1.
  return (1./(12*delta))*dn
```

```
def offset_CPB(Ec, Ej, n, n_offset):
  #set size for hamiltonian
  dims=201
  #create phase space array to evaluate energies over
  phi=np.linspace(-np.pi,np.pi,dims)
  #Call our second order periodic finite difference matrix scaled appropriately wrt
  ↪  phi array
  dndn=periodic_ddn(phi)
  dn=periodic_dn(phi)
  #calculate a diagonal matrix for cos evaluated over phase space
  cosvals=cosmat(phi)
  #create identity
  I=np.diag(np.ones(len(phi)))
  #assemble our uncoupled CPB hamiltonian
  ham_cpb=-Ej*cosvals+4.*Ec*(-dndn-1j*n_offset*dn+I*n_offset**2/4)
  #calculate unsorted eigen energies
  US_eigen_vals=la.eig(ham_cpb)[0]
  #print the first n eigenvalues after offsetting lowest energy to zero and sorting in
  ↪  ascending order
  Eigen_Vals=np.real(sorted(US_eigen_vals-min(US_eigen_vals)))
  return Eigen_Vals[0:n]
```

```
offset_e=np.linspace(-2, 2, 200)
eig_vals=[]
ratio=4
Ec=0.175
Ej=ratio*Ec

fig, ax = plt.subplots(1, figsize=(4, 4))

for n_g in offset_e:
    eig_vals.append(offset_CPB(Ec, Ej, n=4, n_offset=n_g))

eigs=np.array(eig_vals).T[1::]
E_01=eigs[0]
E_12=eigs[1]
E_23=eigs[2]

plt.plot(offset_e, (E_23), label='$E_{2}$')
plt.plot(offset_e, (E_12), label='$E_{1}$')
plt.plot(offset_e, (E_01), label='$E_{0}$')

plt.xlabel('$n_{g}$ ($e^{-}$)', fontsize=20)
plt.ylabel('Energy (GHz)', fontsize=20)
plt.title('$E_{J}/E_{C}=$%i'%ratio, fontsize=20)
ax.tick_params(axis='x', labelsize=15)
ax.tick_params(axis='y', labelsize=15)

plt.legend( fontsize=15, frameon=False, bbox_to_anchor=(1, 1))
```

APPENDIX B

SIMULATING VIBRATIONAL MODES OF CANTILEVER

The below is modified from

https://github.com/apf99/Finite-Element-Modeling/blob/master/fem_beam.py

originally written by Andrew Friedman for solving the 1D Euler-Bernoulli beam equation:

```
import numpy as np
from scipy.linalg import eigh
import math
from matplotlib import pyplot as plt
import time

def beam(num_elems):
        restrained_dofs = [1, 0, -2, -1]

        l = .0165 / num_elems
        Cm = (w)*(t)*3.98E-3*1E6   # rho.A
        Ck = 400E9*1.622E-14   # E.I

        # element mass and stiffness matrices
        m = np.array([[156, 22*l, 54, -13*l],
                                [22*l, 4*l*l, 13*l, -3*l*l],
                                [54, 13*l, 156, -22*l],
                                [-13*l, -3*l*l, -22*l, 4*l*l]]) * Cm * l / 420

        k = np.array([[12, 6*l, -12, 6*l],
                                [6*l, 4*l*l, -6*l, 2*l*l],
                                [-12, -6*l, 12, -6*l],
                                [6*l, 2*l*l, -6*l, 4*l*l]]) * Ck / l**3

        # construct global mass and stiffness matrices
        M = np.zeros((2*num_elems+2,2*num_elems+2))
        K = np.zeros((2*num_elems+2,2*num_elems+2))

        # for each element, change to global coordinates
        for i in range(num_elems):
                M_temp = np.zeros((2*num_elems+2,2*num_elems+2))
                K_temp = np.zeros((2*num_elems+2,2*num_elems+2))
                M_temp[2*i:2*i+4, 2*i:2*i+4] = m
                K_temp[2*i:2*i+4, 2*i:2*i+4] = k
                M += M_temp
                K += K_temp

        # remove the fixed degrees of freedom
        for dof in restrained_dofs:
                for i in [0,1]:
                        M = np.delete(M, dof, axis=i)
                        K = np.delete(K, dof, axis=i)

        evals, evecs = eigh(K,M)
        frequencies = np.sqrt(evals)
        return M, K, frequencies, evecs
```

```
    errors = []
    prev_frequency=0
    for i in range(2,20):     # number of elements
            start = time.time()
            M, K, frequencies, evecs = beam(i)
            time_taken = time.time() - start
            error = (frequencies[0] - prev_frequency) / prev_frequency * 100.0
            prev_frequency=frequencies[0]
            errors.append( (i, error) )
            print ('Num Elems: {} \tFrequency: {}\tError: {}% \tShape: {} \tTime:
            ↪  {}'.format( i, round(frequencies[0],3), round(error, 3), K.shape,
            ↪  round(time_taken*1000, 3) ))

        print('Exact Freq:', round(exact_frequency, 3))

        element  = np.array([x[0] for x in errors])
        error    = np.array([x[1] for x in errors])


        # plot the result
        plt.plot(element, error, 'o-')
        plt.xlim(1, element[-1])
        plt.xlabel('Number of Elements')
        plt.ylabel('Error (%)')
        plt.show()
```

APPENDIX C

SNR ESTIMATION

```
'''
Pertinent functions. Noise power is the primary Friis calculation function,
Power_n is a function that calculates power as a function of photon number,
SNR is a wrapper function that takes in the params object to calculate the
SNR, returning arrays for IFBW in Hz, noise and signal power in dB.
'''
def Noise_Power(params):
    meas_chain=params['meas_chain']
    IFBW=params['IFBW']
    T_dev=params['T_dev']
    PN_NWA=10**((params['PN_NWA']-30.)/10)
    F=np.zeros(len(meas_chain))
    G=np.zeros(len(meas_chain))
    T0=290 #kelvin
    kb=1.38E-23
    amp_list=[]
    for I, dev in enumerate(meas_chain):
        if dev['type']=='atten':
            L=10**(dev['loss']/10.)
            F[I]=1+((L-1)*dev['T_atten'])/T0
            G[I]=L**-1
        elif dev['type']=='amp':
            F[I]=10**(dev['NF']/10.)
            G[I]=10**(dev['gain']/10.)
            amp_list.append(dev['name'])
        else:
            print('ERROR: Not valid device type:'+str(dev['type']))
            break
            return
    F_tot=F[0]
    for i in range(1,len(G)):
        F_tot+=(F[i]-1)/np.prod(G[0:i-1])
    PN_tot=30+10*np.log10(kb*T_dev*np.prod(G)*IFBW*F_tot+PN_NWA)
    return {'PN':PN_tot, 'Gain':G, 'Amps':amp_list}
```

```
def Power_n(params):
    k_i=2*np.pi*params['k_int']
    kc_1=2*np.pi*params['kc_1']
    kc_2=2*np.pi*params['kc_2']
    if params['meas_port']==1:
        k_meas=kc_1
    elif params['meas_port']==2:
        k_meas=kc_2
    else:
        print('ERROR: Invalid port')
        return
    f0=params['f0']
    n_bar=params['n_photon']
    h=scipy.constants.h
    #equation 2.53 from Audrey Bienfait thesis (pg.31)
    power=(n_bar*h*f0*(kc_1+kc_2+k_i)**2)/(4.*k_meas)
    print(30+10*np.log10(power))
    print(power)
    return power
```

```
def SNR(params):
    params_noise={'meas_chain':params['meas_chain'],
                  'IFBW':params['IFBW'],
                  'PN_NWA':params['PN_NWA'],
                  'T_dev':params['T_dev']
                 }
    params_signal={'k_int':params['k_int'],
                   'kc_1':params['kc_1'],
                   'kc_2':params['kc_2'],
                   'meas_port':params['meas_port'],
                   'n_photon':params['n_photon'],
                   'f0':params['f0']
                 }
    noise_output=Noise_Power(params_noise)
    G=noise_output['Gain']
    IFBW=params['IFBW']
    P_n=Power_n(params_signal)
    #print(30+10*np.log10(P_n))
    P_signal=30+10*np.log10(np.prod(G)*P_n)
    P_noise=noise_output['PN']
    if params['plot']==True:
        plt.plot(IFBW, P_signal-P_noise, label="=>".join(noise_output['Amps']))
        plt.xscale('log')
        plt.xlabel("IFBW (Hz)")
        plt.ylabel("SNR (dB)")
        plt.xlim(np.min(IFBW), np.max(IFBW))
        plt.legend(loc=0, prop={"size" : 10}, frameon=False)
        plt.title("n=%s photon SNR"%params['n_photon'])
        plt.grid(which='both')
    return IFBW, P_signal, P_noise
```

APPENDIX D

GAS-GAP HEAT-SWITCHES

D.1 Heat-switch operation and overview

You may be wondering *why* there is a section on the arcane inner workings of gas-gap heat switches—a rather pedestrian cryogenic component—in this thesis on quantum memories. The short answer is because I simply can, but the longer answer is because, before and during the research described above, I was also working as a cryogenics engineer, developing a novel kind of continuous adsorption refrigerator capable of driving a dilution unit. The full description of this device, lovingly called the THeIA refrigerator by myself, could *easily* have been another thesis. In fact, it *was* supposed to be my thesis, but was considered by many to not have enough "newness" to the problem. Being that I did the work to build and characterize it (and to be clear, it does work) I decided to write another entire half to this thesis going over this refrigerator, and the intricacies of its design and operation. Needless to say, I bit off more than I could chew, and givem the sheer amount of work required to write one thesis, I was unable to finish this second half. Instead, I provide you the one chapter I *did* write as an appendix, with the hope that the reader, whoever they may be and with whatever background they may have, may learn something new.

D.1.1 Heat-switch types and background

In Ch 11 the idea of a heat switch, or a device with variable thermal conductivity, was used to couple the pump to the 4K plate during pumping, or thermally decouple the pump during desorption as to not perturb the plate temperature. In this chapter I will discuss the working principles underlying the design, construction, and operation of the gas-gap heat switch. The gas-gap heat switch is just one class of variable thermal conductance devices that are used to control the flow of heat at various points in a cryogenic cycle. The choice

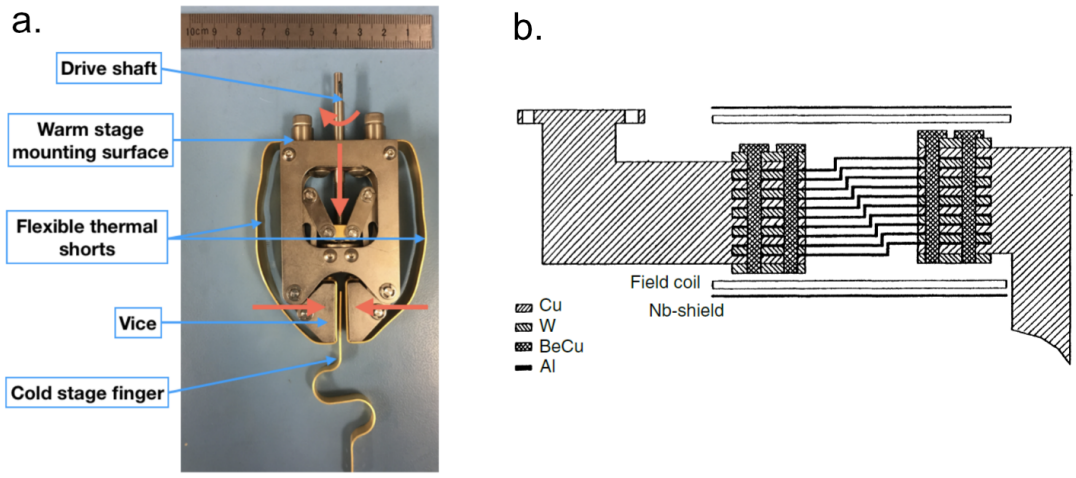

Figure D.1: **Mechanical and superconducting heat switchs**: (a) An HPD Instruments© (now FormFactor©) mechanical heat switch. In ADR applications a high temperature heat switch would connect the paramagnetic salt pill to a thermal sink during the application of a constant magnetic field, lowering the spin entropy of the system. This switch is then opened as the gradient is reduced, while a second switch, attached to the cold-stage finger, is closed, cooling the system. Image taken from [195]. (b) A schematic of a superconducting aluminum heat-switch used to connect a nuclear demagnetization stage to the mixing chamber of the Bayreuth refrigerator [80]. Parallel strips of ultra-high purity aluminum are used due to aluminum's low critical field (10.5mT), high Debye temperature ($\Theta_D \approx 400K$), and availability in high purity form. Given its strong native oxide extra care is taken to reduce the contact resistance of the aluminum to improve on-state conductivity.

of heat switch type is largely determined by the amount of on-state conductance, the heat switch "aspect ratio", or the on-state vs off-state ($\kappa_{\text{on}}/\kappa_{\text{off}}$), and the specific size constraints of the system.

While this section will not be an exhaustive overview of the various types, to provide context for the choice of using gas-gap heat switches for both the Breton and Whistler fridges, I contrast the gas-gap heat switch to other common heat-switch designs. The most intuitive of these designs is the mechanical heat switch Fig. D.1 (a), which uses mechanical pressure to make contact between two high-conductivity pieces of metal. This design has the advantage of very high on-state conductance and an incredibly high aspect ratio. A drawback to this design however, is the large contact force, with the contact conductance scaling $\kappa_c \propto F^{1/3}$ [21, 47, 46]. This means for typical applications, such as cooling a

paramagnetic salt pill for an adiabatic demagnetization refrigerator, forces in excess of 100N are required to achieve conductivities approaching 1W/K. While screws are often employed to increase mechanical advantage, large room temperature motors are still required to make sufficient contact, necessitating room-temperature feedthroughs and adding to the overall complexity of the system.

A mechanically simpler heat switch design is the superconducting heat switch, as shown in Fig. D.1 (b), which takes advantage of the poor electron conductivity of a superconductor when $T \ll T_c$ to limit it's off-state conductivity, while achieving high thermal conductivity in the presence of a magnetic field due to high electron conductivity when the metal is in a normal state. In the limit where $T \ll T_c$ the switching ratio scales as $\kappa_{on}/\kappa_{off} \approx \alpha T^{-2}$, with the proportionality constant $\alpha > 10^2$ for a well optimized switch. As $T \geq 0.1T_c$ then $\kappa_{on}/\kappa_{off} \propto \exp\{(\Delta E(T)/k_b T)\}$, where $\Delta E(T)$ is the superconducting gap of the heat switch metal used in the switch. This means for materials like aluminum, which exhibit good on-state thermal conductivity, the operating temperature is $\leq$ 100mK [141, 79]. In theory materials with higher T_c can be used, however materials like Nb and its alloys have comparatively low normal-state conductivity, and still only operate well below 1K. This makes them unusable for applications at or above 4K, as is the case for adsorption pump applications. This aspect relegates them to sub-kelvin applications such as cooling nuclear demagnetization stages for conducting sub-mK research.

In addition to these there are other, more exotic, heat switches have also been made, which utilize the magnetoresistive effect. These switches however are limited to metals with closed Fermi surfaces and require the use of single crystals in order to ensure a high on/off aspect ratio. This, combined with the need of magnetic fields in excess of 1T, and difficulty machining examples that have high l/w ratios, relegate them to niche applications [36]. There have also been attempts to utilize thermal expansion to make "passive" mechanical heat-switches, however these have been relegated to higher temperature (>80K) applications

where the thermal expansion coefficient of many polymeric materials, such as UHMW-PE, are large enough to achieve the high aspect conductivity ratios necessary while also being easy to manufacture [60, 165].

D.1.2 Gas-gap heat-switch overview

In contrast to these designs, gas-gap heat-switches require no moving parts, and can operate with relatively high-performance, of >100mW/K, at temperatures as low as 3K. This is done by mediating heat transfer through a low-pressure working fluid that spans across two pieces of high thermal conductivity material. Fig. D.2 shows the two states of a gas-gap heat switch. In the off-state (Fig. D.2 (a)) a small piece of adsorbant material is passively cooled via a weak thermal link or radiative cooling, adsorbing the helium and leading to a low equilibrium pressure inside of the gap. At this point the mean-free path of helium atoms is far longer than the gap (free-molecular regime), and heat transfer is dominated by conduction through the outer support shell (grey) or radiation between the hot and cold end. As the adsorber temperature increases (shown in Fig. D.2 (b)) the isosteric equilibrium helium pressure increases. As the mean-free path of helium atoms approaches the gap length (transition regime) the thermal conductivity is pressure dependent, before reaching a maximum as the helium pressure increases to a point where the mean-free path is much shorter than the gap (continuum regime) and thermal conductivity becomes pressure independent.

In the next few sections we will explore the kinetic theory underpinning this operation in more detail. We will see that, because pressure (and subsequently thermal conductivity) are mediated by the isosteric adsorption properties of our adsorber, we can have a high degree of control over the conductivity of our switch, effectively turning it into a "heat valve" instead. From a control perspective this is highly advantageous for dynamical systems such as continuous adsorption refrigerators. This, combined with their small size, high-reliability and performance make them ideally suited for our application.

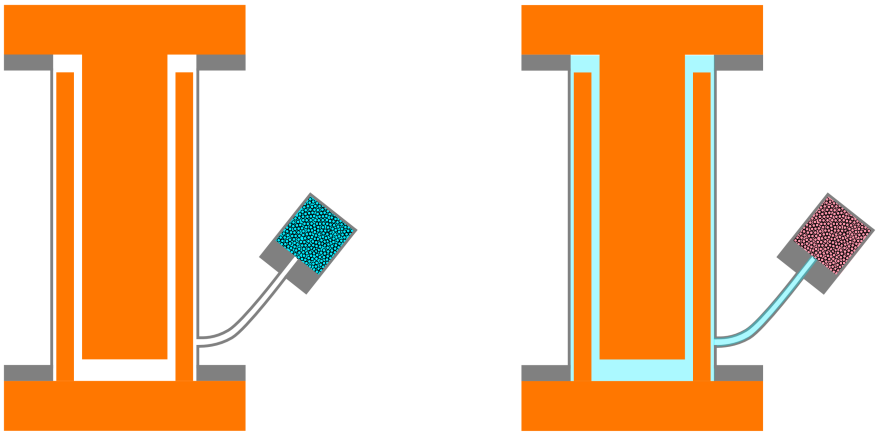

Figure D.2: **Gas-gap heat-switch operation**: The two states of the gas-gap heat switch are mediated by the presence of helium in the gap. An "off" heat switch has the adsorber cooled, and no gas in the gap, leading to thermal conduction only through radiative means, or parasitic conduction through the heat-switch's outer supporting shell.

In this chapter we will explore the technical challenges in building a high-performance gas-gap heat switch, and show how its method of operation makes it ideal for continuous adsorption refrigeration. To do this however will require a more thorough treatment of the switching dynamics than has been previously performed in the existing literature, along with characterization and calibration of these dynamics to take full advantage of them.

D.2 Theory of operation

D.2.1 Knudsen number, free-molecular and continuum heat transfer

In cryogenic systems the transfer of heat is often mediated by a working fluid. In liquids phonon conductivity determines the thermal conductivity of the fluid, however in gases the method in which heat can be transferred depends largely on the density of the fluid. At higher densities energy is transferred through elastic collisions between individual atoms and molecules, which, given a large number of atoms, leads to the gas acting like a homogeneous mass. At low densities and pressures, where the scale of collisions for the gas atoms are longer than the characteristic length of the container, energy is transferred by atoms colliding with

the walls of the container. When the atom collides, it does not achieve thermal equilibrium with the wall, leading to only partial energy transfer between the wall and atom. In the case of a helium atom interacting with a hot surface, it will gain a small amount of energy which it will carry to a lower temperature surface, where it will deposit some of that excess energy, repeating the process. As one can see, the net transfer of energy in this low-density, or "free-molecular" regime, is directly proportional to the number of atoms and subsequently the pressure of our system.

A useful way of determining the flow regime of a given system is the dimensionless Knudsen number, which is simply the ratio of the mean-free path and the characteristic length of the system:

$$K_n = \lambda / L_e \tag{D.1}$$

Where L_e is the characteristic length of our system, which is generally written as:

$$L_e = 4V / A_w \tag{D.2}$$

V is the enclosed volume and A_w is the enclosed area. The gas mean-free path is derived from kinetic theory and is a function of the gas molecule's mean velocity, viscosity, and pressure:

$$\lambda = \frac{\mu}{p}\left(\frac{\pi RT}{2M_g}\right)^{1/2} \tag{D.3}$$

Where μ is the gas viscosity and M_g is the molecular mass of the gas species. R is the universal gas-constant, T is the gas temperature, and p is the gas pressure.

From the above we can determine the Knudsen number for any system. This quantity gives us a useful way of describing what flow regime the system is in, with the division of regimes being given as:

$$\text{Continuum: } K_n < 0.01 \tag{D.4}$$

$$\text{Transition: } 0.01 < K_n < 0.3 \tag{D.5}$$

$$\text{Free-Molecular: } K_n > 0.3 \tag{D.6}$$

In the continuum regime the total power throughput of the switch is independent of the pressure, and can be treated as a simple 1D heat transfer problem:

$$\dot{Q}_{cont} = F_\alpha \frac{A}{l} \int_{T_l}^{T_h} \kappa(T) dT \tag{D.7}$$

Where F_α is the dimensionless accommodation factor, or the degree in which the gas molecule reaches thermal equilibrium with the wall. A is the effective heat transfer area and l is in this case the thickness of the sheet of gas, which is set by the gap length. κ is the thermal conductivity of the gas. The above integral represents the *maximum* effective conductivity of the heat switch, assuming that the switch is operating in a regime where radiation, convection, and parasitic thermal conduction is negligible. Equation D.7 also gives is the two main parameters for maximizing conductivity, the width of the gap and the area of the two halves of the heat switch.

In the free-molecular and transition regimes the thermal conductivity is pressure dependent. From Barron [21] eq. 5-63 we can get the mass flow rate of gas molecules interacting with the surface from the average velocity and density:

$$m/A = \frac{\rho v_a}{4} = \frac{P}{(2\pi RT)^{1/2}} \tag{D.8}$$

To determine the total energy flux along the surfaces of the heat switch we need to determine the exchange in energy per unit mass of the gas as it interacts with the surfaces. To determine this we need to know the change in both internal energy and kinetic energy

per unit mass of the molecules striking the surfaces. This can be written as:

$$\Delta E = c_v \Delta T_{gas} + 1/2 R \Delta T_{gas} \tag{D.9}$$

We can simplify the above into a more convenient form by making the transformation:

$$c_v = \frac{R}{\gamma - 1} \tag{D.10}$$

Where γ is the ratio of heat capacity under constant pressure and constant volume (c_p/c_v). This transforms D.9 into:

$$\Delta E = \frac{R(\gamma + 1)}{(\gamma - 1)} \Delta T_{gas} \tag{D.11}$$

While we know the temperature of the heat-switch's hot and cold surfaces T_h and T_c, to determine the exchange in energy we need to know the gas molecule's temperature after interacting with the two surfaces to then determine ΔT_{gas}. A molecule starting at the cold surface at temperature T_c will come into partial thermal equilibrium with the surface and leave at some new temperature $T_c{}'$, before interacting with the hot surface at temperature T_h, where it will exchange some energy before leaving at a new temperature $T_h{}'$. To make this more explicit we can write the factors for both the hot and cold surfaces as:

$$\text{Cold Surface: } \alpha_c = \frac{T_h{}' - T_h{}'}{T_h{}' - T_h} \tag{D.12}$$

$$\text{Hot Surface: } \alpha_h = \frac{T_h{}' - T_c{}'}{T_h - T_c{}'} \tag{D.13}$$

From the above we can solve for the actual difference in the gas molecule temperature as a function of the wall temperatures and the two accommodation coefficients:

$$T_h - T_c = \left(\frac{1}{\alpha_c} + \frac{1}{\alpha_h} - 1\right)(T_h{}' - T_c{}') = \frac{T_h{}' - T_c{}'}{F_\alpha} \tag{D.14}$$

Where $T_h{}' - T_c{}'$ is ΔT_{gas} in equation D.9, F_α is the accommodation factor as shown above in eq. D.7, written here for a simple case of two parallel walls. F_α however is geometry dependant, and for the case of concentric cylinders takes on the form:

$$F_\alpha = \frac{1}{\alpha_c} + \frac{A_c}{A_h}\left(\frac{1}{\alpha_h} - 1\right) \tag{D.15}$$

This factor is analogous to the emissivity coefficient factor listed in table 5.4 of Barron [21]. With the above in hand, we can combine equations D.8, D.11, and D.14 to arrive at the power conducted through the switch at a given pressure P in the free molecular regime at the cold surface A_c:

$$\dot{Q}_{FM}(P,T) = F_\alpha P A_c \left(\frac{\gamma+1}{\gamma-1}\right)\left(\frac{R}{2\pi T}\right)^{1/2}(T_h - T_c) \tag{D.16}$$

In our heat switch this pressure p is a quantity that is dependent on the equilibrium pressure of our adsorbent at a given temperature. In the next subsection we will connect the isosteric adsorption properties of our heat-switch's adsorber to arrive at a temperature dependent conduction model for our switch.

D.2.2 Adsorption mediated switching

In equation D.16 we showed that the power through a heat switch in the free-molecular regime is linearly dependent on pressure, but in most heat switches the pressure of the gas inside the gap is dependent on the adsorber temperature. In section **??** we discussed the nature of isosteric adsorption. There we saw that, under isosteric conditions, or rather when the ratio of adsorbant to gas by mass is constant, the relationship between pressure and temperature follow the Clausius-Clapyron relationship [177]:

$$Q_{ads} = -R\left(\frac{d(ln(P))}{d(1/T)}\right)_n \tag{D.17}$$

Where Q_{ads} is the isosteric heat of adsorption. This form is analogous to to the coexistence curve of an ideal saturated vapor-liquid system taken that the latent heat is constant over the desired temperature range. Here we can ignore the change in specific volume between gaseous and adsorbed phases, in part because the specific volume of the gas phase $v_g >> v_{ads}$. This approximation holds well due to the difference in energy scales between the binding energy of helium atoms to the adsorbant versus to each other. With this approximation in hand we can rewrite it in the form;

$$P(T) = P_0 \exp(-Q_{ads}/RT_{ads}) \tag{D.18}$$

From the isosteric condition set in equation D.17, equation D.18, only holds for systems where both the mass of adsorbant and helium are fixed. In a gas-gap heat-switch this condition is met. Combining equations D.18 with D.16:

$$\dot{Q}_{FM}(T_{ads}, T_{gap}) = F_\alpha A_c P_0 \exp\left(\frac{-Q_{ads}}{RT_{ads}}\right)\left(\frac{\gamma+1}{\gamma-1}\right)\left(\frac{R}{2\pi T_{gap}}\right)^{1/2}(T_h - T_c) \tag{D.19}$$

In equation D.19 the distinction is made between the adsorber temperature (T_{ads}) and the temperature of the gas *in* the gap (T_{gap}). As we will see in section D.4.2 the gap temperature can be treated as roughly constant if the hot end temperature is held constant while changing the applying power.

D.2.3 Switching performance in the transition regime

In the previous section we created a complete temperature dependent model for gas-gap thermal conduction in the free-molecular regime, and by D.2.1 we showed what the expected

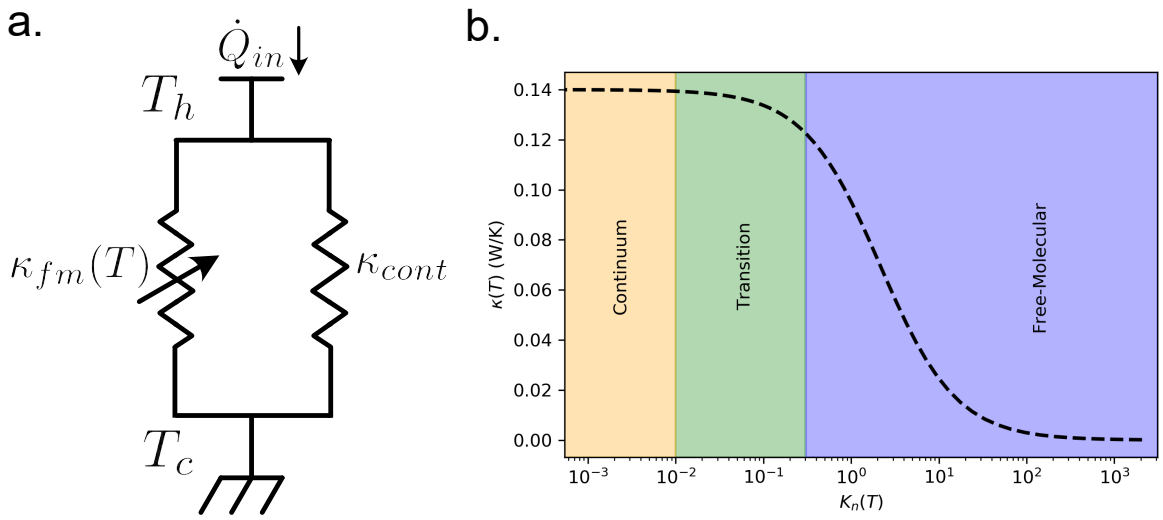

Figure D.3: **Model for gas-gap heat switch operation**: (a) shows the equivalent circuit-model for the flow heat from T_h to T_l given an applied power $\dot{Q}_{in}$. κ_{cont} represents the parasitic pathways for heat to travel from high to low temperature, while $\kappa_{fm}(T)$ represents the variable temperature conductivity of the free-molecular and transition regimes. (b) shows the expected response of thermal conduction for our theoretical heat switch for the temperature dependent Knudson number $K_d(T)$.

max thermal conductivity should be in the continuum regime based on the heat switch geometry and gas conductivity, but what about the transition regime described by eq. D.5? As the name implies this regime in flow is somewhere in between a continuum and free-molecular flow. In this regime the mechanisms that underpin thermal conduction become poorly defined.

To connect the continuum and free-molecular regimes we can devise a simple two-fluid model, with each phase having their own thermal conductance κ_{cont} and κ_{fm}. We can draw this as a parallel circuit model where one path is a variable thermal conductor ($\kappa_{fm}(T)$)and the other has a fixed conductance (κ_{cond}) as shown in fig. D.3 (a). Some power $\dot{Q}_{in}$ is applied leading to a temperature at the hot end of T_h, while the cold end is fixed to a thermal sink at T_c. We can write the total thermal conductance of the circuit as:

$$\kappa(T) = \left(\frac{1}{\kappa_{fm}(T)} + \frac{1}{\kappa_{cont}} \right)^{-1} \tag{D.20}$$

Applying eq. D.19 in place of $\kappa_{fm}(T)$ in eq. D.20, first by dividing the left hand side by $\Delta T = T_h - T_c$ to give the thermal conductivity we arrive at:

$$\kappa(T) = \left(\frac{1}{C_{ads} \exp(-Q_{ads}/RT_{ads})} + \frac{1}{\kappa_{cont}} \right)^{-1} \tag{D.21}$$

Here the constants of eq. D.19 are combined into C_{ads} for compactness.

To see how this function behaves over different flow regimes it is informative to plot the conductivity against the Knudsen number K_n. Combining eqs. D.1, D.3, and D.18 we see that $k_n \propto \exp(Q_{ads}/RT_{ads})$. Fig. D.3 (b) shows the expected conductivity vs K_n for a Q_{ads} of $2700 J/mol$, κ_{cont} or $140 mW/K$, and C_{ads} of $3.5 \times 10^5 W/K$, which are derived in part from the geometry of the heat switch described in section D.3.2 and typical values for adsorbant P_0 and Q_{ads} [177, 91].

Fig. D.3 (b) gives a nice visualization of the switch behavior in the three regimes, as described by eqns. D.4-D.6, with the continuum and free-molecular regimes demonstrating the expected pressure independence and dependence respectively. Likewise the transition regime demonstrates only weak pressure dependence. As we will see in section D.4.2, the two-fluid model described by eq. D.20 shows good agreement with measured gas-gap heat switch performance.

D.3 Heat-switch design and construction

D.3.1 Performance requirements and constraints

For adsorption pumping applications the primary constraint is the maximum conductivity. As shown in the previous section this maximum conductivity occurs when the pressure of the exchange gas pressure is high enough that the mean-free path is much shorter than the gap width. This pressure independent value, κ_{cont} is described by the simple 1D conductivity equation D.7, which is directly proportional to the exchange surface area of the switch, and

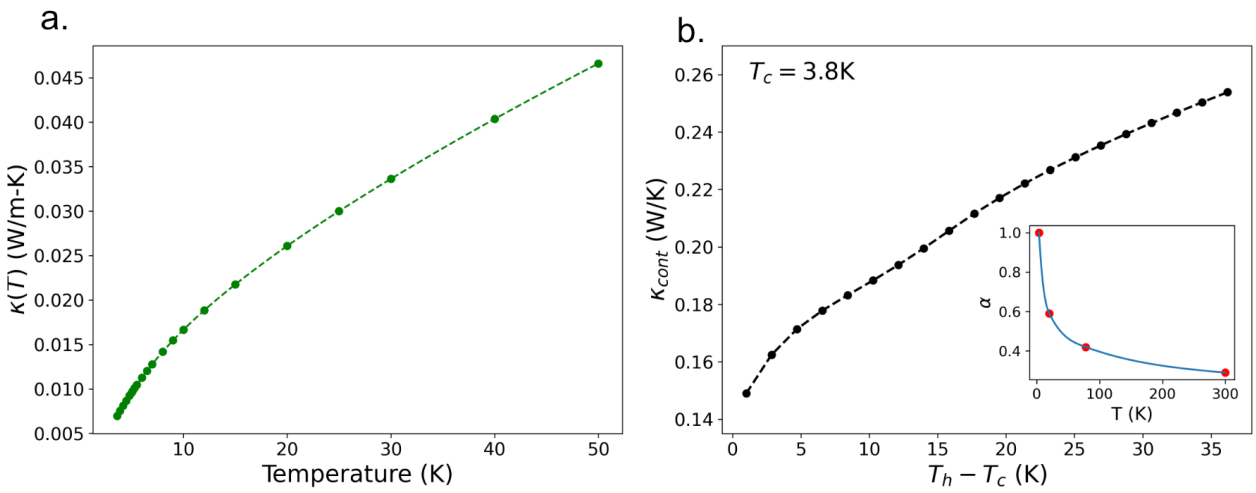

Figure D.4: **Thermal conduction of He^4 gas**: (a) shows the thermal conductance of helium gas from the critical point to 50K. (b) shows the thermal conduction of our heat switch in the continuous regime (i.e. not temperature dependent) for various temperature differentials, with T_c fixed at 3.8K.

inversely proportional to the gap width. One difficulty in estimating the on-state performance is the accommodation coefficients α_c and α_h, and subsequently the accommodation factor F_α for a given system. The accommodation coefficient is dependent on surface composition, temperature, and roughness, and is empirically derived for each surface. Low-temperature data on helium accommodation coefficients is scarce, but for most metals it ranges from $\approx 0.6 - 1$ in the range of $T \approx 20 - 3.5$K [96, 21].

For our fridge, with molar flow rates of $100 - 200\mu$mol/sec, we can expect a dissipated power of $200-400$mW at the 4K plate. For efficient adsorption to occur it is also necessary to reduce the heat lift of the pumps during operation. For this reason a maximum conductivity of $100-200$mW/K is ideal to ensure a lift of less than 4K at full power dissipation. Using this number as a guide we can find a suitable exchange area for our heat-switch by using eq. D.7. Using eq. D.15, we can estimate the accommodation factor making the approximation that, with a gap much smaller than the diameter of our concentric cylinders, $A_c \approx A_h$. Taking a worst-case estimate of $\alpha_c \approx 0.9$ and $\alpha_h \approx 0.6$, $F_\alpha \approx 0.56$.

To calculate the thermal conductivity of the helium in our gap we need to know conductivity as a function of temperature ($\kappa_{He}(T)$). Here we use the conductivity under isobaric conditions, instead of saturated vapor pressure, since at the temperature this switch will operate the saturated vapor pressure is $> 100kPa$, whereas the pressure needed for continuum conduction with a L_e for the heat-switch shown in D.3 is $\sim 4Pa$. The thermal conductivity of helium-4 from 3.5K$-$50K is shown in fig. D.4 a at 10kPa from [10].

With an estimation for the accommodation coefficient of helium taken from table 5.6 of [21] (shown in Fig. D.4 (b) inset), the conductivity of helium, and an $A_c/l_{gap} = 20$, we can expect a continuum conductivity of $\kappa_{cont} \sim 168$mW/K given a $\Delta T \sim 4$K. To obtain a Kn value that gives the desired switching characteristics as shown in fig. D.3 (b), an $l_{gap} \sim 100\mu$m is desired, giving a desired exchange area $A_c \sim 20\text{cm}^2$.

D.3.2 Heat switch design and components

In the previous section we outlined the design parameters required for a target performance. Translating this to a real-world design requires a few additional considerations that have not been discussed. These include reducing parasitic load while the switch is in the *off-state*, the creation of hermetic seals, and heating and cooling the adsorber with a high degree of control. There are additional practical design considerations one has to make concerning the construction of the heat-switch itself. The high-tolerances of the gap and close fitting of the parts, along with the need for helium leak-tight joints makes gas-gap heat-switch construction particularly challenging, so careful design is necessary for successful implementation.

Figure D.5 shows a cutaway of the heat switch design. To make the heat switch easier to manufacture and more compact the switch is a double concentric design, with an inner and outer cylinder with a total surface area of 21cm^2 out of OFHC copper. A stainless steel (304) tube of $\varnothing$22.2mm with a wall thickness of .23mm. Two flanges provide additional support to the thin tube, while also providing a surface for an indium seal between the tube

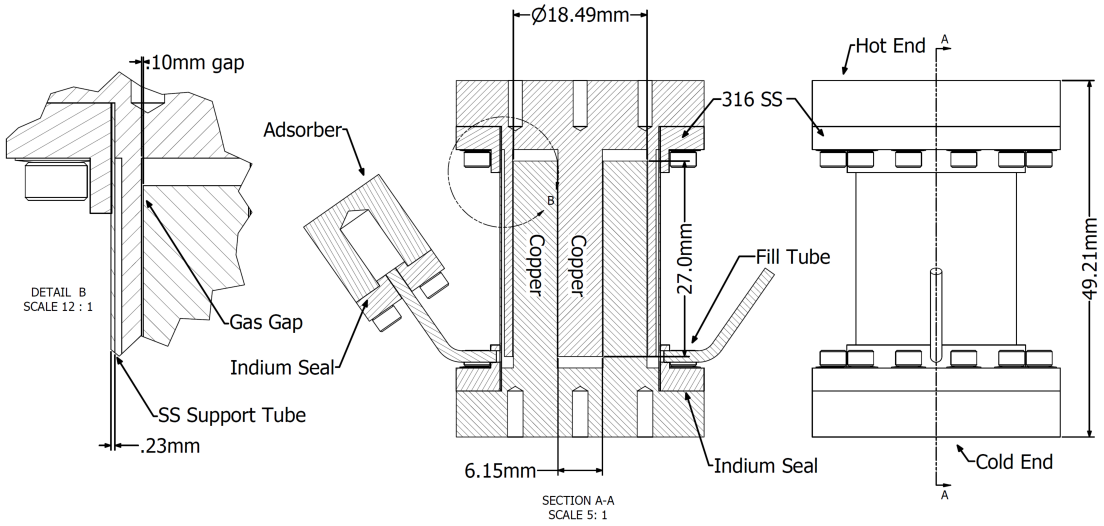

Figure D.5: **Cross-section of THeIA gas-gap heat switches**

and top and bottom heat-switch halves.

From the support tube dimenstions we can compute the parasitic heat load through the tube using the NIST TRC database for 304 stainless steel [95], as shown in fig. D.6. From this and fig. D.4 (b) we can also determine a theoretical aspect ratio of $\kappa_{on}/\kappa_{off} > 240$, for the range $35 > \Delta T > 1$ with a T_c of 3.8K, which is more than acceptable for our desired application.

The adsorber is made of a copper can that is attached to the heat switch via a $\varnothing \sim$1.5mm with a 0.5mm wall thickness. A stainless steel flange with an indium seal allows the adsorber to be made in two pieces. Additional holes in the copper housing are created for a 10kOhm resistor for the heater and a for a DT670 silicon diode from Lakeshore Cryotronics ©.

D.3.3 Construction and gas-charging

The tight tolerances of the heat switch requires extra attention to be paid to the accumulation of small dimensional deviations at every step. The most important aspect of this is the concentricity and parallelism of the various components with respect to one another, and the creation of datum surfaces that properly indicate the location of the two heat-switch

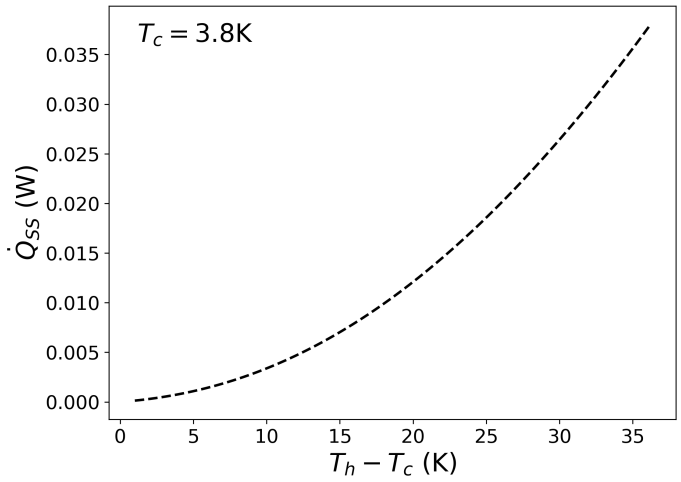

Figure D.6: **Expected parasitic heat load through the stainless steel supporting shell for various temperatures**

halves to one another. This is most important for the supporting shell, which is made of three components that are assembled into a contiguous leak-tight housing.

Fig. D.7 (a-c) shows the construction of the shell. Fig. D.7 (a) shows the jig used during silver brazing [1] of the top and bottom flange ring to the main support tube to maintain parallelism of the sealing surfaces. Fig. D.7 (b) shows a custom mandrel for repeatable indication of the now brazed tube for final machining of the inner datum surfaces that indicate the top and bottom of the heat switch, before Fig. D.7 (c) showing complete assembly with fill and adsorber tube.

In addition to the parts mentioned above, a small 4K radiation shield is placed over the adsorber for passive cooling. ∅0.5mm pure indium wire was used to seal the switch top and bottoms halves to the support tube and seal the adsorber can. A 1/8" tube weld Swagelok VCR©[2] was also brazed onto the fill tube for easy attachment of either a fill or pinch-off tube for helium charging. An amount of ∼0.1g of activated charcoal [3] was sealed into the adsorber volume, and packed, without additional epoxy, to ensure good thermalization to the

1. AWS BAg-36 (P07454) alloy was used for all braze joints

2. Swaglok SS-1-VCR-3

3. Fisher 05-685A

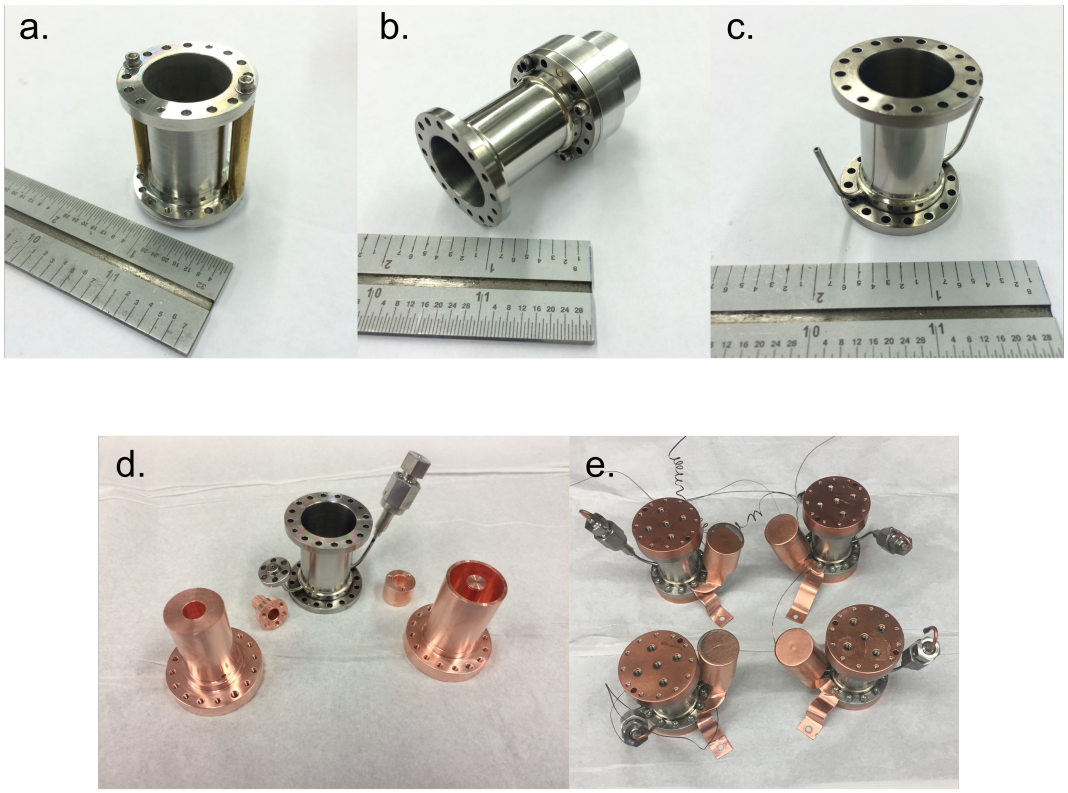

Figure D.7: **Pictures of manufactured parts and steps of construction for gas-gap heat switch**: (a) shows a SS shell assmbly jigged for silver brazing. Brass posts ensure that the end-flanges remain parallel to one another while udnergoing heating. (b) shows a brazed shell on a mandrel to undergo post-braze machining to finalize the dimensions and enser parallelness and concentricity. (c) shows a nearly complete shell with the fill and adsorber tubes in place. (d) shows the complete finalized parts for a version 1 THeIA switch (now in Breton), while (e) shows four version 2 switches following construction and testing.

copper adsorber. For the heater a 10kOhm 0805 size surface mount resistor is soldered to two 43 AWG phosphor-bronze wires before being encapsulated in Stycast© 2850FT (using 23LV catalyst) and finally epoxied into a hole machined into the adsorber housing. Lakeshore© DT-670-SD silicon diode thermometers are also inserted into the housing using GE varnish and connected with 4×43 phosphor bronze ribbon wire.

A crucial step in the production of a gas-gap heat switch is the charging of the switch with he-4 gas. Fig. D.8 shows the charging circuit used. The heat switch fill line attached to a ∅1.5mm OFHC copper tube to a regulated high-pressure tank of UHP (99.999% he-4) via an isolation valve. A second isolation valve is attached to a dry 2-stage ultra-high

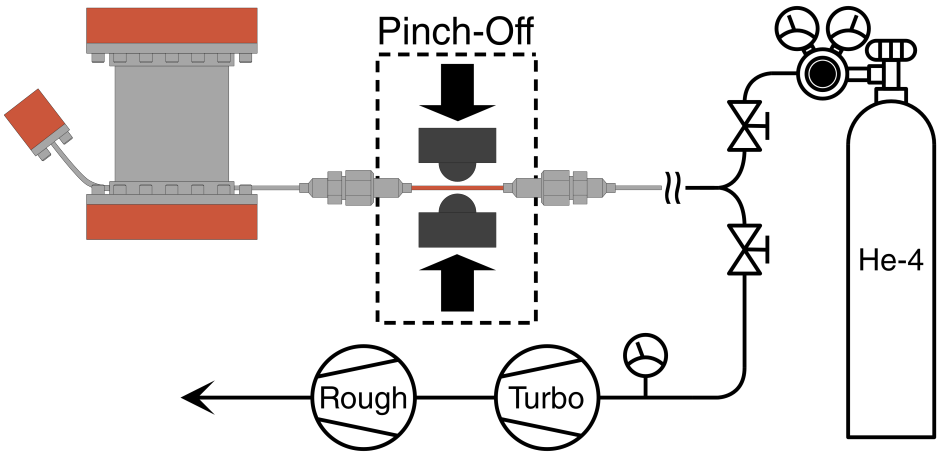

Figure D.8: **Gas filling and pinch-off layout**: The heat-switch is evacuated and filled via a 1/16" (1.5mm) OD copper tube. Two isolation valves allow for the evacuation of the fill-lines via a turbo and dry lobe vacuum pump, while a regulated ultra high-purity UHP He^4 with a compound gauge allows for precision filling of the switch.

vacuum pump capable of pressures $P < 1 \times 10^{-7}$mbar. The switch is initially pumped down to $P < 1 \times 10^{-5}$mbar with the tank isolation valve closed. The heat-switch is also leak-checked at this time. Care should be taken to purge the helium regulator to reduce water or air ingress. Once the base pressure is reached the vacuum isolation valve is closed, before opening the filling side valve. The regulator outlet pressure, measured via a compound gauge, is slowly increased to the desired pressure. Finally a hydraulic pinch off tool[4] is used to crimp the copper fill tube shut, permanently sealing the heat-switch. If refilling is needed the fill tube can be replaced via the attached Swagelok VCR.

D.4 Characterization and performance

D.4.1 Experimental setup

The heat-switch is attached to the 4K plate of the refrigerator. A Lakeshore DT-670-CU thermometer is placed on the 4K plate and the top of the heat-switch for measuring T_c and

4. Solid Sealing Technology KT35046

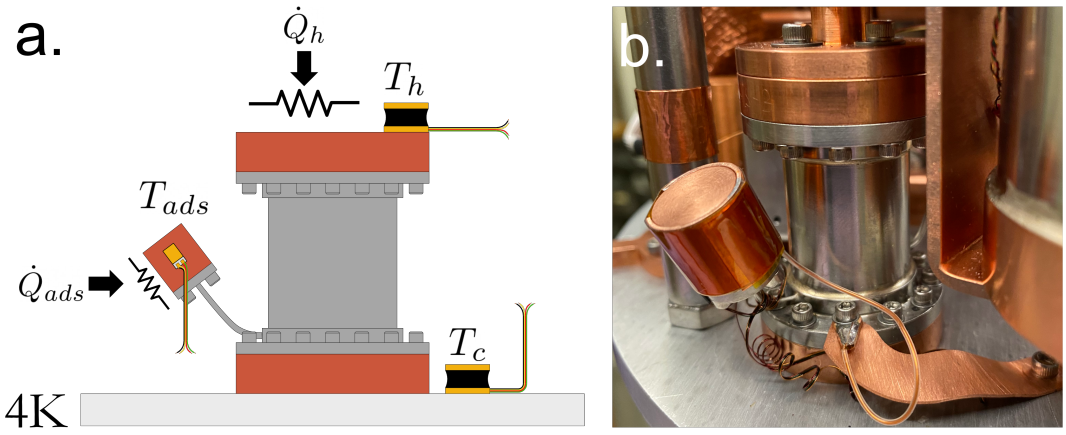

Figure D.9: **Gas-gap heat-switch measurement layout and location in fridge**: (a) the setup of the heat-switch measurement. (b) the location of the pump in the fridge under pump 2 of the THeIA refrigerator. Note the small copper wire shorting the adsorber to the plate to allow for faster reset of the adsorber temperature, allowing for more accurate PID stabilization.

T_h respectively (locations showed in fig. D.9 (a). A ∼36ohm, 15W cartridge heater [5] is attached via a 4-wire configuration for more accurate measurement of the applied power.

For accurate measurement of the adsorber temperature vs the thermal conduction a way of stabilizing T_{ads}. Closed-loop PID control, using proportional-on-measurement (POM) feedback was implemented. A BK Precision 8130B provided the power for the adsorber heater, via a 10kOhm resistor directly mounted to the adsorber. A Lakeshore DT-670-SD thermometer is mounted to the adsorber using GE varnish and wrapped tightly to the adsorber surface with Kapton tape. A small wire attached to a tightly fitting copper sleeve via lead-tin solder is attached at one end to the adsorber and at the other end to the 4K plate. The length of wire is chosen to guarantee a few mW of heat conduction with the adsorber at 15K. This gives the adsorber a "damping" term in the PID control, allowing for more accurate and faster temperature stablization.

5. Thor Labs HT15W

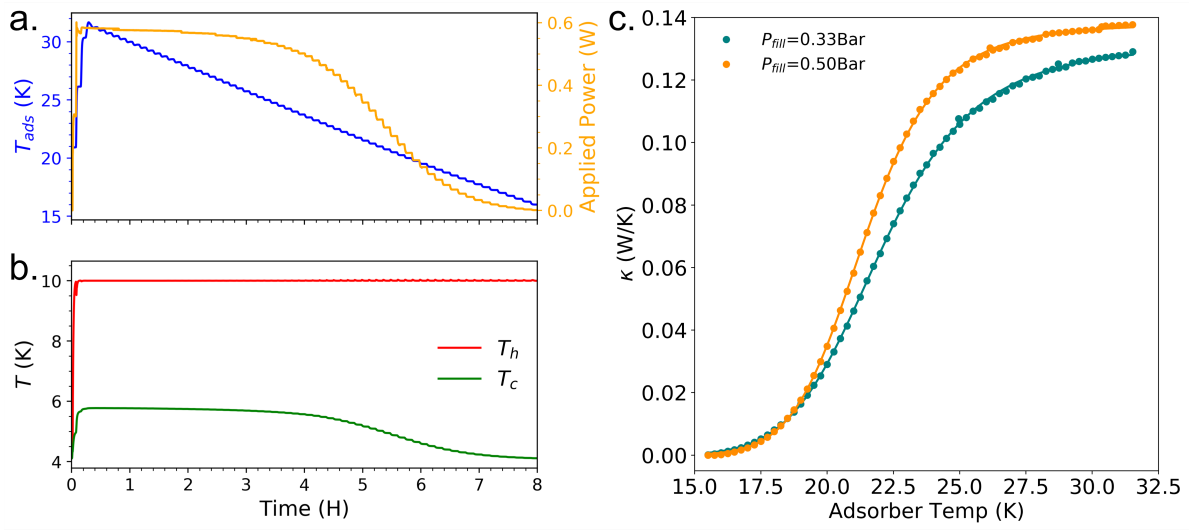

Figure D.10: **Heat switch performance**: (a) shows the variation of adsorber temperature and applied power required to stabilize T_h over time. (b) shows the change in T_h and T_c over the same time span. (c) shows the thermal conductivity of two identical version 2 THeIA heat switches with different fill pressures.

D.4.2 Heat-switch performance and analysis

Two identical heat switches were filled with either 0.33 bar or 0.5 bar. T_h was stabilized at a fixed temperature of 10K, while the the adsorber temperature was varied. Between every set adsorber temperature the system was allowed to stabilize for 5 minutes. Once stable the applied power was determined, and the temperature differential was computed. From this a conductivity for the switch could be determined. Fig D.10 shows this process being done, with (a) showing adsorber power and applied power required to stabilize T_h, (b) showing T_h and T_c, and (c) showing the scaling of thermal conduction to the switch vs the adsorber temperature.

Fig. D.10 (c) shows that the higher fill pressure leads to higher thermal conduction in the continuous regime for the same adsorber temperature, indicating that the equilibrium isosteric adsorption pressure for the given ratio of helium to adsorbant likely favors the larger fill pressure. A second observation is that, for the given fill pressures both perform comparably to the predicted continuous regime conduction performance predicted in Fig. D.4

(b). Another observation is how *closely* our model, outlined in Eq. D.19, fits our data, with only a proportionality factor needed for the mess of fixed variables in our system. This thermodynamic model, and the data above, are important parameters for the complete control of the continous fridge, as they allow us to dynamically control the flow of heat out of our pumps, giving us full closed-loop control over the continuous cycle.

REFERENCES

[1] M. R. Alexander, G. E. Thompson, X. Zhou, G. Beamson, and N. Fairley. Quantification of oxide film thickness at the surface of aluminium using xps. *Surface and Interface Analysis*, 34(1):485–489, 2002.

[2] M. S. Allman, J. D. Whittaker, M. Castellanos-Beltran, K. Cicak, F. da Silva, M. P. DeFeo, F. Lecocq, A. Sirois, J. D. Teufel, J. Aumentado, and R. W. Simmonds. Tunable resonant and nonresonant interactions between a phase qubit and *lc* resonator. *Phys. Rev. Lett.*, 112:123601, Mar 2014.

[3] AMD/Xilinx. *ZCU216 Evaluation Board User Guide(UG1390).*

[4] C. P. Anderson, E. O. Glen, C. Zeledon, A. Bourassa, Y. Jin, Y. Zhu, C. Vorwerk, A. L. Crook, H. Abe, J. Ul-Hassan, T. Ohshima, N. T. Son, G. Galli, and D. D. Awschalom. Five-second coherence of a single spin with single-shot readout in silicon carbide. *Science Advances*, 8(5):eabm5912, 2022.

[5] G. Andersson, A. L. O. Bilobran, M. Scigliuzzo, M. M. de Lima, J. H. Cole, and P. Delsing. Acoustic spectral hole-burning in a two-level system ensemble. *npj Quantum Information*, 7(1):15, Jan 2021.

[6] C. Antoine, A. Aspart, M. Berthelot, J. Poupeau, F. Valin, and Y. Gasser. Morphological and chemical studies of nb samples after various surface treatments. Technical report, SIS-2002-103, 1999.

[7] C. Antoine, B. Aune, B. Bonin, J. Cavedon, M. Juillard, A. Godin, C. Henriot, P. Leconte, H. Safa, A. Veyssiere, et al. The role of atomic hydrogen in q-degradation of niobium superconducting rf cavities: analytical point of view. In *Proc. of 5th Workshop on RF Superconductivity, DESY, Hamburg, Germany*, pages 616–634, 1991.

[8] C. Z. Antoine. How to achieve the best srf performance: (practical) limitations and possible solutions, 2015.

[9] C. Z. Antoine and S. Berry. H in niobium: Origin and method of detection. *AIP Conference Proceedings*, 671(1):176–189, 2003.

[10] V. Arp, R. McCarty, and D. Friend. Thermophysical properties of helium-4 from 0.8 to 1500 k with pressures to 2000 mpa, 1998-02-19 00:02:00 1998.

[11] P. Arrangoiz-Arriola, E. A. Wollack, Z. Wang, M. Pechal, W. Jiang, T. P. McKenna, J. D. Witmer, R. Van Laer, and A. H. Safavi-Naeini. Resolving the energy levels of a nanomechanical oscillator. *Nature*, 571(7766):537–540, Jul 2019.

[12] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney,

M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, Oct 2019.

[13] N. Ashcroft and N. Mermin. *Solid State Physics*. Cengage Learning, 2011.

[14] A. Aspart and C. Antoine. Study of the chemical behavior of hydrofluoric, nitric and sulfuric acids mixtures applied to niobium polishing. *Applied Surface Science*, 227(1):17–29, 2004.

[15] M. Aspelmeyer, T. J. Kippenberg, and F. Marquardt. Cavity optomechanics. *Rev. Mod. Phys.*, 86:1391–1452, Dec 2014.

[16] B. Aune, R. Bandelmann, D. Bloess, B. Bonin, A. Bosotti, M. Champion, C. Crawford, G. Deppe, B. Dwersteg, D. A. Edwards, H. T. Edwards, M. Ferrario, M. Fouaidy, P.-D. Gall, A. Gamp, A. Gössel, J. Graber, D. Hubert, M. Hüning, M. Juillard, T. Junquera, H. Kaiser, G. Kreps, M. Kuchnir, R. Lange, M. Leenen, M. Liepe, L. Lilje, A. Matheisen, W.-D. Möller, A. Mosnier, H. Padamsee, C. Pagani, M. Pekeler, H.-B. Peters, O. Peters, D. Proch, K. Rehlich, D. Reschke, H. Safa, T. Schilcher, P. Schmüser, J. Sekutowicz, S. Simrock, W. Singer, M. Tigner, D. Trines, K. Twarowski, G. Weichert, J. Weisend, J. Wojtkiewicz, S. Wolff, and K. Zapfe. Superconducting tesla cavities. *Phys. Rev. ST Accel. Beams*, 3:092001, Sep 2000.

[17] C. J. Axline. *Building Blocks for Modular Circuit QED Quantum Computing*. PhD thesis, 2018. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2021-09-28.

[18] J. Bardeen, L. N. Cooper, and J. R. Schrieffer. Microscopic theory of superconductivity. *Physical Review*, 106(1):162, 1957.

[19] F. Barkov, A. Romanenko, Y. Trenikhina, and A. Grassellino. Precipitation of hydrides in high purity niobium after different treatments. *Journal of Applied Physics*, 114(16):164904, 2013.

[20] K. Barnes, P. Battaglino, B. J. Bloom, K. Cassella, R. Coxe, N. Crisosto, J. P. King, S. S. Kondov, K. Kotru, S. C. Larsen, J. Lauigan, B. J. Lester, M. McDonald, E. Megidish, S. Narayanaswami, C. Nishiguchi, R. Notermans, L. S. Peng, A. Ryou, T.-Y. Wu, and M. Yarwood. Assembly and coherent control of a register of nuclear spin qubits. *Nature Communications*, 13(1):2779, May 2022.

[21] R. Barron. *Cryogenic Heat Transfer*. Series in Chemical and Mechanical Engineering. Taylor & Francis, 1999.

[22] T. Barron and G. White. *Heat Capacity and Thermal Expansion at Low Temperatures*. International Cryogenics Monograph Series. Springer US, 2012.

[23] K. Berdel, J. Rivas, P. Bolivar, P. de Maagt, and H. Kurz. Temperature dependence of the permittivity and loss tangent of high-permittivity materials at terahertz frequencies. *IEEE Transactions on Microwave Theory and Techniques*, 53(4):1266–1271, 2005.

[24] H. A. Bethe. Theory of diffraction by small holes. *Phys. Rev.*, 66:163–182, Oct 1944.

[25] L. S. Bishop. *Circuit Quantum Electrodynamics*. PhD thesis, 2010.

[26] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff. Circuit quantum electrodynamics. *Rev. Mod. Phys.*, 93:025005, May 2021.

[27] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf. Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation. *Phys. Rev. A*, 69:062320, Jun 2004.

[28] C. Boffo, C. Cooper, A. Rowe, and G. Galasso. Optimization of the BCP processing of elliptical nb srf cavities. *Conf. Proc. C*, 060626:469–471, 2006.

[29] V. Bouchiat. *Quantum fluctuations of the charge in single electron and single cooper pair devices*. PhD thesis, 1998.

[30] S. B. Bravyi and A. Y. Kitaev. Quantum codes on a lattice with boundary. *arXiv preprint quant-ph/9811052*, 1998.

[31] T. L. Brecht. *Micromachined quantum circuits*. PhD thesis, 2017. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2021-09-28.

[32] N. T. Bronn, Y. Liu, J. B. Hertzberg, A. D. Córcoles, A. A. Houck, J. M. Gambetta, and J. M. Chow. Broadband filters for abatement of spontaneous emission in circuit quantum electrodynamics. *Applied Physics Letters*, 107(17):172601, 2015.

[33] N. T. Bronn, E. Magesan, N. A. Masluk, J. M. Chow, J. M. Gambetta, and M. Steffen. Reducing spontaneous emission in circuit quantum electrodynamics by a combined readout/filter technique. *IEEE Transactions on Applied Superconductivity*, 25(5):1–10, oct 2015.

[34] P. Campagne-Ibarcq, A. Eickbusch, S. Touzard, E. Zalys-Geller, N. Frattini, V. Sivak, P. Reinhold, S. Puri, S. Shankar, R. Schoelkopf, et al. Quantum error correction of a qubit encoded in grid states of an oscillator. *Nature*, 584(7821):368–372, 2020.

[35] P. Campagne-Ibarcq, A. Eickbusch, S. Touzard, E. Zalys-Geller, N. E. Frattini, V. V. Sivak, P. Reinhold, S. Puri, S. Shankar, R. J. Schoelkopf, L. Frunzio, M. Mirrahimi, and M. H. Devoret. Quantum error correction of a qubit encoded in grid states of an oscillator. *Nature*, 584(7821):368–372, Aug 2020.

[36] E. R. Canavan, M. J. Dipino, M. Jackson, J. Panek, P. J. Shirron, and J. G. Tuttle. A magnetoresistive heat switch for the continuous adr. *AIP Conference Proceedings*, 613(1):1183–1190, 2002.

[37] S. Chakram, K. He, A. Dixit, A. Oriani, R. Naik, N. Leung, H. Kwon, W.-L. Ma, L. Jiang, and D. Schuster. Multimode interactions and entanglement through photon blockade. *arXiv:2010.15292*, 2020.

[38] S. Chakram, K. He, A. V. Dixit, A. E. Oriani, R. K. Naik, N. Leung, H. Kwon, W.-L. Ma, L. Jiang, and D. I. Schuster. Multimode photon blockade, 2020.

[39] S. Chakram, A. E. Oriani, R. K. Naik, A. V. Dixit, K. He, A. Agrawal, H. Kwon, and D. I. Schuster. Seamless high-q microwave cavities for multimode circuit quantum electrodynamics. *Phys. Rev. Lett.*, 127:107701, Aug 2021.

[40] Y. Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, et al. Qubit architecture with high coherence and fast tunable coupling. *Phys. Rev. Lett.*, 113(22):220502, 2014.

[41] W. C. Chew. Lectures on theory of microwave and optical waveguides, 2021.

[42] W. C. Chew, K.-H. Lin, J. Friedich, and C. H. Chan. Reflection and transmission operators for general discontinuities in waveguides. *Journal of electromagnetic waves and applications*, 5(8):819–834, 1991.

[43] V. Chouhan, S. Kato, Y. Ida, K. Nii, T. Yamaguchi, and H. Hayano. Chemical analysis and field emission study of electropolished niobium surface containing synthesized niobium oxide particles. *Materials Chemistry and Physics*, 259:124044, 2021.

[44] Y. Chu, P. Kharel, T. Yoon, L. Frunzio, P. T. Rakich, and R. J. Schoelkopf. Creation and control of multi-phonon fock states in a bulk acoustic-wave resonator. *Nature*, 563(7733):666–670, Nov 2018.

[45] G. Ciovati, P. Dhakal, and G. R. Myneni. Superconducting radio-frequency cavities made from medium and low-purity niobium ingots. *Superconductor Science and Technology*, 29(6):064002, 2016.

[46] A. Clausing. Heat transfer at the interface of dissimilar metals—the influence of thermal strain. *International Journal of Heat and Mass Transfer*, 9(8):791–801, 1966.

[47] A. M. Clausing and B. T. Chao. Thermal Contact Resistance in a Vacuum Environment. *Journal of Heat Transfer*, 87(2):243–250, 05 1965.

[48] A. A. Clerk and D. W. Utami. Using a qubit to measure photon-number statistics of a driven thermal oscillator. *Physical Review A*, 75(4):042302, 2007.

[49] J. F. Cochran and D. E. Mapother. Superconducting transition in aluminum. *Phys. Rev.*, 111:132–142, Jul 1958.

[50] R. Collin. *Foundations for Microwave Engineering*. IEEE Press Series on Electromagnetic Wave Theory. Wiley, 2001.

[51] L. N. Cooper. Bound electron pairs in a degenerate fermi gas. *Phys. Rev.*, 104:1189–1190, Nov 1956.

[52] R. Corruccini and J. Gneiwek. *Thermal Expansion of Technical Solids at Low Temperatures; a compilation from literature*. United States. National Bureau of Standards. Monograph. National Bureau of Standards, Washington,, 1961.

[53] A. Cottet. *Implémentation d'un bit quantique dans un circuit supraconducteur/Implementation of a quantum bit in a superconducting circuit*. PhD thesis, 2002.

[54] D. L. Creedon, M. Goryachev, N. Kostylev, T. B. Sercombe, and M. E. Tobar. A 3d printed superconducting aluminium microwave cavity. *Applied Physics Letters*, 109:032601, 2016.

[55] M. H. Devoret and R. J. Schoelkopf. Superconducting circuits for quantum information: An outlook. *Science*, 339(6124):1169–1174, 2013.

[56] P. Dhakal. Nitrogen doping and infusion in srf cavities: A review. *Physics Open*, 5:100034, 2020.

[57] P. Dhakal, G. Ciovati, P. Kneisel, and G. R. Myneni. Enhancement in quality factor of srf niobium cavities by material diffusion. *IEEE Transactions on Applied Superconductivity*, 25(3):1–4, 2015.

[58] P. Dhakal, G. Ciovati, U. Pudasaini, S. Chetri, S. Balachandran, and P. J. Lee. Surface characterization of nitrogen-doped high purity niobium coupons compared with superconducting rf cavity performance. *Phys. Rev. Accel. Beams*, 22:122002, Dec 2019.

[59] P. Dhakal, G. Ciovati, W. Rigby, J. Wallace, and G. R. Myneni. Design and performance of a new induction furnace for heat treatment of superconducting radiofrequency niobium cavities. *Review of Scientific Instruments*, 83(6):065105, 2012.

[60] M. Dietrich, A. Euler, and G. Thummes. A lightweight thermal heat switch for redundant cryocooling on satellites. *Cryogenics*, 83:31–34, Apr 2017.

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

[62] A. V. Dixit. *Searching for Dark Matter with Superconducting Qubits*. PhD thesis, The University of Chicago, 2021-06.

[63] A. V. Dixit, S. Chakram, K. He, A. Agrawal, R. K. Naik, D. I. Schuster, and A. Chou. Searching for dark matter with a superconducting qubit. *arXiv:2008.12231*, 2020.

[64] M. Dressel. Electrodynamics of metallic superconductors. *Advances in Condensed Matter Physics*, 2013:104379, Jul 2013.

[65] N. Earnest, S. Chakram, Y. Lu, N. Irons, R. K. Naik, N. Leung, L. Ocola, D. A. Czaplewski, B. Baker, J. Lawrence, J. Koch, and D. I. Schuster. Realization of a Λ system with metastable states of a capacitively shunted fluxonium. *Phys. Rev. Lett.*, 120:150504, Apr 2018.

[66] A. Eickbusch, V. Sivak, A. Z. Ding, S. S. Elder, S. R. Jha, J. Venkatraman, B. Royer, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret. Fast universal control of an oscillator with weak dispersive coupling to a qubit, 2021.

[67] D. K. Finnemore, T. F. Stromberg, and C. A. Swenson. Superconducting properties of high-purity niobium. *Phys. Rev.*, 149:231–243, Sep 1966.

[68] D. C. Ford, L. D. Cooley, and D. N. Seidman. First-principles calculations of niobium hydride formation in superconducting radio-frequency cavities. *Superconductor Science Technology*, 26(9):095002, Sept. 2013.

[69] N. E. Frattini, U. Vool, S. Shankar, A. Narla, K. M. Sliwa, and M. H. Devoret. 3-wave mixing josephson dipole element. *Applied Physics Letters*, 110(22):222603, 2017.

[70] P. Frigola, R. Agustsson, L. Faillace, A. Murokh, G. Ciovati, W. A. Clemens, P. Dhakal, F. Marhauser, R. A. Rimmer, J. K. Spradlin, et al. Advance additive manufacturing method for srf cavities of various geometries. Technical report, Thomas Jefferson National Accelerator Facility (TJNAF), Newport News, VA ..., 2015.

[71] N. Fuschillo, B. Lalevic, and N. Annamalai. Dielectric properties of amorphous nb2o5 thin films. *Thin Solid Films*, 30(1):145–154, 1975.

[72] J. Gambetta, A. Blais, D. I. Schuster, A. Wallraff, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Qubit-photon interactions in a cavity: Measurement-induced dephasing and number splitting. *Phys. Rev. A*, 74:042318, Oct 2006.

[73] J. Gao. *The physics of superconducting microwave resonators*. PhD thesis, 2008.

[74] J. Gao, M. Daal, A. Vayonakis, S. Kumar, J. Zmuidzinas, B. Sadoulet, B. A. Mazin, P. K. Day, and H. G. Leduc. Experimental evidence for a surface distribution of two-level systems in superconducting lithographed microwave resonators. *Applied Physics Letters*, 92(15):152505, 2008.

[75] Y. Y. Gao, B. J. Lester, K. S. Chou, L. Frunzio, M. H. Devoret, L. Jiang, S. Girvin, and R. J. Schoelkopf. Entanglement of bosonic modes through an engineered exchange interaction. *Nature*, 566(7745):509–512, 2019.

[76] J. Gere and S. Timoshenko. *Mechanics of Materials*. Stanley Thornes, 1999.

[77] V. L. Ginzburg and L. D. Landau. On the Theory of superconductivity. *Zh. Eksp. Teor. Fiz.*, 20:1064–1082, 1950.

[78] S. M. Girvin. Circuit qed: superconducting qubits coupled to microwave photons.

[79] K. Gloos, C. Mitschka, F. Pobell, and P. Smeibidl. Thermal conductivity of normal and superconducting metals. *Cryogenics*, 30(1):14–18, 1990.

[80] K. Gloos, P. Smeibidl, C. Kennedy, A. Singsaas, P. Sekowski, R. Mueller, and F. Pobell. The bayreuth nuclear demagnetization refrigerator. *Journal of low temperature physics*, 73(1):101–136, 1988.

[81] A. Grassellino, A. Romanenko, A. Crawford, O. Melnychuk, A. Rowe, M. Wong, C. Cooper, D. Sergatskov, D. Bice, Y. Trenikhina, L. D. Cooley, C. Ginsburg, and R. D. Kephart. Fermilab experience of post-annealing losses in srf niobium cavities due to furnace contamination and the ways to its mitigation: a pathway to processing simplification and quality factor improvement, 2013.

[82] A. Grassellino, A. Romanenko, D. Sergatskov, O. Melnychuk, Y. Trenikhina, A. Crawford, A. Rowe, M. Wong, T. Khabiboulline, and F. Barkov. Nitrogen and argon doping of niobium for superconducting radio frequency cavities: a pathway to highly efficient accelerating structures. *Superconductor Science and Technology*, 26(10):102001, aug 2013.

[83] M. Grundner and J. Halbritter. Xps and aes studies on oxide growth and oxide coatings on niobium. *Journal of Applied Physics*, 51(1):397–405, 1980.

[84] C. T. Hann, C.-L. Zou, Y. Zhang, Y. Chu, R. J. Schoelkopf, S. Girvin, and L. Jiang. Hardware-efficient quantum random access memory with hybrid quantum acoustic systems. *Physical Review Letters*, 123(25), Dec 2019.

[85] A. K. Harman, S. Ninomiya, and S. Adachi. Optical constants of sapphire (-al2o3) single crystals. *Journal of Applied Physics*, 76(12):8032–8036, 1994.

[86] S. Haroche and J.-M. Raimond. *Exploring the Quantum: Atoms, Cavities, and Photons*. Oxford University Press, 2006.

[87] R. W. Heeres, P. Reinhold, N. Ofek, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf. Implementing a universal gate set on a logical qubit encoded in an oscillator. *Nature communications*, 8(1):1–7, 2017.

[88] R. W. Heeres, B. Vlastakis, E. Holland, S. Krastanov, V. V. Albert, L. Frunzio, L. Jiang, and R. J. Schoelkopf. Cavity state manipulation using photon-number selective phase gates. *Physical review letters*, 115(13):137002, 2015.

[89] P. Heidler, C. M. F. Schneider, K. Kustura, C. Gonzalez-Ballestero, O. Romero-Isart, and G. Kirchmair. Non-markovian effects of two-level systems in a niobium coaxial resonator with a single-photon lifetime of 10 milliseconds. *Phys. Rev. Applied*, 16:034024, Sep 2021.

[90] J. Heinsoo, S. Inel, D. Janzso, M. Jenei, J. Kotilahti, A. Landra, C. Ockeloen-Korppi, J. Räbinä, and N. Savola. Kqcircuits, June 2021.

[91] B. Helvensteijn, A. Kashani, R. Wilcox, and A. Spivak. Modelling helium adsorption in activated carbon. Advances in cryogenic engineering, pages 1561–1568, New York, 1994. Cryogenic Engineering Conference; distributed by Plenum Press.

[92] L. Hu, Y. Ma, W. Cai, X. Mu, Y. Xu, W. Wang, Y. Wu, H. Wang, Y. Song, C.-L. Zou, S. M. Girvin, L.-M. Duan, and L. Sun. Quantum error correction and universal gate set operation on a binomial bosonic logical qubit. *Nature Physics*, 15(5):503–508, 2019.

[93] E. Jensen. Rf cavity design. *arXiv preprint arXiv:1601.05230*, 2016.

[94] K. Kao. *Dielectric Phenomena in Solids*. Elsevier Science, 2004.

[95] A. Kazakov, C. Muzny, R. Chirico, V. Diky, and M. Frenkel. Nist/trc web thermo tables - professional edition nist standard reference subscription database 3, 2002-06-30 2002.

[96] W. Keesom and G. Schmidt. Researches on heat conduction by rarefied gases. iii. the thermal accommodation coefficient of helium, neon, and hydrogen at 12–20° k. *Physica*, 4(9):828–834, 1937.

[97] M. Kelly and T. Reid. Surface processing for bulk niobium superconducting radio frequency cavities. *Superconductor Science and Technology*, 30(4):043001, 2017.

[98] M. Kervinen, I. Rissanen, and M. Sillanpää. Interfacing planar superconducting qubits with high overtone bulk acoustic phonons. *Phys. Rev. B*, 97:205443, May 2018.

[99] K. Khodjasteh and D. A. Lidar. Fault-tolerant quantum dynamical decoupling. *Phys. Rev. Lett.*, 95:180501, Oct 2005.

[100] M. L. Kinter, I. Weissman, and W. W. Stein. Chemical polish for niobium microwave structures. *Journal of Applied Physics*, 41(2):828–829, 1970.

[101] E. Knill. Quantum computing with realistically noisy devices. *Nature*, 434(7029):39–44, Mar 2005.

[102] E. Knill, R. Laflamme, and W. H. Zurek. Resilient quantum computation. *Science*, 279(5349):342–345, 1998.

[103] J. Knobloch. The "q disease" in superconducting niobium rf cavities. *AIP Conference Proceedings*, 671(1):133–150, 2003.

[104] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Charge-insensitive qubit design derived from the cooper pair box. *Phys. Rev. A*, 76:042319, Oct 2007.

[105] J. Krupka, K. Derzakowski, M. Tobar, J. Hartnett, and R. G. Geyer. Complex permittivity of some ultralow loss dielectric crystals at cryogenic temperatures. *Measurement Science and Technology*, 10(5):387–392, jan 1999.

[106] M. Kudra, J. Biznárová, A. Fadavi Roudsari, J. Burnett, D. Niepce, S. Gasparinetti, B. Wickman, and P. Delsing. High quality three-dimensional aluminum microwave cavities. *Applied Physics Letters*, 117(7):070601, 2020.

[107] M. Kudra, M. Kervinen, I. Strandberg, S. Ahmed, M. Scigliuzzo, A. Osman, D. P. Lozano, G. Ferrini, J. Bylander, A. F. Kockum, F. Quijandría, P. Delsing, and S. Gasparinetti. Robust preparation of wigner-negative states with optimized snap-displacement sequences, 2021.

[108] V. Kurlov. Sapphire: Properties, growth, and applications. In K. J. Buschow, R. W. Cahn, M. C. Flemings, B. Ilschner, E. J. Kramer, S. Mahajan, and P. Veyssière, editors, *Encyclopedia of Materials: Science and Technology*, pages 8259–8264. Elsevier, Oxford, 2001.

[109] S. V. Kutsaev, K. Taletski, R. Agustsson, P. Carriere, A. N. Cleland, Z. A. Conway, É. Dumur, A. Moro, and A. Y. Smirnov. Niobium quarter-wave resonator with the optimized shape for quantum information systems. *EPJ Quantum Technology*, 7(1):7, Apr 2020.

[110] C. U. Lei, L. Krayzman, S. Ganjam, L. Frunzio, and R. J. Schoelkopf. High coherence superconducting microwave cavities with indium bump bonding. *Applied Physics Letters*, 116(15):154002, 2020.

[111] Q.-H. Liu and W. C. Chew. Analysis of discontinuities in planar dielectric waveguides: an eigenmode propagation method. *IEEE Transactions on Microwave Theory and Techniques*, 39(3):422–430, 1991.

[112] G. S. MacCabe, H. Ren, J. Luo, J. D. Cohen, H. Zhou, A. Sipahigil, M. Mirhosseini, and O. Painter. Nano-acoustic resonator with ultralong phonon lifetime. *Science*, 370(6518):840–843, Nov 2020.

[113] R. Manenti, A. F. Kockum, A. Patterson, T. Behrle, J. Rahamim, G. Tancredi, F. Nori, and P. J. Leek. Circuit quantum acoustodynamics with surface acoustic waves. *Nature Communications*, 8(1):975, Oct 2017.

[114] J. M. Martinis, K. B. Cooper, R. McDermott, M. Steffen, M. Ansmann, K. D. Osborn, K. Cicak, S. Oh, D. P. Pappas, R. W. Simmonds, and C. C. Yu. Decoherence in josephson qubits from dielectric loss. *Phys. Rev. Lett.*, 95:210503, Nov 2005.

[115] D. C. Mattis and J. Bardeen. Theory of the anomalous skin effect in normal and superconducting metals. *Physical Review*, 111(2):412, 1958.

[116] B. A. Mazin. *Microwave kinetic inductance detectors*. PhD thesis, California Institute of Technology, 2005.

[117] B. T. McAllister, J. Bourhill, W. H. J. Ma, T. Sercombe, M. Goryachev, and M. E. Tobar. Characterisation of cryogenic material properties of 3d-printed superconducting niobium using a 3d lumped element microwave cavity. 2020.

[118] D. C. McKay, S. Filipp, A. Mezzacapo, E. Magesan, J. M. Chow, and J. M. Gambetta. Universal gate for fixed-frequency qubits via a tunable bus. *Phys. Rev. Appl.*, 6(6):064007, 2016.

[119] D. C. McKay, R. Naik, P. Reinhold, L. S. Bishop, and D. I. Schuster. High-contrast qubit interactions using multimode cavity QED. *Phys. Rev. Lett.*, 114:080501, Feb 2015.

[120] M. Merio, M. Checchin, A. C. Crawford, A. Grassellino, M. Martinello, A. Rowe, and M. Wong. Furnace n2 doping treatments at fermilab. 2015.

[121] Z. K. Minev, Z. Leghtas, S. O. Mundhada, L. Christakis, I. M. Pop, and M. H. Devoret. Energy-participation quantization of josephson circuits. *npj Quantum Information*, 7(1):131, Aug 2021.

[122] Z. K. Minev, Z. Leghtas, P. Reinhold, S. O. Mundhada, A. Diringer, D. C. Hillel, D. Z.-R. Wang, M. Facchini, P. A. Shah, and M. Devoret. pyEPR: The energy-participation-ratio (EPR) open- source framework for quantum device design, May 2021. https://github.com/zlatko-minev/pyEPR https ://pyepr-docs.readthedocs.io/en/latest/.

[123] Z. K. Minev, T. G. McConkey, J. Drysdale, P. Shah, D. Wang, M. Facchini, G. Harper, J. Blair, H. Zhang, N. Lanzillo, S. Mukesh, W. Shanks, C. Warren, and J. M. Gambetta. Qiskit Metal: An Open-Source Framework for Quantum Device Design & Analysis, 2021.

[124] R. Naik, N. Leung, S. Chakram, P. Groszkowski, Y. Lu, N. Earnest, D. McKay, J. Koch, and D. Schuster. Random access quantum information processors using multimode circuit quantum electrodynamics. *Nature communications*, 8(1):1–7, 2017.

[125] S. B. Nam. Theory of electromagnetic properties of strong-coupling and impure superconductors. ii. *Phys. Rev.*, 156:487–493, Apr 1967.

[126] M. A. Nielsen and I. Chuang. Quantum computation and quantum information, Cambridge University Press, 2000.

[127] S. E. Nigg, H. Paik, B. Vlastakis, G. Kirchmair, S. Shankar, L. Frunzio, M. H. Devoret, R. J. Schoelkopf, and S. M. Girvin. Black-box superconducting circuit quantization. *Physical Review Letters*, 108(24):240502, 2012.

[128] A. Noguchi, R. Yamazaki, Y. Tabuchi, and Y. Nakamura. Qubit-assisted transduction for a detection of surface acoustic waves near the quantum limit. *Phys. Rev. Lett.*, 119:180505, Nov 2017.

[129] T. Noguchi. Characteristics of very high q nb superconducting resonators for microwave kinetic inductance detectors. *IEEE Transactions on Applied Superconductivity*, PP:1–1, 03 2019.

[130] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. Girvin, L. Jiang, M. Mirrahimi, M. J. Devoret, and R. J. Schoelkopf. Extending the lifetime of a quantum bit with error correction in superconducting circuits. *Nature*, 536:441–445, 2016.

[131] C. Owens, A. LaChapelle, B. Saxberg, B. M. Anderson, R. Ma, J. Simon, and D. I. Schuster. Quarter-flux hofstadter lattice in a qubit-compatible microwave cavity array. *Phys. Rev. A*, 97:013818, Jan 2018.

[132] J. C. Owens, M. G. Panetta, B. Saxberg, G. Roberts, S. Chakram, R. Ma, A. Vrajitoarea, J. Simon, and D. Schuster. Chiral cavity quantum electrodynamics, 2021.

[133] H. Padamsee. *RF superconductivity: science, technology, and applications*. John Wiley & Sons, 2009.

[134] H. Padamsee, J. Knobloch, and T. Hays. *RF Superconductivity for Accelerators*. Wiley, 2008.

[135] H. Paik, D. I. Schuster, L. S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. R. Johnson, M. J. Reagor, L. Frunzio, L. I. Glazman, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf. Observation of high coherence in josephson junction qubits measured in a three-dimensional circuit qed architecture. *Phys. Rev. Lett.*, 107:240501, Dec 2011.

[136] D. P. Pappas, M. R. Vissers, D. S. Wisbey, J. S. Kline, and J. Gao. Two level system loss in superconducting microwave resonators. *IEEE Transactions on Applied Superconductivity*, 21(3):871–874, 2011.

[137] M. Pechal, L. Huthmacher, C. Eichler, S. Zeytinoğlu, A. A. Abdumalikov Jr, S. Berger, A. Wallraff, and S. Filipp. Microwave-controlled generation of shaped single photons in circuit quantum electrodynamics. *Physical Review X*, 4(4):041010, 2014.

[138] W. A. Phillips. Tunneling states in amorphous solids. *Journal of Low Temperature Physics*, 7(3):351–360, May 1972.

[139] W. A. Phillips. Two-level states in glasses. *Reports on Progress in Physics*, 50(12):1657–1708, dec 1987.

[140] A. P. M. Place, L. V. H. Rodgers, P. Mundada, B. M. Smitham, M. Fitzpatrick, Z. Leng, A. Premkumar, J. Bryon, A. Vrajitoarea, S. Sussman, G. Cheng, T. Madhavan, H. K. Babla, X. H. Le, Y. Gang, B. Jäck, A. Gyenis, N. Yao, R. J. Cava, N. P. de Leon, and A. A. Houck. New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds. *Nature Communications*, 12(1):1779, Mar 2021.

[141] . Pobell, Frank. *Matter and methods at low temperatures*. Springer, Berlin ; New York, array edition, 2007.

[142] H. Podlech. Superconducting versus normal conducting cavities. *arXiv preprint arXiv:1303.6552*, 2013.

[143] C. Powell and A. Jabolonski. Nist standard reference database 82 - nist electron effective-attenuation-length database, version 1.3. (Version 1.3), 2011-01-01 2011.

[144] D. M. Pozar. *Microwave engineering*. John Wiley & Sons, 2009.

[145] A. Premkumar, C. Weiland, S. Hwang, B. Jäck, A. P. M. Place, I. Waluyo, A. Hunt, V. Bisogni, J. Pelliciari, A. Barbour, M. S. Miller, P. Russo, F. Camino, K. Kisslinger, X. Tong, M. S. Hybertsen, A. A. Houck, and I. Jarrige. Microscopic relaxation channels in materials for superconducting qubits. *Communications Materials*, 2(1):72, July 2021.

[146] A. P. Read, B. J. Chapman, C. U. Lei, J. C. Curtis, S. Ganjam, L. Krayzman, L. Frunzio, and R. J. Schoelkopf. Precision measurement of the microwave dielectric loss of sapphire in the quantum regime with parts-per-billion sensitivity, 2022.

[147] M. Reagor, H. Paik, G. Catelani, L. Sun, C. Axline, E. Holland, I. M. Pop, N. A. Masluk, T. Brecht, L. Frunzio, M. H. Devoret, L. Glazman, and R. J. Schoelkopf. Reaching 10 ms single photon lifetimes for superconducting aluminum cavities. *App. Phys. Lett.*, 102(19):192604, 2013.

[148] M. Reagor, W. Pfaff, C. Axline, R. W. Heeres, N. Ofek, K. Sliwa, E. Holland, C. Wang, J. Blumoff, K. Chou, et al. Quantum memory with millisecond coherence in circuit qed. *Phys. Rev. B*, 94(1):014506, 2016.

[149] M. Reagor, W. Pfaff, C. Axline, R. W. Heeres, N. Ofek, K. Sliwa, E. Holland, C. Wang, J. Blumoff, K. Chou, M. J. Hatridge, L. Frunzio, M. H. Devoret, L. Jiang, and R. J. Schoelkopf. A quantum memory with near-millisecond coherence in circuit QED, 2015.

[150] M. J. Reagor. *Superconducting Cavities for Circuit Quantum Electrodynamics*. PhD thesis, Yale University, 2016.

[151] M. D. Reed, L. DiCarlo, B. R. Johnson, L. Sun, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf. High-fidelity readout in circuit quantum electrodynamics using the Jaynes-Cummings nonlinearity. *Phys. Rev. Lett.*, 105:173601, Oct 2010.

[152] P. Reinhold, S. Rosenblum, W.-L. Ma, L. Frunzio, L. Jiang, and R. J. Schoelkopf. Error-corrected gates on an encoded qubit. *Nature Physics*, pages 1–5, 2020.

[153] C. Rigetti, J. M. Gambetta, S. Poletto, B. L. T. Plourde, J. M. Chow, A. D. Córcoles, J. A. Smolin, S. T. Merkel, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen. Superconducting qubit in a waveguide cavity with a coherence time approaching 0.1 ms. *Phys. Rev. B*, 86:100506, Sep 2012.

[154] A. Romanenko, A. Grassellino, A. C. Crawford, D. A. Sergatskov, and O. Melnychuk. Ultra-high quality factors in superconducting niobium cavities in ambient magnetic fields up to 190 mg. *Applied Physics Letters*, 105(23):234103, 2014.

[155] A. Romanenko, R. Pilipenko, S. Zorzetti, D. Frolov, M. Awida, S. Belomestnykh, S. Posen, and A. Grassellino. Three-dimensional superconducting resonators at t¡ 20 mk with photon lifetimes up to τ= 2 s. *Physical Review Applied*, 13:034032, Mar 2020.

[156] A. Romanenko, R. Pilipenko, S. Zorzetti, D. Frolov, M. Awida, S. Belomestnykh, S. Posen, and A. Grassellino. Three-dimensional superconducting resonators at $tlt; 20$ mk with photon lifetimes up to $\tau = 2$ s. *Phys. Rev. Applied*, 13:034032, Mar 2020.

[157] A. Romanenko and D. I. Schuster. Understanding quality factor degradation in superconducting niobium cavities at low microwave field amplitudes. *Phys. Rev. Lett.*, 119:264801, Dec 2017.

[158] S. Rosenblum, Y. Y. Gao, P. Reinhold, C. Wang, C. J. Axline, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, , and R. Schoelkopf. A cnot gate between multiphoton qubits encoded in two cavities. *Nature communications*, 9(1):1–6, 2018.

[159] S. Rosenblum, P. Reinhold, M. Mirrahimi, L. Jiang, L. Frunzio, and R. J. Schoelkopf. Fault-tolerant detection of a quantum error. *Science*, 361(6399):266–270, 2018.

[160] K. J. Satzinger, Y. P. Zhong, H.-S. Chang, G. A. Peairs, A. Bienfait, M.-H. Chou, A. Y. Cleland, C. R. Conner, É. Dumur, J. Grebel, I. Gutierrez, B. H. November, R. G. Povey, S. J. Whiteley, D. D. Awschalom, D. I. Schuster, and A. N. Cleland. Quantum control of surface acoustic-wave phonons. *Nature*, 563(7733):661–665, Nov 2018.

[161] R. T. Schermer and T. H. Stievater. Millimeter-wave dielectric properties of highly refractive single crystals characterized by waveguide cavity resonance. *IEEE Transactions on Microwave Theory and Techniques*, 67(3):1078–1087, 2019.

[162] F. Schmid and D. C. Harris. Effects of crystal orientation and temperature on the strength of sapphire. *Journal of the American Ceramic Society*, 81(4):885–893, 1998.

[163] D. Schuster, A. Houck, J. Schreier, A. Wallraff, J. Gambetta, A. Blais, L. Frunzio, J. Majer, B. Johnson, M. Devoret, S. Girvin, and R. Schoelkopf. Resolving photon number states in a superconducting circuit. *Nature*, 445(7127):515–518, 2007.

[164] D. I. Schuster. *Circuit quantum electrodynamics*. PhD thesis, 2007. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2021-09-28.

[165] Q. S. Shu, J. A. Demko, and J. E. Fesmire. Heat switch technology for cryogenic thermal management. *IOP Conference Series: Materials Science and Engineering*, 278:012133, dec 2017.

[166] W. Singer. Fabrication of elliptical SRF cavities. *Superconductor Science and Technology*, 30(3):033001, jan 2017.

[167] W. Singer, S. Aderhold, A. Ermakov, J. Iversen, D. Kostin, G. Kreps, A. Matheisen, W.-D. Möller, D. Reschke, X. Singer, K. Twarowski, H. Weise, and H.-G. Brokmeier. Development of large grain cavities. *Phys. Rev. ST Accel. Beams*, 16:012003, Jan 2013.

[168] K. Soda, H. Kondo, M. Kato, T. Shiraki, K. Niwa, K. Kusaba, M. Hasegawa, and E. Ikenaga. Characterization of nb hydrides synthesized in high-pressure supercritical water by micro-beam hard x-ray photoelectron spectroscopy. *Journal of Electron Spectroscopy and Related Phenomena*, 186:54–57, 2013.

[169] A. Somoroff, Q. Ficheux, R. A. Mencia, H. Xiong, R. V. Kuzmin, and V. E. Manucharyan. Millisecond coherence in a superconducting qubit. 2021.

[170] S. J. Srinivasan, A. J. Hoffman, J. M. Gambetta, and A. A. Houck. Tunable coupling in circuit quantum electrodynamics using a superconducting charge qubit with a v-shaped energy level diagram. *Phys. Rev. Lett.*, 106:083601, Feb 2011.

[171] A. M. Steane. Efficient fault-tolerant quantum computing. *Nature*, 399(6732):124–126, May 1999.

[172] M. Steer. *Fundamentals of Microwave and RF Design*. NC State University, 2019.

[173] L. Stefanazzi, K. Treptow, N. Wilcer, C. Stoughton, S. Montella, C. Bradford, G. Cancelo, S. Saxena, H. Arnaldi, S. Sussman, A. Houck, A. Agrawal, H. Zhang, C. Ding, and D. I. Schuster. The qick (quantum instrumentation control kit): Readout and control for qubits and detectors. 2021.

[174] Y. Tereshkin, C. Boffo, G. Davis, H. Edwards, and A. Rowe. BCP processing facility. In *11th Workshop on RF Superconductivity (SRF 2003)*, 10 2003.

[175] M. Tinkham. *Introduction to superconductivity*. Courier Corporation, 2004.

[176] Y. Uzel, K. Schnitzke, and N. Krause. Improvement of the polishing treatment for niobium surfaces of superconducting cavity resonators. *Applied Physics A*, 30(3):185–187, Mar 1983.

[177] I. Vázquez, M. P. Russell, D. R. Smith, and R. Radebaugh. *Helium Adsorption on Activated Carbons at Temperatures between 4 and 76 K*, pages 1013–1021. Springer US, Boston, MA, 1988.

[178] J. C. Vickerman and I. S. Gilmore. *Surface analysis: the principal techniques*. John Wiley & Sons, 2011.

[179] L. Viola, E. Knill, and S. Lloyd. Dynamical decoupling of open quantum systems. *Phys. Rev. Lett.*, 82:2417–2421, Mar 1999.

[180] L. Viola and S. Lloyd. Dynamical suppression of decoherence in two-state quantum systems. *Phys. Rev. A*, 58:2733–2744, Oct 1998.

[181] L. Viola, S. Lloyd, and E. Knill. Universal control of decoupled quantum systems. *Phys. Rev. Lett.*, 83:4888–4891, Dec 1999.

[182] V. P. Vlasov, A. E. Muslimov, A. V. Butashin, and V. M. Kanevsky. Sapphire evolution of the vicinal (0001) sapphire surface upon annealing in air. *Crystallography Reports*, 61(1):58–62, Jan 2016.

[183] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf. Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics. *Nature*, 431(7005):162–167, 2004.

[184] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Approaching unit visibility for control of a superconducting qubit with dispersive readout. *Phys. Rev. Lett.*, 95:060501, Aug 2005.

[185] D. F. Walls and G. J. Milburn. *Quantum optics*. Springer Science & Business Media, 2007.

[186] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potočnik, Y. Salathé, M. Pechal, M. Mondal, M. Oppliger, C. Eichler, and A. Wallraff. Rapid high-fidelity single-shot dispersive readout of superconducting qubits. *Phys. Rev. Applied*, 7:054020, May 2017.

[187] C. Wang, Y. Y. Gao, P. Reinhold, R. W. Heeres, N. Ofek, K. Chou, C. Axline, M. Reagor, J. Blumoff, K. Sliwa, et al. A schrödinger cat living in two boxes. *Science*, 352(6289):1087–1091, 2016.

[188] C. Wang, X. Li, H. Xu, Z. Li, J. Wang, Z. Yang, Z. Mi, X. Liang, T. Su, C. Yang, G. Wang, W. Wang, Y. Li, M. Chen, C. Li, K. Linghu, J. Han, Y. Zhang, Y. Feng, Y. Song, T. Ma, J. Zhang, R. Wang, P. Zhao, W. Liu, G. Xue, Y. Jin, and H. Yu.

Towards practical quantum computers: transmon qubit with a lifetime approaching 0.5 milliseconds. *npj Quantum Information*, 8(1):3, Jan 2022.

[189] C. S. Wang, J. C. Curtis, B. J. Lester, Y. Zhang, Y. Y. Gao, J. Freeze, V. S. Batista, P. H. Vaccaro, I. L. Chuang, L. Frunzio, L. Jiang, S. M. Girvin, and R. J. Schoelkopf. Efficient multiphoton sampling of molecular vibronic spectra on a superconducting bosonic processor. *Phys. Rev. X*, 10:021060, Jun 2020.

[190] H. Wang, M. Mariantoni, R. C. Bialczak, M. Lenander, E. Lucero, M. Neeley, A. D. O'Connell, D. Sank, M. Weides, J. Wenner, T. Yamamoto, Y. Yin, J. Zhao, J. M. Martinis, and A. N. Cleland. Deterministic entanglement of photons in two superconducting microwave resonators. *Phys. Rev. Lett.*, 106:060401, Feb 2011.

[191] Y. Wang, M. Um, J. Zhang, S. An, M. Lyu, J.-N. Zhang, L.-M. Duan, D. Yum, and K. Kim. Single-qubit quantum memory exceeding ten-minute coherence time. *Nature Photonics*, 11(10):646–650, Oct 2017.

[192] W. C. Young, R. G. Budynas, and A. M. Sadegh. *Roark's formulas for stress and strain; 8th ed.* McGraw Hill, New York, NY, 2012.

[193] E. Zakka-Bajjani, F. Nguyen, M. Lee, L. R. Vale, R. W. Simmonds, and J. Aumentado. Quantum superposition of a single microwave photon in two different 'colour' states. *Nat. Phys.*, 7(8):599–603, Aug 2011.

[194] B. Zhou and W. F. Ramirez. Kinetics and modeling of wet etching of aluminum oxide by warm phosphoric acid. *Journal of The Electrochemical Society*, 143(2):619–623, feb 1996.

[195] N. Zhu, J. L. Orlowski-Scherer, Z. Xu, A. Ali, K. S. Arnold, P. C. Ashton, G. Coppi, M. J. Devlin, S. Dicker, N. Galitzki, P. A. Gallardo, S. W. Henderson, S.-P. P. Ho, J. Hubmayr, B. Keating, A. T. Lee, M. Limon, M. Lungu, P. D. Mauskopf, A. J. May, J. McMahon, M. D. Niemack, L. Piccirillo, G. Puglisi, M. S. Rao, M. Salatino, M. Silva-Feaver, S. M. Simon, S. Staggs, R. Thornton, J. N. Ullom, E. M. Vavagiakis, B. Westbrook, and E. J. Wollack. Simons observatory large aperture telescope receiver design overview, 2018.