

**Enhanced Cooperativity in a Near-Quantum
Microwave-to-Optical Transducer**

by

Sarang Mittal

B.S., California Institute of Technology, 2018

M.S., University of Colorado, Boulder, 2021

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Physics
2024

Committee Members:

Konrad W. Lehnert, Chair

Cindy A. Regal

John D. Teufel

Shuo Sun

Taylor W. Barton

Mittal, Sarang (Ph.D., Physics)

Enhanced Cooperativity in a Near-Quantum Microwave-to-Optical Transducer

Thesis directed by Professor Konrad W. Lehnert

One of the main thrusts of quantum science over the past few decades has been the development of quantum networks for the purposes of secure communication, enhanced detector sensitivity, and advanced computing. Realizing such a network of superconducting quantum processors that communicate via optical fibers would leverage the high fidelity quantum signal processing of superconducting circuits and the thermal robustness of infrared light but requires a transducer capable of connecting these two sections of the electromagnetic spectrum separated by five orders of magnitude in energy. This thesis explores the optimization of a transducer architecture where the mechanical mode of a Si_3N_4 membrane mediates the coupling of a superconducting lumped-element circuit and a Fabry-Perot optical cavity. We aim to maximize the coupling of these three harmonic oscillator to each other while shielding them from noisy processes that would decohere quantum signals. This architecture has led to transducers with unparalleled efficiency and continuous operation. Enhanced cooperativity between the optical cavity and mechanical oscillator enabled optically-detected readout of a superconducting qubit and optomechanical ground state cooling with negligible laser-induced heating of the superconducting qubit and microwave circuit. To surpass the threshold for quantum-enabled operation, we subsequently improved the cooperativity between the microwave circuit and mechanical oscillator by reducing the microwave loss and noise from two-level-system-like defects in the Si_3N_4 dielectric. When combined with enhanced coupling between the circuit and membrane or improved mechanical isolation, we project that this architecture will be capable of transducing quantum signals between the microwave and optical regimes with a signal-to-noise greater than one.

Dedication

To the past that shaped us, the present that nurtures us, and the future that awaits us

Acknowledgements

I really held off on writing this until the last minute, because the thought of thanking every person who contributed in some way to my graduate student experience feels more impossible than writing this thesis in the first place. Here goes nothing.

Thank you to my advisor, Konrad Lehnert, for taking me under your wing. I learned an immense amount of physics and also your detailed-oriented, no-stone-left-unturned attitude towards problem-solving. I'm grateful to have been your student, and most importantly, your friend and colleague. In that same vein, this project would not have been possible without the advice and insight of Cindy Regal. Your ability to see right to the heart of the problem, while ignoring all of the fluff, is something I aspire to develop for myself.

I've also had the privilege to work with some amazing post-docs and graduate students during my time. Peter Burns was my first friend in Boulder, and showed me how to do science with joy and curiosity. Benjamin Brubaker and Jon Kindem taught me how to properly document my work and analyze data with a fine-toothed comb. Maxwell Urney helped me develop an intuitive understanding for the experiment and became someone I look to for technical and life advice. Rob Delaney's can-do attitude and expansive knowledge of cryogenics were resources I often leaned on in my first few years. We all suffered through the COVID pandemic together while still producing groundbreaking research, for which I will always be grateful.

In the later years of my PhD, I was very fortunate to work with new faces, who brought their own valuable insights. Nicholas Frattini showed me how to be an actual microwave engineer and helped develop a lab culture that I will sorely miss. Kazemi Adachi is a wizard with plots and color

scales and helped me execute many complicated ideas. Luca Talamo, Sarah Dickson, Sheng-Xiang Lin, and Max Olberding have brought fresh perspectives and I'm sure the experiment is in good hands.

Wow, how is this supposed to be one page? Much of my time was spent at the NIST BMF cleanroom, developing and perfecting fabrication skills under the guidance of the cleanroom staff. Peter Lowell, John Nibarger, Jim Beall, Mike Vissers, Kat Cicak, Dan Schmidt, and Susan Schima have a major hand in the realizing the difficult fabrication of the transducers. I have also really enjoyed working with Tasshi Dennis, Jake Davidson, and Aly Artusio-Glimpse at NIST who are working to extend many of the ideas in this thesis. Reevaluating my long-held assumptions about how things should work or need to work pushed the boundaries of my knowledge and the boundaries of the experiment.

The rest of Konrad's group have been excellent sounding boards for ideas and confusions over the years. In particular, Alec Emser and Cyril Metzger were integral in understanding the TLS model in chapter 5. I'd also like to thank Eric, Maxime, Liz, Joyce, Kyle, Ali, Ziyi, Eva, Brendon, Pablo, Brad, Dan, Lucas, Caesar, and Kelly. Around JILA, the machine shop, electronics shop, and computing have also made many contributions, all with tremendous patience when I didn't really know what problem I had. Cindy's other optomechanics experiment were phenomenal resources for questions about mechanical membranes, in particular Chris and Ravid.

Running out of real estate. My parents, Riju and Shuchi. My sister, Somya. My cousin, Shivan, and the rest of my family. My girlfriend, Rachel, and the loyal, thoughtful friends I've made in Colorado. Whether that be my weekly Dungeons and Dragons group, Colorado Mamabird, Inception, or ISO Atmo. My roommates Anish and Lange. Thank you all for keeping my sane outside of the lab and letting me be a part of your lives. My experience is defined as much by those relationships as it is by my research.

Lastly, a huge shoutout to my therapists, who helped me learn to accept myself and find peace amongst the ups and downs of this PhD.

Contents

Chapter

1	Introduction	1
1.1	Platforms for microwave-to-optical transduction	2
1.2	Mechanically-mediated doubly-parametric transducer	3
1.3	Thesis overview	4
2	Transduction and Optomechanical Theory	7
2.1	Metrics for transducer performance	7
2.2	Cavity opto-/electro-mechanics for efficient transduction	10
2.2.1	Hamiltonian formulation	10
2.2.2	Input-output theory for cavity electromechnics	12
2.2.3	Transducer scattering parameters	14
2.2.4	Measuring transducer efficiency	17
2.3	Opto/electromechanical sideband cooling	18
2.3.1	Membrane thermometry	21
2.4	Transduction added noise	25
2.4.1	Microwave cavity occupation from phase/amplitude noise	27
2.4.2	Transducer added noise from phase/amplitude noise	28
2.4.3	Corrections to membrane thermometry from phase/amplitude noise	30
2.4.4	Frequency fluctuations and heating of the microwave circuit	32
2.5	Understanding the parameter space	33

3	Membrane mode engineering	42
3.1	Phononic crystal	43
3.1.1	Infinite PNC design	45
3.1.2	Finite PNC effects	47
3.1.3	PNC performance	49
3.2	Oxide-Oxide bonding	52
3.2.1	Circuit chip-mirror chip bonding	54
3.2.2	Membrane chip-circuit chip bonding	56
3.3	Pinned membrane	58
3.3.1	Pinning post design	59
3.3.2	Experimental realization	62
4	Transducer Demonstrations and Limits	66
4.1	Optically-detected superconducting qubit readout	70
4.2	Optomechanical ground-state cooling	73
4.3	Investigating microwave loss and noise	79
4.3.1	Trapped flux	80
4.3.2	Quasiparticle loss	84
4.3.3	NbTiN surface and bulk analysis	87
4.3.4	Identifying Si_3N_4 -related dielectric loss	89
5	Microwave properties of Si_3N_4	93
5.1	Identifying bulk Si_3N_4	94
5.2	Generalized model of circuit loss and frequency shift	97
5.2.1	Susceptibility of a single TLS	100
5.2.2	Resonant interaction	102
5.2.3	Relaxation interaction	105
5.2.4	Extracting material properties	111

5.2.5	Self-heating	113
5.3	Annealing reduces microwave-frequency dielectric loss	118
5.3.1	Tensile stress of annealed Si_3N_4	121
5.4	Microscopic origin of TLS	123
5.5	TLS energy distribution	126
5.6	TLS description of $\gamma_m(T)$	128
6	Characterization of microwave noise	130
6.1	Noise calibrations	133
6.1.1	Generator phase noise	133
6.1.2	Coherent pump cancellation	134
6.2	Impact of annealing on microwave noise	138
6.3	Modeling internal cavity noise	140
6.4	Upgrading microwave architecture	147
6.4.1	Pump and signal mode isolation	147
6.4.2	Incorporating near-quantum limited amplifier	150
	References	153
 Appendix		
A	Transducer Fabrication	167
A.1	Circuit and mirror wafer	167
A.2	Membrane wafer	174
A.3	Transducer construction	182
A.3.1	Glued construction	182
A.3.2	Bonded construction	184
A.3.3	Determining gap separation	186

B	Numerical Fitting Techniques and Results	189
B.1	Resonator characterization	189
B.2	Tables of fit parameters	191

Tables

Table

2.1	Example transducer parameters	37
3.1	Pinned membrane mode frequencies	63
4.1	Electro-optic transducer parameters	69
4.2	Contributions to quantum efficiency of qubit readout	74
4.3	Contributions to minimum added noise in upconversion	78
4.4	Surface composition from XPS	89
4.5	RMS Roughness from AFM scans	90
5.1	Electric field energy participation	95
5.2	Annealing reduces tensile stress of Si_3N_4	121
B.1	Fitted resonator parameters	192
B.2	TLS model parameters	194

Figures

Figure

2.1	Generalized transducer model	8
2.2	Doubly-parametric mechaniacly-mediated microwave-to-optical transducer	11
2.3	Efficiency calibration	17
2.4	Electromechanical sideband cooling	19
2.5	Measurement efficiency calibration	23
2.6	Sideband asymmetry thermometry	24
2.7	Transduction performance versus microwave loss	38
2.8	Transduction performance versus mechanical loss	39
2.9	Transduction performance versus electromechanical coupling	40
2.10	Transduction performance versus microwave noise	41
3.1	Evolution of membrane chip	43
3.2	Spectator modes of Si frame	46
3.3	PNC unit cell	47
3.4	Improved Mechanical Properties	50
3.5	Oxide-Oxide bonding for transducer construction	53
3.6	Stress compensation of mirror coating	55
3.7	Bonded transducer etalons	57
3.8	Pinning post design	61

3.9	Electromechanical properties of pinned membrane	64
3.10	Duffing nonlinearity of some pinned membranes	64
4.1	Transducer description	68
4.2	Single-shot optically-detected qubit readout	71
4.3	Transducer backaction on qubit	72
4.4	Optomechanical ground state cooling	75
4.5	Minimum added noise sweep	77
4.6	Effects of laser on microwave circuit	78
4.7	Pump-induced microwave noise and loss	79
4.8	Trapped flux in LC circuit	81
4.9	NbTiN T_c and Thermal Quasiparticles	85
4.10	Depth profile of $\text{Nb}_x\text{Ti}_{1-x}\text{N}$ composition	88
4.11	Inductive and Capacitive Si_3N_4	91
4.12	Measurements of inductive and capacitive Si_3N_4	92
5.1	Power-dependent loss comparison	96
5.2	Loss versus F_{SiN}	98
5.3	TLS Hamiltonian	99
5.4	TLS energies sampled by resonant interaction	105
5.5	TLS energies sampled by relaxation interaction	110
5.6	Power- and temperature-dependent properties of as-deposited Si_3N_4	112
5.7	Evidence of self-heating	114
5.8	TLS and self-heating model	115
5.9	Power- and temperature-dependent properties of annealed Si_3N_4	119
5.10	TLS and self-heating in transducer improved by annealing	122
5.11	Infrared spectroscopy of Si_3N_4	125
5.12	Multimode CPW resonator	129

6.1	Total circuit occupation in electromechanical experiments	131
6.2	Microwave filter cavity	134
6.3	Pump cancellation schematic	135
6.4	Measured phase noise correction	137
6.5	Internal cavity noise in planar resonators	141
6.6	Internal cavity noise in transducers	141
6.7	Self-heating based noise model	145
6.8	Annealing reduces pump-induced internal circuit noise	147
6.9	Two-mode circuit design	148
6.10	KIT setup and gain	152
A.1	Circuit and mirror wafer fabrication flow	168
A.2	Membrane wafer fabrication flow	175
A.3	Trion Fl etch non-uniformity	176
A.4	KOH etch angle	177
A.5	Filling membrane hole with resist	180
A.6	DRIE damage	181
A.7	Transducer construction	183
A.8	Example of epoxied transducer	184
A.9	One-dimensional model of chip bowing	185
A.10	Applying force for circuit-membrane chip bonding	186
A.11	Capacitor gap spacing from color	188
B.1	Example of numerical fit	190

Chapter 1

Introduction

*It's a dangerous business, Frodo, going out your door. You step onto the road,
and if you don't keep your feet, there's no knowing where you might be swept off
to.*

- J.R.R. Tolkein

In the summer of 1880, Alexander Graham Bell received a voice message carried on a beam of light from his assistant, Charles Sumner Tainter, over a distance of about two football fields [Bell 1880; Polishuk 1981]. Sound waves from Tainter's voice vibrated a mirror, which was reflecting sunlight to a selenium cell at the receiver. Modulation of the intensity of the sunlight incident on the receiver modulated the current through the photoconductive selenium cell, which was transduced back to sound with an earphone. Using their newly invented photophone, the two demonstrated a primitive version of modern optical communication technology. Nearly 150 years later, one can find traces of this hybrid optical, electrical, and acoustic system in efforts to tackle a modern communication problem: long-range quantum networks of superconducting circuits.

Superconducting circuits are a mature platform for quantum computing, demonstrating long coherence times [Somoroff et al. 2023; Ganjam et al. 2024], high fidelity operations [Ding et al. 2023; Moskalenko et al. 2022], and preparation of exotic states with error protection [Grimm et al. 2020; Sivak et al. 2023]. These systems must be operated at milliKelvin temperatures in a dilution refrigerator in order to suppress thermal decoherence at the microwave operating frequencies ($h \times 6 \text{ GHz}/k_B \sim 300 \text{ mK}$). While small scale networks based on cryogenic transmission have been demonstrated [Storz et al. 2023], the low temperature requirement poses a fundamental challenge

in building long-range networks of superconducting circuits. Room temperature transmission will decohere the quantum state and cryogenic systems are unlikely to grow to kilometer-length scales.

To surmount these challenges, signal propagation at optical frequencies over fiber is a popular modern proposal. Optical fields do not suffer the same thermal decoherence at room temperature ($\hbar \times 200 \text{ THz}/k_B \sim 10^4 \text{ K}$) and modern fiber technology is remarkably low-loss and cost-efficient. In fact, metropolitan-scale all-optical quantum networks have already been demonstrated [Dynes et al. 2019; Chen et al. 2021]. Establishing a quantum link between superconducting circuits mediated by optical fiber would enable applications in secure communications [Bennett and Brassard 1984], enhancing telescope sensitivity [Gottesman et al. 2012], and distributed quantum computing [Bennett, Brassard, et al. 1993; L. Jiang et al. 2007].

A transducer capable of converting quantum signals between microwave and optical frequencies is a requirement for this proposed network architecture. Such a device would need to bridge the five order of magnitude energy difference between these sections of the electromagnetic spectrum, while preserving the fragile quantum states. The additional energy is provided by an optical pump tone, introducing a significant technical challenge. A single optical-frequency photon incident on the microwave circuit will destroy the superconductivity. Often thousands to millions of optical photons are required to optimize the interaction rates in these transducers, further exacerbating the problem. Integrating the microwave and optical elements into one device is thus the key hurdle in realizing a quantum microwave-to-optical transducer.

1.1 Platforms for microwave-to-optical transduction

In recent years, there has been enormous progress on a number of microwave-to-optical transduction experiments, leveraging a wide variety of physical phenomena. Quantum-enabled conversion between millimeter wave and optical photons using Rydberg states of neutral atoms has been achieved [Kumar et al. 2023], and there are experiments to extend this down to the microwave regime [Verdú et al. 2009; Vogt et al. 2019; Smith et al. 2023; Borówka et al. 2024]. Other experiments leverage the energy levels of solid state emitters such as nitrogen-vacancy centers [Imamoğlu 2009;

Marcos et al. 2010] or rare-earth ions [Williamson et al. 2014; Bartholomew et al. 2020; Rochman et al. 2023]. Early work with magnons suggests another potential pathway [Hisatomi et al. 2016; Zhu et al. 2020].

One can also forgo the use of an intermediate mode and utilize the electro-optic (Pockels) effect of lithium niobate [C. Wang et al. 2022]. Implementations using bulk crystals in superconducting three-dimensional cavities have achieved quantum-enabled performance [Sahu et al. 2022; Qiu et al. 2023], and work is progressing rapidly on more compact thin-film lithium niobate implementations [Holzgrafe et al. 2020; M. Xu et al. 2024].

One of the more successful approaches is to utilize an mechanical oscillator as the mediator of transduction. In this platform, microwave circuits can be resonantly coupled to gigahertz-frequency (GHz) mechanical oscillators, such as bulk acoustic wave resonators [Doeleman et al. 2023] or silicon optomechanical crystals [W. Jiang, Mayor, et al. 2023; Weaver et al. 2024], using piezoelectricity or the electromechanical interaction. The mechanical oscillator is also parametrically-coupled to the optical cavity via radiation pressure. The most advanced of these experiments have entered the quantum-regime [Mirhosseini et al. 2020; Zhao et al. 2023] and have observed non-classical correlation between microwave and optical fields [Meesala et al. 2024].

In many of these implementations, optical heating of the transducer via absorption in the transduction medium or scattering onto the superconductor adds significant noise. This necessitates pulsing or lowering the power of the optical pump, such as to reduce the average thermal load. Lower noise operation is achieved at the expense of reduced efficiency or eventual entangled bit generation rates. A transducer architecture more robust to optical illumination would ideally avoid this trade-off of interaction rate and noise.

1.2 Mechanically-mediated doubly-parametric transducer

The transduction platform explored in my research, and the focus of this thesis, achieves this robustness to optical light by engineering the optical mode to live predominantly in vacuum and physically separating the microwave and optical resonators. This avoids absorption in a medium and

limits scattering of infrared photons onto the superconductor. To realize this, a megahertz-frequency (MHz) mechanical mode of a tensioned silicon nitride (Si_3N_4) membrane is simultaneously coupled to a superconducting circuit and a high-finesse optical cavity [Regal and Lehnert 2011; R. W. Andrews et al. 2014; Planz et al. 2022]. The low-frequency mechanical oscillator is relatively large compared to the GHz-frequency mechanical oscillators, allowing the microwave and optical resonators to both couple to the mechanics while being separated by many times the Gaussian beam width. In combination with excellent optical mode control due to robust cavity design, photons from the optical mode are unlikely to disturb the superconducting circuit.

As the microwave and mechanical resonators are not resonant (as in the experiments leveraging piezoelectricity described earlier), an additional pump at microwave-frequencies is required to mediate the electromechanical interaction. In contrast to other transducer implementations, noise and heating from electromechanical pump has, up to now, been the dominant limiting factor in the transducers performance. This noise often manifests as heating of the microwave circuit, a phenomena observed in many superconducting electromechanical experiments over the past two decades [Suh 2011; Teufel et al. 2011; Massel et al. 2012; Wollman 2015; Lei 2017]. As we will explore in Sec. 2.3, the microwave pump also cools the mechanical mode by coupling to a low-temperature microwave-frequency bath. Heating of the microwave circuit will serve as an ultimate limit for the electromechanical cooling, as the membrane cannot get colder than the bath to which it is coupled.

A significant part of this thesis is dedicated to understanding, modeling, and reducing this source of noise. As a result of these efforts, the quantum operation of a doubly-parametric mechanically-mediated microwave-to-optical transducer is within sight, opening the door for long-distance networking demonstrations with superconducting circuits.

1.3 Thesis overview

This thesis follows the significant progress made to improve the performance of the JILA transducer during my tenure. This was a large team effort, and throughout, I will focus on my personal contributions while referring to the theses of my colleagues when appropriate.

In [Ch. 2](#), I introduce the standard metrics for transducer performance, as well as the cavity optomechanical theory that underlies the operation of our transducer. I will show that in addition to mediating the transduction, the optical and microwave pumps cool the mechanical mode towards its ground state, but can also add technical noise. Balancing these contributions is crucial to minimizing the transduction noise. When I joined the experiment, the team had just demonstrated a transducer with exceptionally low noise and high-efficiency for the time [[Higginbotham et al. 2018](#)], but still limited by tens of photons of added noise, enough to completely drown out any single-quantum signal. [Ch. 3](#) will cover some of the subsequent major improvements made to the device at the beginning of my tenure. A phononic crystal reduced the mechanical resonators coupling to the hot, low-frequency thermal bath, and an improved optical cavity assembly reduced noise effects from the microwave and optical pumps. In addition, I will introduce a novel technique for achieving large electromechanical coupling rates, which can be implemented in future transducers.

In [Ch. 4](#), I will review two experiments that showcased the continuous, low-noise operation of our transducer. First, we demonstrated optically-detected single-shot readout of a superconducting qubit by upconverting a classical readout pulse with minimal backaction on the superconducting qubit [[R. D. Delaney, Urney, et al. 2022](#)]. This demonstrated the compatibility of sensitive quantum circuits with microwave-to-optical transduction. Second, reduced noise and improved optomechanical coupling enabled optomechanical ground-state cooling of the mechanical oscillator and transduction noise at the few photon level [[Brubaker et al. 2022](#)]. The limiting factor in both of these experiments was microwave-pump-induced loss and noise in the superconducting resonator of the transducer. The majority of this chapter is dedicated to tracking down the physical mechanism of this effect, a two-year process that involved many dead ends.

[Ch. 5](#) unmask the mysterious physical mechanism of the microwave-pump-induced loss: dielectric loss in the Si_3N_4 membrane governed by the two-level-systems (TLS) model [[Mittal et al. 2024](#)]. This result was initially surprising as it clashed with the common understanding of dielectric loss from so-called "resonant" damping saturation. We were able to show that the increased circuit loss can be understood within the TLS model using the "relaxation" damping term, in concert

with a simple self-heating assumption. Most importantly, post-deposition annealing of the Si_3N_4 reduced both damping terms, especially the relaxation damping. This chapter will also discuss the implications our observations could have for the details of the TLS model.

With a model for the pump-induced loss in hand, we turn to understanding the noise of the microwave circuit in [Ch. 6](#). There are two contributions to this effect: internally generated noise in the circuit and technical noise on the microwave pump. We find that annealing also reduces the microwave-pump power dependent noise generated inside of the circuit, after correctly accounting for the contribution from generator technical noise to the measured noise spectrum. We also review strategies for mitigating the effect of technical noise on the microwave pump. In the appendices, I outline the details of transducer fabrication, which was carried out at the NIST Boulder Microfabrication (BMF) cleanroom, as well as computer-aided techniques for resonator characterization.

Chapter 2

Transduction and Optomechanical Theory

It's still magic even if you know how it's done
- Terry Pratchett

In this chapter, I will discuss the theoretical underpinnings for the operation and optimization of the transducer. In [Sec. 2.1](#), I'll introduce the key performance metrics for a microwave-to-optical transducer and [Sec. 2.2](#) will outline the operating principle for the doubly-parametric mechanically-mediated transducer. In [Sec. 2.3](#), I'll show how the strong pump tones cool the mechanical mode and outline the techniques for extracting the thermal population. Lastly, in [Sec. 2.4](#), I'll describe how technical noise on the microwave and optical pumps can limit cooling and complicate membrane thermometry. I'll also explore how asymmetry in the noise affects upconversion and downconversion.

2.1 Metrics for transducer performance

There are four key metrics for a transducer: efficiency, added noise, bandwidth, and duty cycle [[Zeuthen et al. 2020](#)]. To understand these quantities, let's consider a generalized model of a transducer as a beamsplitter, as shown in [Fig. 2.1](#). Microwave input \hat{b}_{in} and output \hat{b}_{out} modes are on the left and top ports, while optical input \hat{a}_{in} and output \hat{a}_{out} modes are on the bottom and right ports. All transducers require at least one pump, and a general transducer could require many tones at frequencies $\omega_{\text{p},i}$ to mediate the interaction. Finally, noise is represented by a stochastic force $\hat{\mathcal{F}}$ that accompanies the signals on the output ports.

The efficiency in up-conversion (microwave-to-optics) is defined $\eta_{\text{up}} = |\langle \hat{a}_{\text{out}} \rangle / \langle \hat{b}_{\text{in}} \rangle|^2$ and in




Figure 2.1: **Generalized transducer model.** Microwave-to-optical transduction can be modeled as a beamsplitter, with microwave ports and optical ports. Noise added to the signal is shown on the output ports and the pumps which mediate the transduction are internal to the beamsplitter. The efficiency of the transducer η is represented by the transmission coefficient of the beamsplitter.

downconversion (optics-to-microwaves) is defined $\eta_{\text{down}} = |\langle \hat{b}_{\text{out}} \rangle / \langle \hat{a}_{\text{in}} \rangle|^2$. The bidirectional conversion efficiency is defined as the geometric mean, $\eta = \sqrt{\eta_{\text{up}} \eta_{\text{down}}}$, and is analogous to the transmission coefficient in the beamsplitter model. In practice, the transducer is embedded in some measurement setup, with unknown or difficult to accurately calibrate transmissions between the sources and the detectors. To extract the efficiency of the transducer itself, we measure all four scattering coefficients, and use their ratios in a way that isolates η from gain and loss in the measurement apparatus [R. W. Andrews et al. 2014] (more details in Sec. 2.2.4).

The added noise in transduction at a single frequency ω , referenced to the input port, is given by

$$N_{\text{add}}(\omega) = \frac{\langle \hat{\mathcal{F}}^\dagger(\omega) \hat{\mathcal{F}}(\omega) \rangle}{\eta(\omega)} \quad (2.1)$$

in units of average photon flux per Hz of bandwidth (photons/s/Hz). The threshold for quantum performance is considered $N_{\text{add}} < 1$, as this is the point at which a single photon quantum signal will be transduced with a signal-to-noise ratio greater than 1. Equation 2.1 highlights an important design constraint. Quantum transducers that are strongly driven by noise must operate with high efficiency, whereas those weakly driven by noise can tolerate more loss. We will see that our implementation falls squarely into the first category, while many of the other architectures outlined in Sec. 1.1 fall in the latter.

The transduction bandwidth is defined as the full-width half-maximum of the transducer's efficiency as a function of frequency. The duty cycle is the fraction of the time the transducer pumps are turned on over some period of time. The product of efficiency, bandwidth, and duty cycle is proportional to the achievable average communication rate, neglecting noise effects. Comparing different transduction platforms, many operate in the low duty cycle, MHz-scale bandwidth regime, as heating from the pumps require pulsed operation. Our transducer is in the opposite limit, with continuous operation but kHz-scale bandwidth. As a result, the average expected communication rate of our transducer is comparable, if not surpassing, other platforms, once the noise can be reduced below the quantum level. For a numerical comparative analysis of different transducer experiments

that considers various networking protocols, see Dr. Maxwell Urmeý's thesis [Urmeý 2024]. There is also work to explore these trade offs using the tools of quantum channel capacity [C.-H. Wang et al. 2022].

2.2 Cavity opto-/electro-mechanics for efficient transduction

2.2.1 Hamiltonian formulation

A simplified picture of our transducer is depicted in Fig. 2.2(a). The motion of a highly-tensioned Si_3N_4 square membrane (purple) modulates the resonance frequency of a superconducting LC circuit (blue) and a high-finesse optical cavity (red). The Hamiltonian for this system is that of three harmonic oscillators, where the position of one of the oscillators leads to a frequency shift of the other two:

$$\hat{H}/\hbar = \omega_o \hat{a}^\dagger \hat{a} + \omega_e \hat{b}^\dagger \hat{b} + \omega_m \hat{c}^\dagger \hat{c} + g_o \hat{a}^\dagger \hat{a} (\hat{c} + \hat{c}^\dagger) + g_e \hat{b}^\dagger \hat{b} (\hat{c} + \hat{c}^\dagger). \quad (2.2)$$

The first three terms are the energies of the optical, microwave, and mechanical resonators, respectively. The fourth term couples the position $\hat{x} \sim \hat{c} + \hat{c}^\dagger$ of the membrane to the frequency of the optical cavity, at a vacuum coupling rate g_o . The last term is the analogous coupling to the microwave circuit, with a vacuum coupling rate g_e .

Optical- and microwave-frequency coherent pump tones displace the electromagnetic modes to \bar{a} and \bar{b} . Redefining $\hat{a} \rightarrow \bar{a} + \hat{a}$, $\hat{b} \rightarrow \bar{b} + \hat{b}$, we can linearize Eq. 2.2 about these new coherent amplitudes [Aspelmeyer et al. 2014]:

$$\hat{H}_{\text{lin}}/\hbar = -\Delta_o \hat{a}^\dagger \hat{a} - \Delta_e \hat{b}^\dagger \hat{b} + \omega_m \hat{c}^\dagger \hat{c} + g_o \bar{a} (\hat{a}^\dagger + \hat{a}) (\hat{c} + \hat{c}^\dagger) + g_e \bar{b} (\hat{b}^\dagger + \hat{b}) (\hat{c} + \hat{c}^\dagger). \quad (2.3)$$

where $\Delta_o = \omega_{p,o} - \omega_o$ and $\Delta_e = \omega_{p,e} - \omega_e$ are the detunings of the pumps from their respective electromagnetic resonators. Now in the rotating frame of the pump fields, \hat{a} and \hat{b} represent the

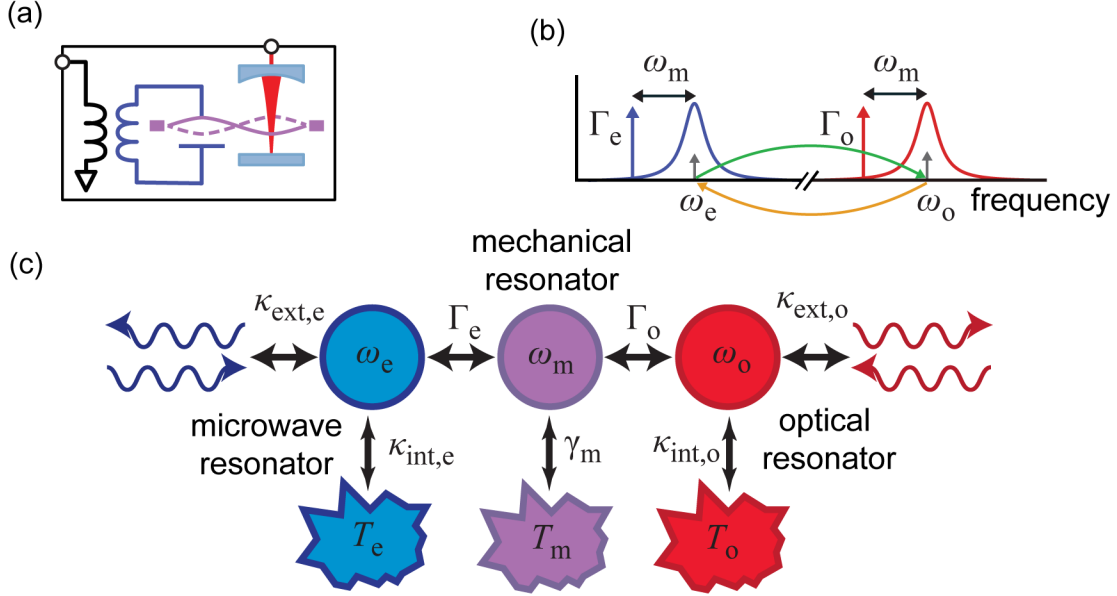


Figure 2.2: **Doubly-parametric mechanically-mediated microwave-to-optical transducer.**

(a) Schematic diagram. The mechanically-vibrating membrane (purple) modulates the frequency of an optical cavity (red) and a microwave-frequency superconducting circuit (blue). (b) Frequency-domain diagram. Two strong pump tones, red-detuned by the mechanical frequency ω_m from the microwave and optical cavities at frequencies ω_e and ω_o , enhance electromechanical and optomechanical interactions to rates Γ_e and Γ_o . A signal incident on one electromagnetic mode is then up-converted (green arrow) or down-converted (yellow arrow) to the other electromagnetic mode. (c) Coupled mode diagram. Each resonant mode is also coupled to other baths. The microwave (optical) modes are coupled to a flying photonic mode at rates $\kappa_{\text{ext},e}$ ($\kappa_{\text{ext},o}$) and the thermal bath at temperature T_e (T_o) at the cavity's intrinsic loss rate $\kappa_{\text{int},e}$ ($\kappa_{\text{int},o}$). Lastly, the mechanical oscillator is driven by a thermal bath at temperature T_m and the intrinsic mechanical damping rate γ_m .

fluctuations of the modes from the coherent displacements. Note, we have ignored constant average displacements of the membrane due to the average radiation pressure forces proportional to $|\bar{a}|^2$ or $|\bar{b}|^2$, which just redefine the origin of the membranes position. As in Fig. 2.2(b), if we set both pumps to be red-detuned by the mechanical frequency ($\Delta_o, \Delta_e = -\omega_m$), then we have the Hamiltonian of three oscillators at the same frequency, that exchange energy via beamsplitter terms like $\hat{a}\hat{c}^\dagger$ or $\hat{b}^\dagger\hat{c}$ at rates given $g_o|\bar{a}|$ and $g_e|\bar{b}|$.

2.2.2 Input-output theory for cavity electromechanics

The last pieces of the puzzle are the decay rates of each oscillator, which are detailed in Fig 2.2(c). The optical and microwave cavities are coupled to traveling photonic modes at rates $\kappa_{\text{ext},o}$ and $\kappa_{\text{ext},e}$, as well as to thermal baths due to some intrinsic dissipation rates $\kappa_{\text{int},o}$ and $\kappa_{\text{int},e}$. The mechanical oscillator is also coupled to its own thermal bath at its intrinsic damping rate γ_m . Using the tools of input-output theory, we can calculate the pump-enhanced optomechanical and electromechanical damping rates Γ_o and Γ_e .

For simplicity, we will calculate only the electromechanical effects, as the optomechanical effects can be derived by straightforward substitution into the final expressions. The set of coupled linear equations of motion describing the coupled circuit and membrane in the time domain is:

$$\begin{aligned}\dot{\hat{b}} &= (i\Delta_e - \kappa_e/2)\hat{b} + ig_e\bar{b}(\hat{c} + \hat{c}^\dagger) + \sqrt{\kappa_{\text{ext},e}}\hat{b}_{\text{in}} + \sqrt{\kappa_{\text{int},e}}\hat{f}_{\text{in},e} \\ \dot{\hat{c}} &= (-i\omega_m - \gamma_m/2)\hat{c} + ig_e\bar{b}(\hat{b} + \hat{b}^\dagger) + \sqrt{\gamma_m}\hat{f}_{\text{in},m}\end{aligned}\tag{2.4}$$

where $\kappa_e = \kappa_{\text{ext},e} + \kappa_{\text{int},e}$ is the total coupling rate of the microwave circuit, \hat{b}_{in} represents the input state from the transmission, and $\hat{f}_{\text{in},e}$ and $\hat{f}_{\text{in},m}$ are the stochastic forces from the thermal baths coupled to the microwave circuit and membrane, respectively.

Initially, we will ignore the mechanical mode and calculate the intracavity field for just the

microwave cavity, while also assuming no noise from the loss port. In the steady-state,

$$\hat{b} = \frac{\sqrt{\kappa_{\text{ext},e}}}{\kappa_e/2 - i\Delta_e} \hat{b}_{\text{in}} \quad (2.5)$$

where \hat{b}_{in} is supplied by the microwave pump field, which has a power $P_e = \hbar\omega_{p,e} \langle \hat{b}_{\text{in}}^\dagger \hat{b}_{\text{in}} \rangle$. For a monochromatic input field $\hat{b}_{\text{in}} = \bar{b}_{\text{in}} e^{-i\omega_{p,e}t}$ (such as that provided by the microwave generator or laser), the average intracavity photon number is

$$\bar{n}_e = |\bar{b}|^2 = \frac{P_e}{\hbar\omega_{p,e}} \frac{\kappa_{\text{ext},e}}{(\kappa_e/2)^2 + \Delta_e^2}. \quad (2.6)$$

Now, let's consider the effect of the electromechanical coupling on the susceptibility of the mechanical oscillator. The signal transduction (excluding noise effects) is governed by the linear response of this system, so we can begin by excluding these noise terms. We can solve for the response of the membrane in the presence of the pump by taking the Fourier transform of Eq. 2.4 and solving for the susceptibility of the mechanical oscillator. In the limit of small coupling ($g_e \ll \kappa_e$), the microwave pump induces a frequency shift $\delta\omega_{m,e}$ and adds additional damping at a rate Γ_e to the mechanical oscillator:

$$\delta\omega_{m,e} = g_e^2 \bar{n}_e \left(\frac{\Delta_e + \omega_m}{\kappa_e^2/4 + (\Delta_e + \omega_m)^2} + \frac{\Delta_e - \omega_m}{\kappa_e^2/4 + (\Delta_e - \omega_m)^2} \right) \quad (2.7)$$

$$\Gamma_e = g_e^2 \bar{n}_e \left(\frac{\kappa_e}{\kappa_e^2/4 + (\Delta_e + \omega_m)^2} - \frac{\kappa_e}{\kappa_e^2/4 + (\Delta_e - \omega_m)^2} \right) \quad (2.8)$$

The first term in either equation arises from the beamsplitter-type terms in the interaction Hamiltonian, and the second term is due to the two-mode squeezing or amplification terms ($\hat{b}^\dagger \hat{c}^\dagger + \hat{b} \hat{c}$). Eq. 2.7 and Eq. 2.8 highlight the fact that the microwave cavity acts as a filter for these two interactions, such that the choice of pump detuning changes the relative strength of the beamsplitter and two-mode squeezing effects.

For a pump perfectly red-detuned from the microwave cavity by the mechanical frequency

($\Delta_e = -\omega_m$), the dominant effect is an increase in the apparent loss rate of the membrane due to its coupling to the microwave circuit. In the resolved sideband limit ($4\omega_m \gg \kappa_e$), the pump-enhanced electromechanical damping rate is:

$$\Gamma_e = \frac{4g_e^2 \bar{n}_e}{\kappa_e}. \quad (2.9)$$

Substituting in the corresponding optical parameters, the average optical intracavity photon number from a pump with power P_o and the optomechanical damping rate Γ_o are:

$$\bar{n}_o = |\bar{a}|^2 = \frac{P_o}{\hbar\omega_{p,o}} \frac{\kappa_{\text{ext},o} \epsilon_{\text{PC}}}{(\kappa_o/2)^2 + \Delta_o^2} \quad (2.10)$$

$$\Gamma_o = g_o^2 \bar{n}_o \left(\frac{\kappa_o}{\kappa_o^2/4 + (\Delta_o + \omega_m)^2} - \frac{\kappa_o}{\kappa_o^2/4 + (\Delta_o - \omega_m)^2} \right). \quad (2.11)$$

In [Eq. 2.10](#), we need to include an additional factor ϵ_{PC} that captures how well the laser pump field is matched to the cavity mode. Otherwise, the optical cavity and microwave circuit couple to the membrane in *exactly* the same way, despite the five orders of magnitude energy difference between them. The coupling rate between one of the electromagnetic cavities and the mechanical oscillator can be cast into a single dimensionless quantity known as the electromechanical or optomechanical cooperativity C_e or C_o :

$$C_{e,o} = \frac{\Gamma_{e,o}}{\gamma_m} = \frac{4g_{e,o}^2 \bar{n}_{e,o}}{\gamma_m \kappa_{e,o}}. \quad (2.12)$$

In the next few sections, we will see that $C_e, C_o \gg 1$ leads to cooling of the mechanical mode and high efficiency transduction.

2.2.3 Transducer scattering parameters

The full-system with all three oscillators and their couplings to external ports, to each other, and to their respective thermal baths is represented by a system of first-order, linear differential

equations efficiently described as a state-space model [R. W. Andrews et al. 2014]:

$$\begin{aligned}\dot{\mathbf{a}}(t) &= A\mathbf{a}(t) + B\mathbf{a}_{\text{in}}(t) \\ \mathbf{a}_{\text{out}}(t) &= C\mathbf{a}(t) + D\mathbf{a}_{\text{in}}(t)\end{aligned}\tag{2.13}$$

where $\mathbf{a} = (\hat{a}, \hat{b}, \hat{c}, \hat{a}^\dagger, \hat{b}^\dagger, \hat{c}^\dagger)$ is a vector representing the internal resonator modes of the transducer. $\mathbf{a}_{\text{in}} = (\hat{a}_{\text{in}}, \hat{f}_{\text{in,o}}, \hat{b}_{\text{in}}, \hat{f}_{\text{in,e}}, \hat{f}_{\text{in,m}}, \hat{a}_{\text{in}}^\dagger, \hat{f}_{\text{in,o}}^\dagger, \hat{b}_{\text{in}}^\dagger, \hat{f}_{\text{in,e}}^\dagger, \hat{f}_{\text{in,m}}^\dagger)^T$ describes the state of the input fields, either from the external ports or the thermal baths, and $\mathbf{a}_{\text{out}} = (\hat{a}_{\text{out}}, \hat{b}_{\text{out}}, \hat{a}_{\text{out}}^\dagger, \hat{b}_{\text{out}}^\dagger)^T$ is a vector containing the output fields. The matrices A, B, C , and D contain the terms that appear in Eq. 2.4 and the appropriate complex conjugates. We can define a transfer function $\Xi(\omega)$ that describes how the output fields are related to the input fields:

$$\begin{aligned}\mathbf{a}_{\text{out}}(\omega) &= \Xi(\omega)\mathbf{a}_{\text{in}}(\omega) \\ \Xi(\omega) &= C(-i\omega I - A)^{-1}B + D.\end{aligned}\tag{2.14}$$

With these tools in hand, the microwave-to-optical scattering parameter $\mathcal{S}_{\text{eo}}(\omega) = \hat{a}_{\text{out}}/\hat{b}_{\text{in}}$ is given:

$$\mathcal{S}_{\text{eo}}(\omega) = \sqrt{\mathcal{A}_{\text{e}}\mathcal{A}_{\text{o}} \frac{\kappa_{\text{ext,e}}}{\kappa_{\text{e}}} \frac{\kappa_{\text{ext,o}}}{\kappa_{\text{o}}}} \frac{\sqrt{\Gamma_{\text{e}}\Gamma_{\text{o}}}}{\Gamma_{\text{T}}/2 - i(\omega - \omega'_{\text{m}})} e^{i\phi}\tag{2.15}$$

where we have defined the total mechanical damping rate $\Gamma_{\text{T}} = \Gamma_{\text{e}} + \Gamma_{\text{o}} + \gamma_{\text{m}}$ and the optomechanically- and electromechanically-shifted membrane frequency $\omega'_{\text{m}} = \omega_{\text{m}} + \delta\omega_{\text{m,e}} + \delta\omega_{\text{m,e}}$. Γ_{T} also sets the characteristic linewidth of the transduction scattering parameter, ie the transduction bandwidth, which is centered on the shifted mechanical frequency ω'_{m} . In addition to frequency transduction, there can also be signal amplification, which is captured by $\mathcal{A}_{\text{e}}\mathcal{A}_{\text{o}}$, and a phase shift ϕ of the signal :

$$\mathcal{A}_{\text{e}}\mathcal{A}_{\text{o}} = \left[\frac{(\Delta_{\text{e}} - \omega'_{\text{m}})^2 + (\kappa_{\text{e}}/2)^2}{-4\Delta_{\text{e}}\omega'_{\text{m}}} \right] \times \left[\frac{(\Delta_{\text{o}} - \omega'_{\text{m}})^2 + (\kappa_{\text{o}}/2)^2}{-4\Delta_{\text{o}}\omega'_{\text{m}}} \right]\tag{2.16}$$

$$\phi = \arctan\left(\frac{\kappa_{\text{e}}/2}{\Delta_{\text{e}} + \omega'_{\text{m}}}\right) + \arctan\left(\frac{\kappa_{\text{o}}/2}{\Delta_{\text{o}} + \omega'_{\text{m}}}\right).\tag{2.17}$$

While amplification may seem like a desirable feature in order to overcome any noise added in transduction, this amplification unavoidably adds noise [Caves 1982]. For a purely classical transducer with some other dominant noise source, this optomechanical gain could improve the signal-to-noise ratio of the subsequent measurement [Metelmann and Clerk 2014]. For a quantum transducer, however, the associated backaction noise is on par with the other noise sources and could result in decoherence of the quantum signal. For optimally red-detuned pumps ($\Delta_e = \Delta_o = -\omega'_m$), the gain becomes $\mathcal{A}_e \mathcal{A}_o = \left[1 + \left(\frac{\kappa_e}{4\omega'_m}\right)^2\right] \left[1 + \left(\frac{\kappa_o}{4\omega'_m}\right)^2\right]$, revealing that the gain (and the noise it adds) is minimized in the resolved sideband limit ($4\omega_m \gg \kappa_e$).

The expression for optical-to-microwave scattering parameter $\mathcal{S}_{oe}(\omega) = \hat{b}_{\text{out}}/\hat{a}_{\text{in}}$ is the same as Eq. 2.15. Absent from these expressions are the optical modematchings, which are in general different for up- and down-conversion. Following the prescription of [Brubaker et al. 2022], we can redefine the input and output ports to contain these modematching factors. In down-conversion, the damp-cavity mode matching is relevant ($\mathcal{S}_{oe} \rightarrow \sqrt{\epsilon_{\text{PC}}} \mathcal{S}_{oe}$), while in up-conversion, the modematching of the cavity to the local oscillator (LO) of our detection chain is relevant ($\mathcal{S}_{eo} \rightarrow \sqrt{\epsilon_{\text{CL}}} \mathcal{S}_{eo}$).

The efficiency in upconversion and downconversion for a signal incident on resonance can be defined $\eta_{\text{up}} = |\mathcal{S}_{oe}|^2/\mathcal{A}_e \mathcal{A}_o$ and $\eta_{\text{down}} = |\mathcal{S}_{eo}|^2/\mathcal{A}_e \mathcal{A}_o$. By dividing the gain, we are defining a conservative estimate of the efficiency in order to compare with an ideal quantum transducer which operates without gain [R. Andrews 2015]. The bidirectional conversion efficiency, introduced in Sec. 2.1, is simply:

$$\eta = \epsilon \frac{\kappa_{\text{ext},e}}{\kappa_e} \frac{\kappa_{\text{ext},o}}{\kappa_o} \frac{4\Gamma_e \Gamma_o}{\Gamma_{\text{T}}^2} \quad (2.18)$$

where $\epsilon = \sqrt{\epsilon_{\text{PC}} \epsilon_{\text{CL}}}$ is a "bidirectional" mode matching. Γ_e and Γ_o are controlled by the strength of the microwave or optical pump strength and can be made equal and much larger than the intrinsic mechanical damping rate in our system ($\Gamma_e = \Gamma_o \sim \text{kHz}$; $\gamma_m \sim \text{Hz}$). In this limit of $\Gamma_e = \Gamma_o \gg \gamma_m$, the "internal" transduction efficiency $\eta_{\text{int}} = \frac{4\Gamma_e \Gamma_o}{\Gamma_{\text{T}}^2}$ approaches unity. The total efficiency, which describes the transduction from input port to output port, is only limited by the external coupling ratios of the two electromagnetic cavities, as well as the modematching.

2.2.4 Measuring transducer efficiency

Our transducer will necessarily be embedded in some measurement network, with losses on the input ports and gain on the output ports. This system is shown schematically in [Fig. 2.3](#). A transmission measurement of upconversion and downconversion will then give $\alpha\mathcal{S}_{oe}\delta$ and $\gamma\mathcal{S}_{eo}\beta$, so the measured efficiency is altered by the measurement chain. In principle, one could measure these gains and losses, but they can drift over time. Instead, by in-situ measuring the reflections from the transducer off-resonance where it nearly perfectly reflects the signal, we can extract $\alpha\beta$ and $\gamma\delta$. Then, we can calculate a quantity that is insensitive to the imperfect individual measurements:

$$\mathcal{S}_{eo}\mathcal{S}_{oe} = \frac{(\gamma\mathcal{S}_{eo}\beta)(\alpha\mathcal{S}_{oe}\delta)}{(\alpha\beta)(\gamma\delta)}. \quad (2.19)$$

The magnitude squared of [Eq. 2.19](#) is then proportional to the bidirectional conversion efficiency [\[R. W. Andrews et al. 2014\]](#). This process assumes that the losses and gains are the same on- and off-resonance of the transducer, which is generally true for detunings on the scale of a few cavity linewidths.




Figure 2.3: **Efficiency Calibration.** The transducer is embedded in some network, with losses and gains between the source and detectors. By measuring upconversion, downconversion, and the off-resonance reflections, we can extract the bidirectional transducer efficiency without having to independently calibrate those losses and gains. Figure adapted from Fig. 2 of [\[R. W. Andrews et al. 2014\]](#).

2.3 Opto/electromechanical sideband cooling

Up to this point, we have largely ignored the noise ports in our model of the transducer. Our transfer function formalism is also applicable to these noise terms, whose impact on the output spectral density of the output fields $\mathcal{S}(\omega)$ is given:

$$2\pi\mathcal{S}(\omega)\delta(\omega - \omega') = \Xi^*(\omega') \left\langle [\mathbf{a}_{\text{in}}(\omega')]^\dagger \mathbf{a}_{\text{in}}^T(\omega) \right\rangle \Xi^T(\omega). \quad (2.20)$$

The autocorrelators of the input modes encode the details of the temperature and quantum noise that the thermal bath imparts on the transducer. If we assume that a mode $\hat{d}_i \in \{\hat{a}_{\text{in}}, \hat{f}_{\text{in,o}}, \hat{b}_{\text{in}}, \hat{f}_{\text{in,e}}, \hat{f}_{\text{in,m}}\}$ can be represented by a thermal bath at temperature T_i , then:

$$\begin{aligned} \left\langle [\hat{d}_i(\omega')]^\dagger \hat{d}_i(\omega) \right\rangle &= 2\pi\delta(\omega - \omega') n_{\text{th},i} \\ \left\langle \hat{d}_i(\omega) [\hat{d}_i(\omega')]^\dagger \right\rangle &= 2\pi\delta(\omega - \omega') (n_{\text{th},i} + 1) \end{aligned} \quad (2.21)$$

where $n_{\text{th},i}$ is the thermal bath population given by the Planck distribution $1/(e^{\hbar\omega_i/k_{\text{B}}T_i} - 1)$ for a mode at frequency ω_i . The asymmetry in these autocorrelators represents the quantum nature of the bath, and we will show how to exploit this fact to extract the temperature of the mechanical oscillator in [Sec. 2.3.1](#).

The input optical mode \hat{a}_{in} and optical cavity bath $\hat{f}_{\text{in,o}}$ are deep into their quantum ground state $n_{\text{th},i} \ll 1$, even at room temperature. By operating at the ~ 20 mK base plate of a dilution refrigerator and with appropriate filtering and attenuation, we can ensure that both the input microwave mode \hat{b}_{in} and microwave cavity bath $\hat{f}_{\text{in,e}}$ are also cold. The MHz-frequency bath of the mechanical oscillator $\hat{f}_{\text{in,m}}$ is however far from its ground state, with hundreds of phonons of occupancy at 20 mK. This was the main source of added noise in an early iteration of our transducer operated at a higher temperature of 4 K [[R. W. Andrews et al. 2014](#)].

The transducer architecture naturally provides a solution to this problem. The pump tones mediating the beamsplitter interaction that enables frequency transduction also cool the membrane

motion by coupling the "hot" mechanical oscillator to the "cold" electromagnetic resonators. Cooling the phonon occupancy of the membrane mode n_m to below 1 (the threshold of ground-state cooling) is a necessary condition for quantum-enabled transduction, as both must happen at rates faster than environmental phonons enter via the mechanical bath.




Figure 2.4: **Electromechanical sideband cooling.** Suppression of the Stokes $\Gamma_{e,-}$ relative to the anti-Stokes $\Gamma_{e,+}$ by the microwave cavity response results in electromechanical sideband cooling of the membrane mode when the pump is red-detuned. The Stokes process adds phonons, while the anti-Stokes process removes them, depicted by the ladder on the right.

For concreteness, let's consider electromechanical cooling, shown schematically in Fig. 2.4 as the results are analogous for the optomechanical interaction. Electromechanical damping is a competition of the Stokes and anti-Stokes processes, which occur at rates $(n_m + 1)\Gamma_{e,-}$ and $n_m\Gamma_{e,+}$. $\Gamma_{e,-}$ and $\Gamma_{e,+}$ are the negative and positive contributions to Eq. 2.8. As shown in Fig. 2.4, the Stokes process involves the scattering of a pump photon to lower energies after the creation of a phonon in the mechanical mode, effectively heating the membrane. The complementary anti-Stokes process cools the motion, and is enhanced relative to the Stokes process for a pump red-detuned from the cavity resonance because the resultant scattered pump photons benefit from buildup in the cavity [Aspelmeyer et al. 2014].

Due to the finite linewidth of the microwave circuit, the Stokes process will always contribute a non-zero heating rate. The ultimate limit of electromechanical cooling is thus the ratio of the

Stokes process to the total electromechanical damping rate:

$$n_{\min,e} = \frac{\Gamma_{e,-}}{\Gamma_{e,+} - \Gamma_{e,-}} = \frac{(\kappa_e/2)^2 + (\Delta_e + \omega'_m)^2}{-4\Delta_e\omega'_m}. \quad (2.22)$$

For an optimally red-detuned pump tone ($\Delta_e = -\omega'_m$), this reduces to $n_{\min,e} = \left(\frac{\kappa_e}{4\omega'_m}\right)^2$ which approaches zero in the resolved-sideband regime. This limit is often referred to as the quantum backaction limit. For the optomechanical interaction, $n_{\min,o} = \left(\frac{\kappa_o}{4\omega'_m}\right)^2$. The noise from these backaction baths is exactly the noise associated with non-unity transducer gain, which can also be seen by noticing that $\mathcal{A}_e = 1 + n_{\min,e}$ and $\mathcal{A}_o = 1 + n_{\min,o}$.

The backaction limit can also be interpreted as the temperature of the bath to which the membrane is electromechanically coupled, which suggests that the mechanical mode cannot get colder than the source of cooling, reminiscent of the second law of thermodynamics. Using this picture, it's then straightforward to calculate the mechanical occupancy under the effects of optomechanical and electromechanical cooling. The final occupancy is a weighted average of the membranes coupling to various baths, where the weights are the coupling rates.

$$n_m = \frac{\gamma_m n_{\text{th}} + \Gamma_e n_e + \Gamma_o n_o}{\Gamma_T} \quad (2.23)$$

In this expression n_{th} is the phonon occupancy of the MHz-frequency mechanical bath. The occupation of the microwave (optical) cavity $n_e = n_{\min,e} + n_{\text{eff},e}$ ($n_o = n_{\min,o} + n_{\text{eff},o}$) has a contribution from quantum backaction, and also any technical noise that leads to an increased effective occupancy $n_{\text{eff},e}$ ($n_{\text{eff},o}$). Phase noise on the pump tones and elevated temperature of either cavity (especially the microwave cavity) will contribute to this effective occupation from technical noise. With sufficiently low technical noise and large sideband resolution, the mechanical mode can be cooled to a phonon of occupation significantly less than one, approaching its quantum ground state.

2.3.1 Membrane thermometry

Any residual thermal excitations in the membrane mode will contribute to noise in transduction, as we will not be able to distinguish the signal and noise photons. We have two techniques to quantify the residual occupancy of the mechanical membrane. The first requires calibration of the measurement chain, while the latter exploits the asymmetry due to the quantum nature of the cooling interaction. As in the previous sections, we will mainly give the expressions for the electromechanical measurements, while commenting on how the optical measurement is related.

The power spectral density of the microwave output mode on the upper (anti-Stokes) $S_{e,+}[\omega]$ and lower (Stokes) $S_{e,-}[\omega]$ sidebands is given [Teufel et al. 2011; Brubaker et al. 2022]:

$$S_{e,+}[\omega] = 1/2 + n_{\text{add}} + \beta \frac{\kappa_{\text{ext},e}}{\kappa_e} \frac{\Gamma_e \Gamma_T (1 + n_{\text{min},e}) n_m}{\Gamma_T^2/4 + (\omega - \omega'_m)^2} \quad (2.24)$$

$$S_{e,-}[\omega] = 1/2 + n_{\text{add}} + \beta \frac{\kappa_{\text{ext},e}}{\kappa_e} \frac{\Gamma_e \Gamma_T n_{\text{min},e} (n_m + 1)}{\Gamma_T^2/4 + (\omega - \omega'_m)^2} \quad (2.25)$$

where n_{add} is the added noise of the detection chain and β is a generalized mode matching ($\beta = 1$ for microwave measurement, $\beta = \epsilon_{\text{CL}}$ for an optical measurement). Ideal heterodyne will necessarily add 1/2 a photon of noise [Yuen and Chan 1983] due to quantum fluctuations of the signal. In real optical and microwave measurements, there will be loss between the device and the detection, which will mix in vacuum and increase n_{add} . The microwave measurement also suffers the input referred added noise of the first amplifier, assuming the noise of this amplifier dominates the noise of measurement chain. This amplification is required to elevate the signal above the noise of the room temperature detection. While quantum-limited amplifiers in the microwave regime do exist [Castellanos-Beltran and Lehnert 2007; Roy and Devoret 2016], most of the measurements in this work were performed with a standard high-electron mobility transistor (HEMT) amplifier, which adds around 10-20 photons of noise.

In practice, we normalize the individual electromechanical spectrum measurements by repeating

them with the pump turned off, resulting in the 'background-normalized' upper sideband spectrum:

$$S_{\text{sn,e},+}[\omega] = 1 + \eta_{\text{meas}} \beta \frac{\kappa_{\text{ext,e}}}{\kappa_{\text{e}}} \frac{\Gamma_{\text{e}} \Gamma_{\text{T}} (1 + n_{\text{min,e}}) n_{\text{m}}}{\Gamma_{\text{T}}^2/4 + (\omega - \omega'_{\text{m}})^2} \quad (2.26)$$

with η_{meas} describing the efficiency of the measurement (not including the inefficiency of ideal heterodyne detection). For real measurements, $\eta_{\text{meas}} \neq 1$ highlighting that Eq. 2.26 is not actually normalized to true shot noise.

To calibrate this measurement efficiency, we can use the mechanics as a calibrated noise source, shown in Fig. 2.5. We simplify the system by only using the microwave pump and operating in a regime such that $\Gamma_{\text{e}} n_{\text{e}} \ll \gamma_{\text{m}} n_{\text{th}}$. The thermomechanical peak height above the background in the shot-noise normalized spectrum is then proportional to $\eta_{\text{meas}} n_{\text{th}}$. By changing the temperature of the dilution refrigerator while holding Γ_{e} fixed, we can control the temperature of the mechanical mode and thus change the thermomechanical noise in a known manner (assuming the mechanics is well-thermalized to the base plate). The peak height above the background $N_{\text{sn,e}} = S_{\text{sn,e},+}[\omega'_{\text{m}}] - 1$ as a function of n_{th} will have a slope equal to the measurement efficiency times a quantity that can be calculated from the linear electromechanical system parameters at each temperature. In this convention, $\eta_{\text{meas}} = 1/(n_{\text{add}} + 1/2)$, which goes to 1 for $n_{\text{add}} = 1/2$, so our η_{meas} does not include the inefficiency of ideal heterodyne detection.

The overhead of this calibration method grows quickly if any of the mechanical or microwave parameters change as a function of temperature. In [Brubaker et al. 2022], we observed that both κ_{e} and γ_{m} changed with temperature. The second method of membrane thermometry, known as sideband asymmetry [Jayich et al. 2012; Safavi-Naeini et al. 2013; Weinstein et al. 2014], does not involve any temperature sweeps, and instead uses the asymmetry of the noise from the quantum bath as the calibration standard. The amplitude of the upper- and lower-sidebands above the background $N_{\text{out,e},\pm} = S_{\text{e},\pm}[\omega_{\text{m}}] - (1/2 + n_{\text{add}})$ are proportional to the anti-Stokes and Stokes scattering rates,

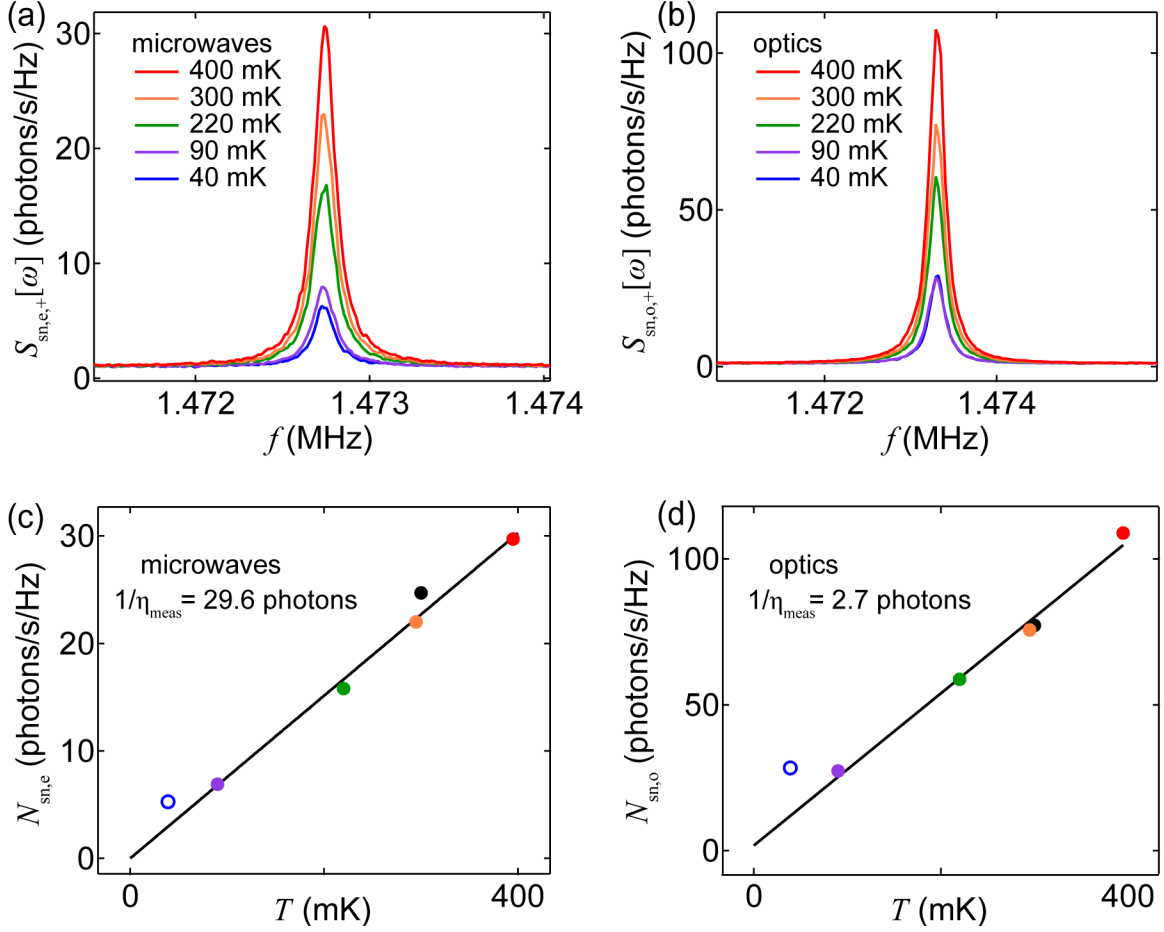


Figure 2.5: **Measurement efficiency calibration.** Sweeping the base temperature of the fridge turns the mechanical oscillator into a calibrated noise source. (a) Microwave and (b) optically-detected background-normalized upper sideband spectrum. The peak height increases with increasing temperature, and the slope of this dependence gives the added noise of (c) microwave and (b) optical measurement chains. The lowest temperature points for both measurements (open circles) are excluded from the fit, as we expect imperfect thermalization of the membrane to the base plate temperature. This figure is adapted from the supplement of [Higginbotham et al. 2018].

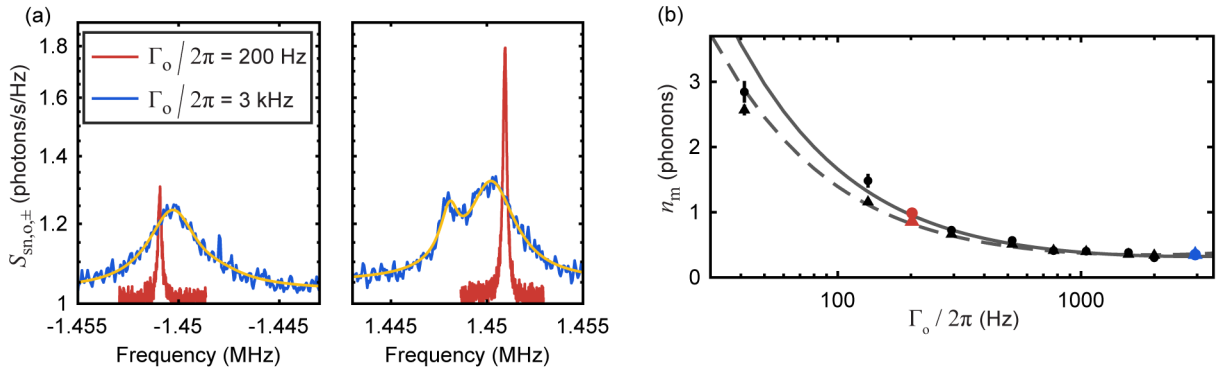


Figure 2.6: **Sideband Asymmetry Thermometry** (a) Background-normalized optical output spectrum with only optomechanical damping. The lower sideband is on the right, and the upper sideband on the left. At low damping rates, the peak heights are asymmetric due to the response of the optical cavity, and become more symmetric at high power as the membrane is cooled towards the backsciton limit. (b) Extracted membrane occupation using sideband asymmetry (circles are data, solid line is fit) and using the calibrated optical measurement efficiency (triangles are data, dashed line is fit). The red and blue data points correspond to the traces in (a). Figure adapted from [Brubaker et al. 2022].

which we saw earlier had an inherent asymmetry. The ratio of these sideband amplitudes:

$$r_e = \frac{N_{\text{out},e,-}}{N_{\text{out},e,+}} = \frac{n_{\text{min},e}}{n_{\text{min},e} + 1} \frac{n_m + 1}{n_m} \quad (2.27)$$

can be used to infer the mode occupancy. Only three measurements are required: the upper and lower sideband spectra and the linear response of the output cavity. At high temperatures, $n_m \approx n_m + 1$, and as the mode cools, r_e asymptotes to 1 as n_m approaches the backaction limit $n_{\text{min},e}$, as shown by the data in [Fig. 2.6](#)

There is a related technique that uses the quantum cross-correlations as the calibration standard in order to measure thermal cross-correlations which was not implemented in this thesis, but could provide yet another method of membrane thermometry [[Purdy et al. 2017](#)].

2.4 Transduction added noise

The thermal noise from the hot mechanical bath, the quantum backaction noise, and any noise injected by the pump tones will not only increase the membrane mode occupancy, but also add noise to incident signals. In upconversion, signals on resonance with the microwave circuit will be transduced to outgoing signals resonant with the optical cavity. The noise added is exactly the amplitude of the optomechanical upper sideband (in units of photons/s/Hz). Dividing by the transducer gain and efficiency gives the input-referred added noise:

$$N_{\text{add,up}} = \frac{N_{\text{out,o,+}}}{\mathcal{A}_e \mathcal{A}_o \eta_{\text{up}}} = \frac{n_m}{\mathcal{A}_e \frac{\kappa_{\text{ext},e}}{\kappa_e}} \frac{\Gamma_T}{\Gamma_e}. \quad (2.28)$$

This expression invites a perhaps surprising interpretation. Up to near unity factors of transducer gain and cavity overcoupling ratios, one can achieve $N_{\text{add,up}} < 1$ by ground-state cooling the mechanical element with the microwave pump tone and applying a small amount of optomechanical damping. In this limit of $\Gamma_o \ll \Gamma_e$ and $n_m \ll 1$, the transducer efficiency will be small, but the portion of the signal that does get upconverted will acquire minimal additional noise.

The situation is identical for downconversion, except for a modematching factor:

$$N_{\text{add,down}} = \frac{N_{\text{out,e,+}}}{\mathcal{A}_e \mathcal{A}_o \eta_{\text{down}}} = \frac{n_m}{\epsilon_{\text{PC}} \mathcal{A}_o \frac{\kappa_{\text{ext,o}}}{\kappa_o}} \frac{\Gamma_{\text{T}}}{\Gamma_o}. \quad (2.29)$$

In the absence of technical noise, optomechanical ground state cooling is necessary and sufficient for quantum downconversion. All of the thermal noise is coupled out of the input port due to the highly mismatched damping rates, so the little signal that is transduced sees little noise.

These intuitions change in the presence of technical noise, which is broadly defined as noise sources that scale with the electromechanical and optomechanical pump powers. Excluded from this definition is the heating of the mechanical mode by its coupling to a thermal environment with occupation n_{th} and backaction from the electro/optomechanical interactions. We also exclude backaction from an auxiliary beam used to lock the laser to the optical cavity. The power of this beam is much smaller than that of the optomechanical pump, but it is tuned near resonance, where backaction is maximal. Characterization and further discussion of this noise source can be found in Max Urmeý's thesis [Urmeý 2024]. We exclude it from the technical noise category as it does not scale with optomechanical pump power, instead entering as an extra term $\gamma_{\text{lock}} n_{\text{lock}}$ in the numerator of Eqn. 2.23.

Technical noise does encompass amplitude or phase modulation/noise on the pumps at ω_m , fluctuations of the cavity parameters induced by the pump, and heating of the microwave cavity by either of the pumps. In principle, the optical cavity could also be heated by the pumps, but is unobserved for the same reasons that optical frequencies are good carriers of quantum information. Phase and amplitude on the pump tones will directly contribute to occupation of the mechanical mode, complicate the thermometry techniques, and add noise to the transduction. Frequency noise and heating of the microwave circuit will also directly impact the membrane occupation. Many of these effects can be cast into an effective occupancy of the microwave or optical cavities $n_{\text{eff,e}}$ and $n_{\text{eff,o}}$.

2.4.1 Microwave cavity occupation from phase/amplitude noise

The effects of generator noise can be included in the input-output formalism introduced in [Sec. 2.2.2](#) by modification of the input fields to include classical noise terms. Following the formalism of [\[Jayich et al. 2012\]](#), the microwave input field with an photon flux of $\dot{N}_e = \frac{P_e}{\hbar\omega_{p,e}}$ is now:

$$\hat{b}_{\text{in}}(t) = e^{-i\omega_{p,e}t} \left[\sqrt{\dot{N}_e} + \frac{1}{2} (\delta x(t) + i\delta y(t)) \right] + \hat{\xi}_{\text{in},e}. \quad (2.30)$$

The first term is simply the coherent amplitude of the pump and the last term is the microwave vacuum noise. We will assume the microwave input is cold, so $\langle \hat{\xi}_{\text{in},e}(t) \hat{\xi}_{\text{in},e}^\dagger(t') \rangle = \delta(t - t')$ and $\langle \hat{\xi}_{\text{in},e}^\dagger(t) \hat{\xi}_{\text{in},e}(t') \rangle = 0$. This is a good approximation given appropriate design of the cryogenic microwave readout chain [\[Kriner et al. 2019\]](#).

The second and third term are classical variables describing technical generator amplitude and phase noise. We will assume a white noise model, as we are only interested in the noise in narrow bandwidth around the mechanical frequency.

$$\begin{aligned} \langle \delta x(t) \delta x(t') \rangle &= C_{xx} \delta(t - t') \\ \langle \delta y(t) \delta y(t') \rangle &= C_{yy} \delta(t - t') \\ \langle \delta x(t) \delta y(t') \rangle &= C_{xy} \delta(t - t') \end{aligned} \quad (2.31)$$

C_{xx} and C_{yy} are spectral densities in units of photons/s/Hz describing the amplitude and phase noise of the generator, and scale with the power in the pump tone. C_{xy} describes any correlations between the amplitude and phase noise and together, they obey the Cauchy-Bunyakovsky-Schwarz inequality $C_{xy}^2 \leq C_{xx}C_{yy}$. The laser classical noise is equal to its quantum noise when $C_{xx} = 1$ or $C_{yy} = 1$.

We can make the power dependence of the noise spectral densities explicit by relating them to the one-sided, double-sideband noise spectral densities relative to the carrier $S_{\delta A, \delta A, e}(\omega)$, $S_{\delta \phi, \delta \phi, e}(\omega)$, and $S_{\delta A, \delta \phi, e}(\omega)$. These are the units that are often specified by laser and microwave generator

manufacturers and thus familiar to experimentalists. To account for this convention (compared to that often employed by theorists), an extra factor of 1/2 is needed:

$$\begin{aligned} C_{xx} &= 4\dot{N}_e[S_{\delta A, \delta A, e}(\omega)/2] \\ C_{yy} &= 4\dot{N}_e[S_{\delta\phi, \delta\phi, e}(\omega)/2] \\ C_{xy} &= 4\dot{N}_e[S_{\delta A, \delta\phi, e}(\omega)/2] \end{aligned} \tag{2.32}$$

If the generator has classical noise fluctuations at ω'_m , when the pump is optimally red-detuned, these fluctuations will drive the mechanical mode. Thus, classical noise looks like an additional bath to which the membrane is coupled via the electromechanical interaction. The occupation of the microwave circuit from classical noise is then:

$$n_{\text{eff},e} = \frac{1}{4} \frac{\kappa_{\text{ext},e}}{\kappa_e} \mathcal{A}_e \frac{\kappa_e^2}{4} (|B_x(\omega'_m)|^2 C_{xx} + |B_y(\omega'_m)|^2 C_{yy} + 2\text{Im}[B_x(\omega'_m)B_y^*(\omega'_m)]C_{xy}) \tag{2.33}$$

where the noise susceptibilities are $B_x(\omega) = e^{-i\phi_e}\chi_e(\omega) + e^{i\phi_e}\chi_e^*(-\omega)$ and $B_y(\omega) = e^{-i\phi_e}\chi_e(\omega) - e^{i\phi_e}\chi_e^*(-\omega)$. $\chi_e(\omega) = 1/(\kappa_e/2 - i(\omega_e + \Delta))$ is the susceptibility of the microwave circuit and $\phi_e = \tan^{-1}(2\Delta/\kappa_e)$ is the phase of the intracavity field relative to the incident field.

In the resolved sideband limit with an optimally red-detuned pump, the amplitude and phase noise of the generator contribute equally to the microwave occupation as $|B_x(\omega'_m)|^2, |B_y(\omega'_m)|^2 \approx 4/\kappa_e^2$. For both the microwave generator and laser used in this thesis, phase noise dominated amplitude noise by at least 10 dB and there is little correlation between the two sources. The optical cavity occupation from noise on the laser is similar, except with an extra factor of pump-cavity modematching ϵ_{PC} .

2.4.2 Transducer added noise from phase/amplitude noise

As we are modeling the technical noise of the generator as white noise, it's natural that technical noise will also contribute to a background offset in the upper and lower sideband output spectrum [Eqn. 2.24](#) and [Eqn. 2.25](#). As we will see, the presence of the cavity will also make the sidebands asymmetrically sensitive to amplitude and phase noise. This excess background in the

output at the upper (lower) sideband $\tilde{S}_{e,+}$ ($\tilde{S}_{e,-}$) consists of three terms: the promptly reflected input field, the cavity field that leaks out, and the interference of these two terms. In total,

$$\begin{aligned} \tilde{S}_{e,\pm} = & \\ \frac{1}{4} \left((|\rho_e|^2 + |\kappa_{\text{ext},e}\chi_e(\pm\omega'_m) - 1|^2) (C_{xx} + C_{yy}) - 2\text{Re} [\rho_e^* [\kappa_{\text{ext},e}\chi_e(\pm\omega'_m) - 1](C_{xx} + 2iC_{xy} - C_{yy})] \right) \end{aligned} \quad (2.34)$$

where we have defined $\rho_e = 1 - \kappa_{\text{ext},e}/(\kappa_e/2 - i\Delta_e)$, the attenuation and phase shift of the pump reflected from the cavity. The optical case is the same, except again a factor of $\epsilon_{\text{PC}}\epsilon_{\text{CL}}$ to account for modematchings in the limit that all the modematchings are approximately the same ($\sqrt{\epsilon_{\text{PL}}} - \sqrt{\epsilon_{\text{PC}}\epsilon_{\text{CL}}} \sim 0$) where ϵ_{PL} is the pump-LO modematching.

In the absence of cavity ($\rho_e \rightarrow 1, \kappa_e \rightarrow 0$), amplitude noise C_{xx} will still contribute to an excess background. C_{xy} and C_{yy} change the background because the cavity transduces fluctuations in the phase quadrature to the amplitude quadrature. Under the ideal conditions for microwave-to-optical transduction with a very overcoupled cavity ($\kappa_{\text{ext},e} \gg \kappa_{\text{int},e}$) in the resolved sideband regime and optimal detuning, the upper sideband mainly sees contributions from phase noise, while the lower sideband mainly sees contributions from amplitude noise.

This excess white-noise background also contributes noise to the transduction, introduced on the output field by the output pump. Translating this to input-referred added noise requires dividing by the efficiency and gain, thus changing our intuition about how to achieve low-added noise from [Eqn. 2.28](#) and [Eqn. 2.29](#). We must add a term to account for this output noise on the upper sideband in downconversion $\tilde{N}_{\text{out},e} = \tilde{S}_{e,+}$ or upconversion $\tilde{N}_{\text{out},o} = \tilde{S}_{o,+}$, so the expressions for added noise now read:

$$N_{\text{add},\text{up}} = \frac{N_{\text{out},o,+} + \tilde{N}_{\text{out},o}}{\mathcal{A}_e \mathcal{A}_o \eta_{\text{up}}} = \frac{n_m}{\mathcal{A}_e \frac{\kappa_{\text{ext},e}}{\kappa_e}} \frac{\Gamma_{\text{T}}}{\Gamma_e} + \frac{\tilde{N}_{\text{out},o}}{\mathcal{A}_e \mathcal{A}_o \eta_{\text{up}}} \quad (2.35)$$

$$N_{\text{add},\text{down}} = \frac{N_{\text{out},e,+} + \tilde{N}_{\text{out},e}}{\mathcal{A}_e \mathcal{A}_o \eta_{\text{down}}} = \frac{n_m}{\epsilon_{\text{PC}} \mathcal{A}_o \frac{\kappa_{\text{ext},o}}{\kappa_o}} \frac{\Gamma_{\text{T}}}{\Gamma_o} + \frac{\tilde{N}_{\text{out},e}}{\mathcal{A}_e \mathcal{A}_o \eta_{\text{down}}}. \quad (2.36)$$

The first term in either expression can be minimized by ground-state cooling with the input pump, but low efficiency operation will increase the second term in the presence of technical noise on the laser or microwave generator. To minimize this background term, we can either reduce the phase noise of the pump or operate with high efficiency. We will explore some of these trade offs in [Sec. 2.5](#).

2.4.3 Corrections to membrane thermometry from phase/amplitude noise

The excess background noise added from the technical noise on the pumps will directly effect the measurement efficiency based thermometry in [Fig. 2.5](#). We can either correct for this by accounting for the excess background, or work at pump powers at which the excess background is negligible. In practice, we often choose the latter, as the upper sideband peak height from [Eqn. 2.24](#), and thus the signal-to-noise ratio (SNR) of the measurement, is maximized for $\Gamma_e = \gamma_m$. The pump powers required to achieve this matching are small enough that we can neglect the classical noise on the pump.

Sideband asymmetry thermometry ([Fig. 2.6](#)) is insensitive to the excess background, but technical noise also modifies the heights of the upper and lower sideband peaks [[Safavi-Naeini et al. 2013](#)]. This effect, known as squashing (anti-squashing), reduces (increases) the thermomechanical peak height due to correlations between phase/amplitude noise and the motional sidebands. In order to extract the true mechanical occupation, we must correct for these interference effects. The susceptibility of the upper and lower sidebands to technical noise is:

$$\begin{aligned} \tilde{B}_{\pm}[\omega] = & \frac{e^{-i\phi}}{4} (\kappa_{\text{ext},e} |\chi_e(\pm\omega)|^2 [B_x(\mp\omega)(C_{xx} + iC_{xy}) + B_y(\mp\omega)(iC_{xy} - C_{yy})] \\ & - \chi_e^*(\pm\omega) [(B_x(\mp\omega)C_{xx} + iB_y(\mp\omega)C_{xy})(1 + \rho_e) + (iB_x(\mp\omega)C_{xy} - B_y(\mp\omega)C_{yy})(1 - \rho_e)]). \end{aligned} \quad (2.37)$$

The real part of this susceptibility corresponds to a symmetric Lorentzian contribution to the peak height, while the imaginary part adds a Lorentzian contribution that is asymmetric about the mechanical frequency. The sign of the individual terms depends on the relative scale of the technical

noise spectral densities.

The microwave output spectrum in the presence of technical noise on the generator, accounting for the measurement chain added noise and technical noise effects is:

$$S_{e,+}[\omega] = 1/2 + n_{\text{add}} + \tilde{S}_{e,+} + \beta \frac{\kappa_{\text{ext},e}}{\kappa_e} \mathcal{A}_e \frac{\Gamma_e}{\Gamma_T^2/4 + (\omega - \omega'_m)^2} \\ \times \left[\Gamma_T \left(n_m - \beta' \frac{\kappa_e^2}{4} \text{Re}[\tilde{B}_+] \right) - 2\beta'(\omega - \omega'_m) \frac{\kappa_e^2}{4} \text{Im}[\tilde{B}_+] \right] \quad (2.38)$$

and

$$S_{e,-}[\omega] = 1/2 + n_{\text{add}} + \tilde{S}_{e,-} + \beta \frac{\kappa_{\text{ext},e}}{\kappa_e} \mathcal{A}_e \frac{\Gamma_e}{\Gamma_T^2/4 + (\omega + \omega'_m)^2} \\ \times \left[\Gamma_T \left(\frac{n_{\text{min},e}}{\mathcal{A}_e} (n_m + 1) + \beta' \frac{\kappa_e^2}{4} \text{Re}[\tilde{B}_-] \right) - 2\beta'(\omega + \omega'_m) \frac{\kappa_e^2}{4} \text{Im}[\tilde{B}_-] \right] \quad (2.39)$$

where β' is a generalized output modematching. For microwave measurements $\beta' = 1$. The optical output spectrum is the same except $\beta = \epsilon_{\text{CL}}, \beta' = \epsilon_{\text{PC}}$ and appropriately replacing microwave parameters with their optical counterparts. n_m is given by [Eqn. 2.23](#), accounting for both the backaction limit and the occupancy of the electromagnetic cavities due to noise on the corresponding pump tones.

While the full expression is quite complex, there are some simplifying assumptions that arise from the experimental parameters. Firstly, in both the optical and microwave domains, phase noise is often the main contribution to technical noise so we can let $C_{xx} = C_{xy} \rightarrow 0$. With this assumption, any pump power dependent change to the background of the output spectrum is a result of phase noise, providing an *in situ* measurement of the phase noise spectral density. This can be complicated by drifting modematchings in optical measurements or microwave circuit heating in microwave measurements.

Additionally, sideband asymmetry only compares the heights of the peaks above the background. On resonance, the anti-symmetric contribution to the spectrum vanishes. The ratio of the upper

and lower sidebands:

$$\tilde{r}_e = \frac{n_m - \beta' \frac{\kappa_e^2}{4} \text{Re}[\tilde{B}_+(\omega'_m)]}{\frac{n_{\text{min},e}}{\mathcal{A}_e} (n_m + 1) + \beta' \frac{\kappa_e^2}{4} \text{Re}[\tilde{B}_-(-\omega'_m)]} \quad (2.40)$$

is modified by the presence of technical noise. The mechanical occupation can still be accurately extracted using the measurement of the phase noise spectral density from the background and parameters of the microwave cavity. The measurement efficiency based thermometry must also account for the modified peak heights, while being careful to consider the excess background from technical noise if performed with pump powers that carry appreciable noise.

2.4.4 Frequency fluctuations and heating of the microwave circuit

The other source of technical noise is heating of the microwave circuit induced by the electromechanical and optomechanical pumps. Unlike the technical noise on the pumps, which can be managed by filtering [Joshi et al. 2021], this noise source originates from inside of the transducer and thus must be addressed by the design. Due to the large energy of a single optical photon, laser heating of the microwave circuit could potentially be the dominant effect. While this is true for many of the other transducer platforms discussed in Sec. 1.1, our transducer is largely insensitive to this effect due to light propagation predominantly in vacuum, spatial separation of the electrical and optical resonators, excellent optical mode control from robust cavity design [Urmev 2024], and a choice of superconductor resilient to high frequency radiation ([Burns 2019]).

Heating of the microwave circuit due to the microwave pump, however, is a significant effect for our transducer. In fact, we will see in Sec. 4.1 and Sec. 4.2 that microwave-pump induced noise and loss are the main limiting factors in the transducer's performance. Similar effects have been observed in other electromechanical experiments [Suh 2011; Teufel et al. 2011; Massel et al. 2012; Wollman 2015; Lei 2017] using silicon nitride membranes or aluminum drumheads. This noise and loss has been attributed to the presence of two-level system (TLS) type defects in the dielectrics, whether that be the amorphous Si_3N_4 that makes the membrane or oxides at the metal-air, metal-substrate, and substrate-air interfaces. The fractional frequency noise and internal microwave loss in linear

microwave circuits in the weakly driven, low temperature regime can be explained by the so-called "resonant" damping contribution to the TLS model. [Gao 2008; Paik and Osborn 2010; Burnett et al. 2014; Behunin et al. 2016].

Early work in understanding these effects in our transducers found that the TLS model indeed explained the noise and loss at low-power [Higginbotham et al. 2018; Burns 2019], yet there was a deviation from the model at high powers. In this thesis, particularly in Ch. 5, I will show that the microwave loss at all powers can be understood within the TLS model if one considers a self-heating effect that activates the temperature-dependent loss of the TLS in the Si_3N_4 . The so-called "relaxation" damping term becomes important, although it is usually negligible and thus ignored in most superconducting circuit experiments. Most importantly, I'll show that annealing the Si_3N_4 improves the material's loss rate and eliminates any power-dependent loss activated by self-heating in both planar test circuits and the transducer.

In Ch. 6, I'll explore how the self-heating effect paired with the relaxation damping term could also explain the noise observed in our transducers, suggesting that a more detailed TLS model can explain the noise that is limiting the transducer performance. As with the loss, annealing is found to improve the power-dependent cavity heating, and technical noise from the generator is improved by filtering. With the intuition gained in the next section, it is clear that these improvements to the microwave pump power-dependent circuit loss and noise should push our transducer architecture into quantum-enabled operation.

2.5 Understanding the parameter space

The expressions for added noise and efficiency are quite cumbersome in their full forms. This is especially true when the added noise is a sum of many small, approximately equal contributions. Additionally, the resolved sideband limit becomes invalid in the presence of sufficient pump-induced microwave loss. In this section, we will explore how the transduction noise, efficiency, and bandwidth depend individually on changes to the intrinsic mechanical damping rate γ_m , electromechanical vacuum coupling rate g_e , internal microwave loss rate $\kappa_{\text{int,e}}$, and occupancy of the microwave circuit

due to technical noise $n_{\text{eff,e}}$. This is key to understanding which parameters are most important to improve in order to achieve quantum transduction.

For each of the parameters studied below, we find the electromechanical and optomechanical damping rates that minimize the added noise in up- or down-conversion given a set of fixed transducer parameters and report the efficiency, bandwidth, and mechanical occupancy under those conditions. The optimal conditions will arise when the thermomechanical noise from the membrane roughly matches the technical noise added by the pumps. The transducer parameters we use in this section are summarized [Table 2.1](#). We will also assume that both pumps are perfectly red-detuned. For simplicity, the technical noise occupations of the optical and microwave cavities $n_{\text{eff,o}}$ and $n_{\text{eff,e}}$ are parameterized as linear functions of the optomechanical damping rate Γ_o or the microwave intracavity photon number \bar{n}_e . In the optical case, the only contribution is phase noise on the laser, which scales linearly with the laser power. Γ_o also scales linearly with laser power, so this is a good parameterization. For the microwave case, there is both phase noise from the generator and heating of the circuit, which are roughly approximated by this linear relation (for more about the scaling of the microwave noise, see [Ch. 3](#)). We use the optical and microwave technical noise scalings extracted for the device reported in [\[Brubaker et al. 2022\]](#).

The transducer parameters at optimized added noise as a function of microwave circuit loss $\kappa_{\text{int,e}}$ ([Fig. 2.7](#)), intrinsic mechanical loss γ_m ([Fig. 2.8](#)), vacuum electromechanical coupling g_e ([Fig. 2.9](#)), and relative microwave noise level ([Fig. 2.10](#)) are studied, with upconversion in green and downconversion in orange. We first point out a few general observations. In these figures, we notice that the noise in downconversion is generally smaller than the noise in upconversion, which is a result of less technical noise per unit optomechanical damping when compared to that of the electromechanical interaction. With technical noise, ground state cooling with the input pump is no longer a sufficient condition from quantum-enabled transduction, but still necessary. Optomechanical cooling is more effective than electromechanical cooling for the given system parameters (subpanel (d) of the [Fig. 2.7-Fig. 2.10](#)), and as a result, downconversion noise is smaller than upconversion noise. The asymmetry in the optical and microwave noise also explains the generally smaller efficiency in

downconversion compared to upconversion. Lowest downconversion noise operation is achieved with a large optical damping rate, and a small microwave damping rate. The damping rates are more matched in optimized upconversion because the optical pump is needed to reduce the mechanical occupancy past what the microwave pump can do alone.

Increased microwave loss has two primary effects: reduced transduction efficiency (Eqn. 2.18) and increased backaction due to decreased sideband resolution (Eqn. 2.22). These effects are evident in the input-referred added noise Fig. 2.7(a) and transduction efficiency Fig. 2.7(b). The decreased sideband resolution also places a limit on the optimal damping rates, which results in a decreased transduction bandwidth Fig. 2.7(c) and larger mechanical occupation Fig. 2.7(d) (both from thermal noise and backaction). In Ch. 5, we will explore the main source of loss in our transducers as well as review the materials analysis that enabled reduction of the microwave loss to the $\kappa_{\text{int,e}}/2\pi \approx 50 \text{ kHz}$ level.

Now, we explore the influence of mechanical dissipation on the transducer performance (Fig. 2.8). Increased γ_m couples the membrane more strongly to its thermal environment, requiring more pump power to achieve the same cooperative and cool the mechanical oscillator to the same occupation. The noise in transduction and mechanical occupancy thus increase with larger mechanical dissipation. Another side effect of the greater pump powers required for optimal noise performance is increased transducer bandwidth at the expense of the increased noise. The transduction efficiency is less sensitive to γ_m as $\Gamma_e, \Gamma_o \gg \gamma_m$ for the range discussed here. In Ch. 3, we will discuss the physics of mechanical dissipation in tensioned membranes and describe a phononic filter which significantly decreased the mechanical dissipation.

Increased vacuum electromechanical coupling improves every transducer performance metric, as shown in Fig. 2.9, especially because the $\Gamma_e \propto g_e^2$. We can achieve larger cooperativity for the same technical noise levels, which reduces the membrane occupation and thus the added noise as well as permitting higher bandwidth operation. Increasing g_e reduces asymmetry in the upconversion and downconversion performance, especially in the transduction efficiency. As electromechanical ground-state cooling becomes possible, we push into a regime of quantum-enabled upconversion with

high efficiency. In [Ch. 3](#), we will introduce a method to reduce the motional capacitor gap spacing, thus increasing g_e .

We parameterized the effective occupation of the microwave circuit as a linear function of the intracavity photon number. In [Fig. 2.10](#), the slope of this linear parametrization is scaled by 0.1 to 10, such that the noise increases along the x-axis of these plots. Naturally, the added noise and membrane occupation increase as the microwave pump contributes more noise. To manage this effect, the optimal electromechanical damping rate and cooperativity is also smaller, corresponding to a larger thermal occupation of the membrane and a reduced transduction bandwidth. In [Ch. 6](#), we will present potential models for this noise as well as strategies to reduce the noise or isolate the signal from it. These plots, all together, suggest that order of magnitude improvements must be made to one of these parameters in order to enable quantum transduction. In this thesis, we generally attempt to make these large improvements, but many smaller improvements can also push the transducer under the quantum threshold.

$\kappa_{\text{ext,e}}$	$2\pi \times 1 \text{ MHz}$
$\kappa_{\text{int,e}}$	$2\pi \times 200 \text{ kHz}$
$\kappa_{\text{ext,o}}$	$2\pi \times 1.5 \text{ MHz}$
$\kappa_{\text{int,o}}$	$2\pi \times 200 \text{ kHz}$
g_e	$2\pi \times 2 \text{ Hz}$
ω_m	$2\pi \times 1.5 \text{ MHz}$
T_m	100 mK $n_{\text{th}} \approx 1400 \text{ phonons}$
γ_m	$2\pi \times 0.2 \text{ Hz}$
$\epsilon_{\text{PC}}, \epsilon_{\text{CL}}, \epsilon_{\text{PL}}$	0.8
$\gamma_{\text{lock}} n_{\text{lock}}$	$2\pi \times 40 \text{ photons/Hz}$
$n_{\text{eff,o}} = a_o \Gamma_o$	$a_o = 2.8 \times 10^{-6} \text{ photons/Hz}$
$n_{\text{eff,e}} = a_e \bar{n}_e + b_e$	$a_e = 2.93 \times 10^{-8} \text{ photons/photon}$ $b_e = 0.09 \text{ photons}$

Table 2.1: **Example transducer parameters.** Typical transducer parameters used to calculate performance in this section. The parameters are similar to the device reported in [Brubaker et al. 2022]. The technical noise behavior is parameterized by a linear function of either the optomechanical damping rate or microwave pump intracavity photon number.

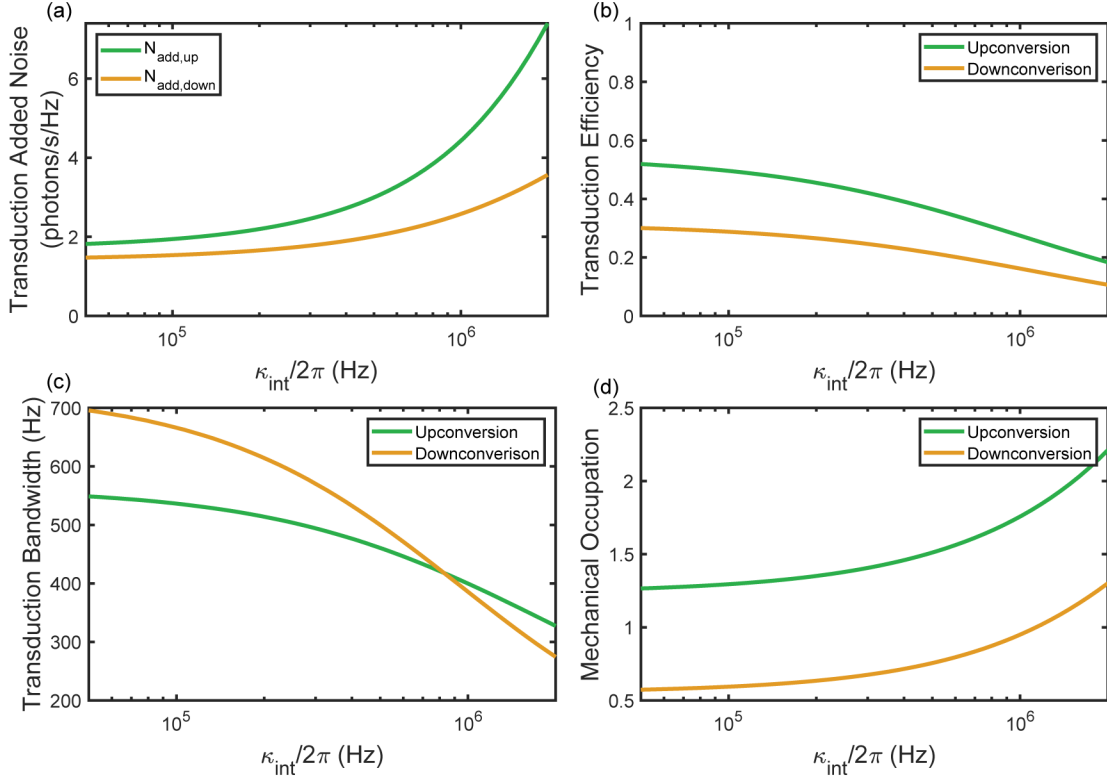


Figure 2.7: **Transduction performance versus microwave loss.** (a) Transduction Added Noise at optimal damping rates. Transduction efficiency (b), bandwidth (c) and occupation of the membrane (d) under the same parameters that minimize the added noise. The two dominant effects of increase microwave loss are decreased efficiency and increased backaction noise.

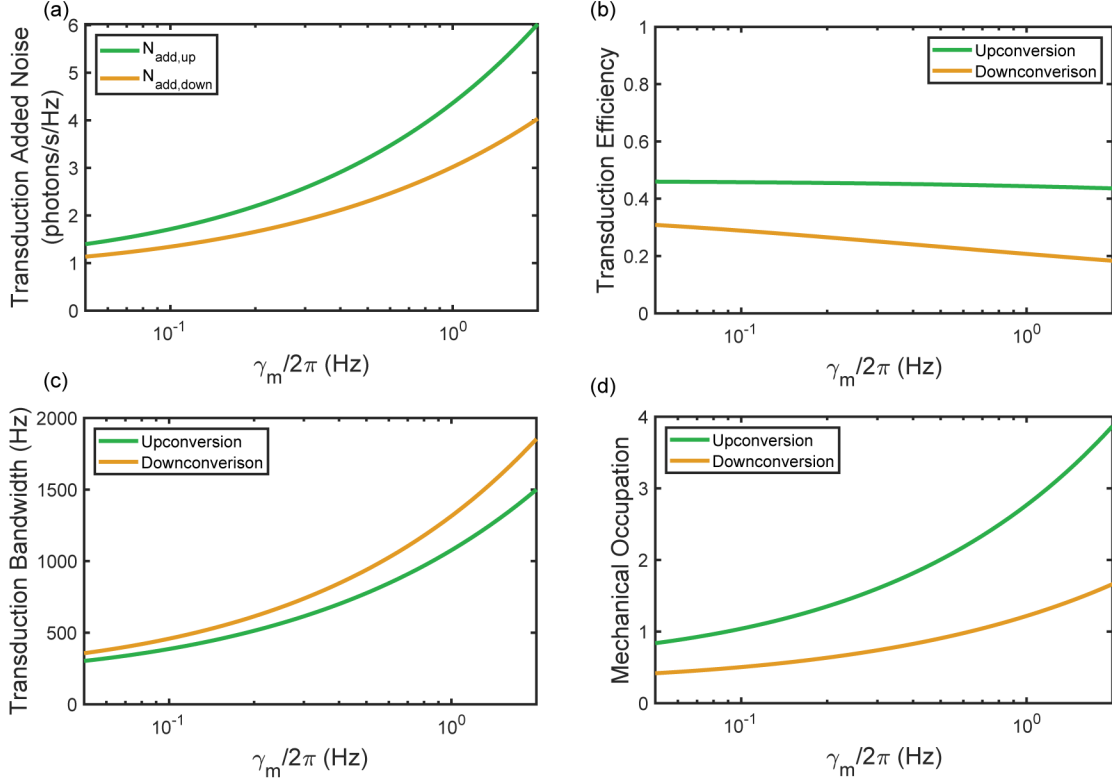


Figure 2.8: **Transduction performance versus mechanical loss.** (a) Transduction Added Noise at optimal damping rates. Transduction efficiency (b), bandwidth (c) and occupation of the membrane (d) under the same parameters that minimize the added noise. Mechanical dissipation couples the membrane to the highly occupied thermal bath, increasing the membrane occupation and requiring stronger pump tones to cool the motion.

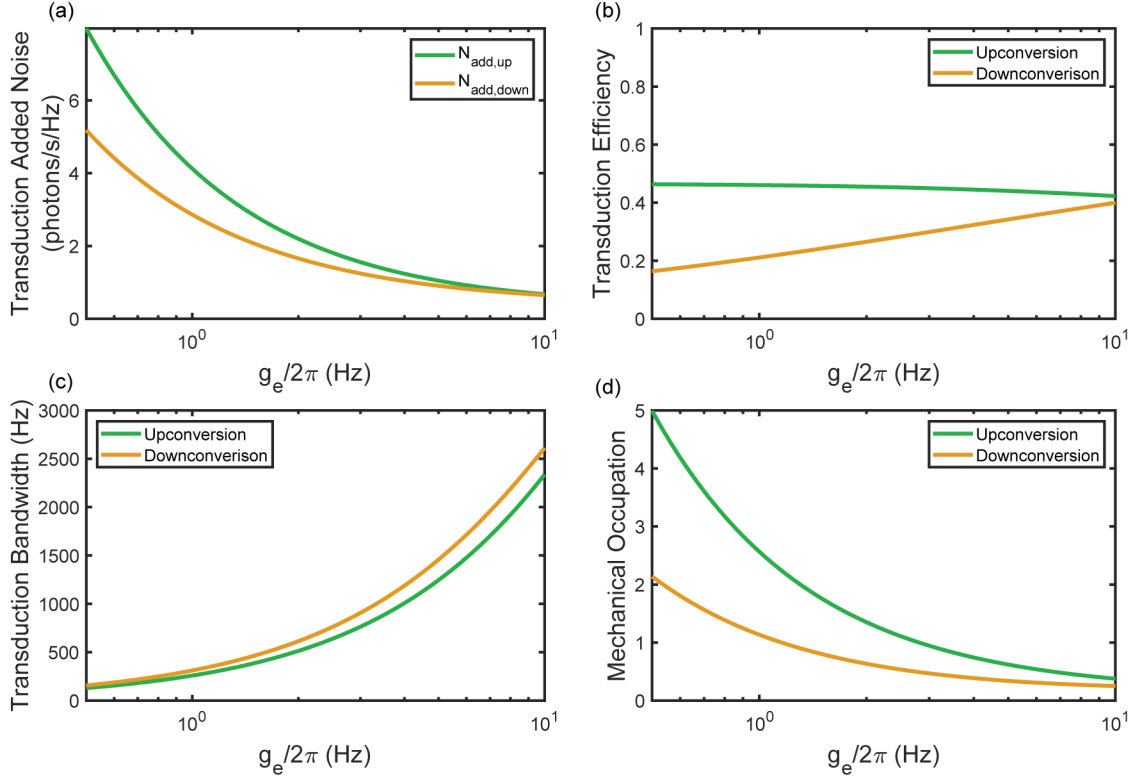


Figure 2.9: **Transduction performance versus electromechanical coupling** (a) Transduction Added Noise at optimal damping rates. Transduction efficiency (b), bandwidth (c) and occupation of the membrane (d) under the same parameters that minimize the added noise. Larger coupling between the microwave circuit and mechanical membrane improves all the transduction metrics.

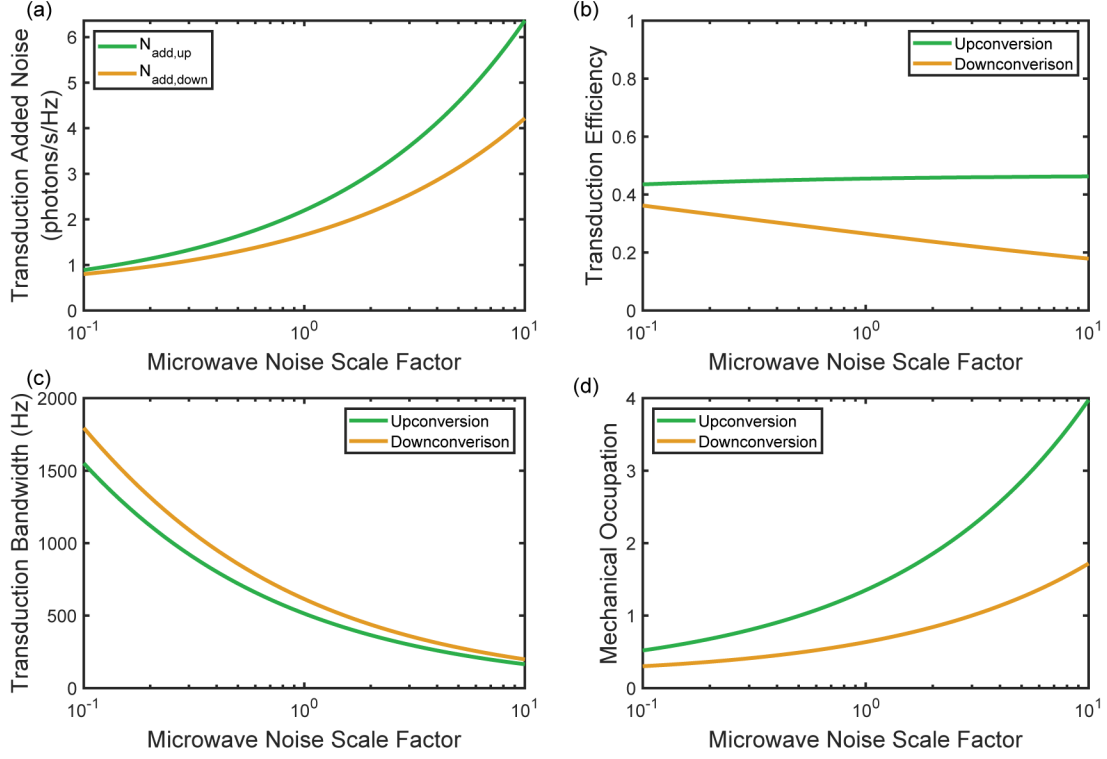


Figure 2.10: **Transduction performance versus microwave noise** (a) Transduction Added Noise at optimal damping rates. Transduction efficiency (b), bandwidth (c) and occupation of the membrane (d) under the same parameters that minimize the added noise. The x-axis is a scale factor multiplied to the linear parameterization of the $n_{\text{eff,e}} = a_e \bar{n}_e + b_e$

Chapter 3

Membrane mode engineering

Every flight begins with a fall.
- G.R.R. Martin

As the intermediary mode in our microwave-to-optical transducer, the mechanically-compliant membrane is at the heart of the experiment. The membrane is a highly-tensioned square sheet of Si_3N_4 whose lowest energy eigenmodes are at MHz-scale frequencies. This structure has approximately sinusoidal eigenmodes, and in the high tension limit, the eigenfrequencies are:

$$\omega_{l,n}^2 \approx \frac{T\pi^2}{\rho L^2}(l^2 + n^2) \quad (3.1)$$

where T is the film tension, ρ is the film density, and L is the side length of the square membrane. A silicon frame in the top chip of our flip-chip architecture supports the membrane on the edges, as shown in [Fig. 3.1](#). This style of mechanical oscillator has long been a tool for cavity optomechanical experiments [[Jayich et al. 2012](#); [Wilson et al. 2009](#)]. In this thesis, we work with the $(l, n) = (2, 2)$ mode of the membrane, with the optical spot and capacitor pad on separate corners of the membrane. The presence of a thin superconducting capacitor pad perturbs the mechanical mode shape, but the mechanical modes are still roughly that of a uniform thin square sheet.

In an ideal device, the desired mechanical mode is weakly coupled to its thermal environment. The motion of the eigenmode of interest is localized to the Si_3N_4 , with little motion in the mechanically lossy Si frame [[Wilson-Rae et al. 2011](#); [Yu et al. 2014](#)]. In any real device, the Si frame will have its own mechanical modes, which could hybridize with the membrane mode if there is significant




Figure 3.1: **Evolution of Membrane Chip.** Pictures of the transducer devices used in (a) [R. W. Andrews et al. 2014] (b) [Higginbotham et al. 2018], and (c) [Brubaker et al. 2022] with a penny for scale. The membrane is the green-tinted square in the center of the device. Patterning of the silicon frame (blue) that supports the membrane evolved to control the capacitor gap spacing and implement a mechanical filter.

motion at either the capacitor or optical spot. Any techniques for mechanically isolating the Si_3N_4 from its Si frame must maintain large electromechanical and optomechanical coupling while being robust to cryogenic operation.

In Sec. 3.1, we will detail the design and implementation of a mechanical filter in the Si frame to isolate a particular membrane modes. In Sec. 3.2, we will introduce a chemical bonding process which increased the reliability of the transducer’s construction, especially once cooled to mK temperatures. Both of these techniques were implemented in the devices discussed in Ch. 4. In the last section of this chapter (Sec. 3.3), we will demonstrate a new technique for achieving larger electromechanical coupling rates by reliably reducing the capacitor gap separation.

3.1 Phononic crystal

A phononic crystal (PNC) effectively implements a mechanical band stop filter by reducing the density of states for phonons in a desired frequency band [Yu et al. 2014; Kushwaha et al. 1993; Eichenfield et al. 2009; Hatanaka et al. 2014]. In analogy to a stepped impedance filter in the microwave regime, periodic modulation of the mechanical impedance can be engineered to prevent propagation of phonons within a desired frequency range. Implementation of a PNC in the Si frame that supports the membrane has two beneficial effects. By reducing the density of states near the

mechanical oscillator's frequency, the intrinsic mechanical dissipation is improved and the probability of hybridization with a mode of the substrate is reduced.

Mechanical dissipation in silicon nitride membranes has two main contributions. The first is bending or clamping loss, where energy is lost via higher order scattering processes at locations of high curvature. This is related to the intrinsic quality factor of the material, which is set by the imaginary part of Young's modulus. The second source is radiation loss, related to phonons entering the support structure. The quality factor of the mechanical resonator $Q_m = \omega_m/\gamma_m$ is given:

$$1/Q_m = 1/Q_{\text{bend}} + 1/Q_{\text{rad}} \quad (3.2)$$

where Q_{bend} is the bending-loss limited Q and Q_{rad} is the radiation-loss limited Q . Using the method introduced by [Reetz et al. 2019], the bending-loss limit Q for the $(l, n) = (2, 2)$ mode of a bare Si_3N_4 membrane with the same dimensions as our devices is $Q_{\text{bend}} \sim 4.5 \times 10^7$ assuming an intrinsic quality factor $Q_{\text{int}} = 6600$. Similarly, the radiation loss can be calculated using theory developed by [Wilson-Rae et al. 2011; Wilson-Rae 2008]. Taking the asymptotic limit of a semi-infinite substrate and square membrane [Villanueva and Schmid 2014], $Q_{\text{rad}} \approx \alpha(2.4 \times 10^7)$, where α is a correction factor to account for substrate imperfections result from chip mounting. Empirically, $\alpha \lesssim 0.1$ can explain the Q_m of a wide-range of experiments. For comparison, the Q_m in [R. W. Andrews et al. 2014] and [Higginbotham et al. 2018] was 3.1×10^5 and 1.3×10^5 , respectively. This suggests that Q_m was radiation-loss limited in those devices with $\alpha \sim 0.01$. The PNC directly improves the radiation loss by reducing the density of states into which phonons in the membrane can decay. Said another way, the PNC reduces the effect of chip mounting (increases α) by reducing the pathways for phonons in the membrane to see the mounting points.

The PNC will reduce the membrane mode occupation from thermomechanical noise, and also prevent thermomechanical noise from other mechanical modes from contributing to the transducer's added noise. These "spectator" modes are primarily modes of the Si frame that supports the membrane. If they have appreciable motion at either the capacitor pad or optical spot and their

frequency lies within the transducer's bandwidth, noise from the spectator modes can interfere with the noise from the membrane mode. We observed this exact effect in a transducer device with a Si frame structure similar to the 2018 device, shown in [Fig. 3.2](#).

The optically-detected thermomechanical spectrum is plotted in [Fig. 3.2a](#) for a fixed Γ_o and swept Γ_e . These spectrum follow the Lorentzian shape described by the optical version of [Eq. 2.24](#), with the peak height decreasing and linewidth increasing with large microwave power. For the largest value of Γ_e , the spectrum acquires a non-Lorentzian shape, and is best fit with the complex sum of two Lorentzians, one from the membrane mode and one from a spectator mode of the Si chip. Interference of these two noise sources artificially reduces the measured peak height, as shown in (b). Using COMSOL to simulate the eigenmodes of the full Si chip, we find the membrane mode and a likely candidate for the spectator mode in (c).

Importantly, the spectator mode has appreciable motion at the capacitor pad and optical spot and is spectrally close to the membrane mode. Consequently, thermal phonons from the spectator mode will couple to the microwave and optical resonators and can be indistinguishable from signal photons. We have found experimentally that g_e and g_o of the spectator modes are too small to be effectively cooled by the pumps, yet large enough to contribute to the transducer added noise. A PNC will reduce the density of states of spectator modes and reduce the electro/optomechanical coupling of the spectator modes that remain by isolating the membrane from the substrate.

3.1.1 Infinite PNC design

Optimization of the PNC design is performed in COMSOL, beginning with a unit cell. The band diagram of the unit cell can be calculated by assigning Floquet periodic boundary conditions and finding the eigenmodes of the infinite structure. The phase vector between adjacent cells is swept along the boundary of the first Brillouin zone [[Yu 2016](#)]. We use a square lattice to match the membrane symmetry, resulting in a network of alternating "pads" (thick sections) and "tethers" (thin sections). The effective length of these sections sets the center of the band gap, while the mass contrast between the pads and tethers defines the width of the band gap [[Reetz et al. 2019](#)].

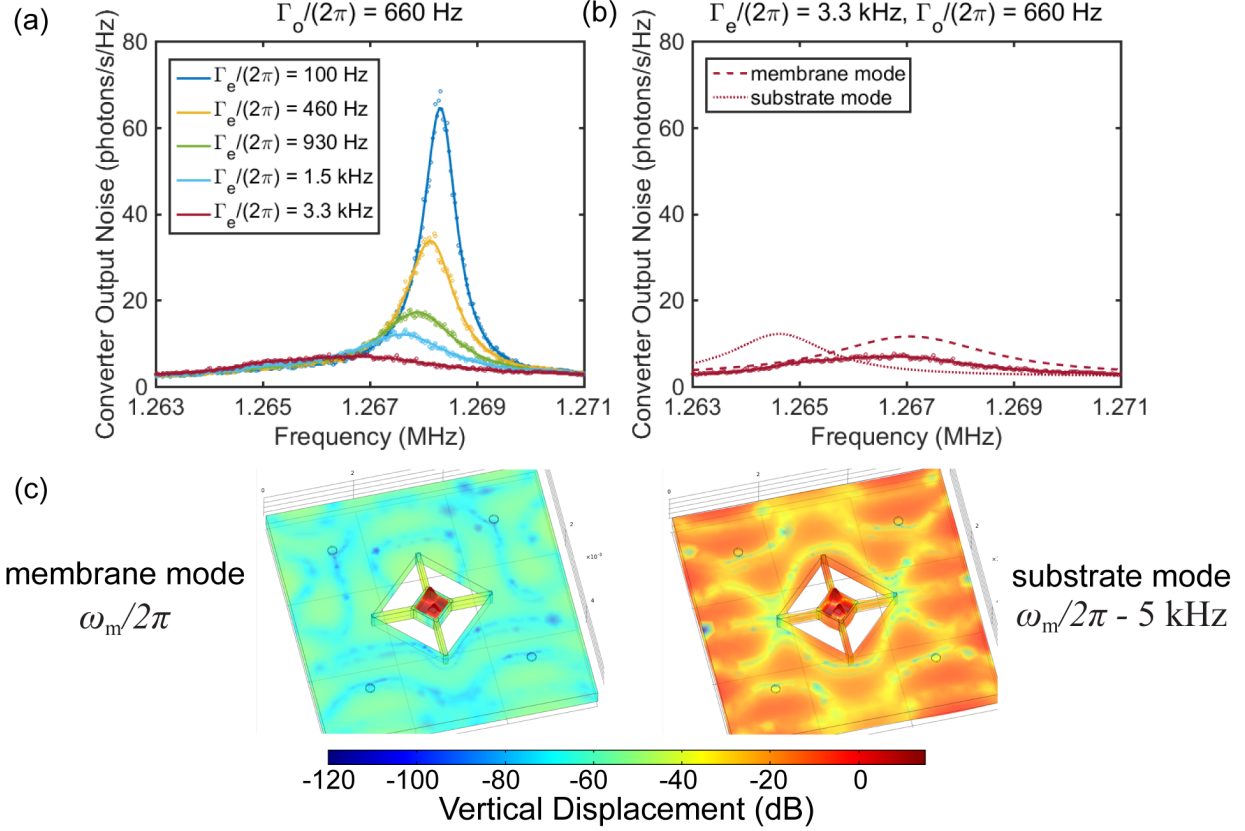


Figure 3.2: **Spectator modes of Si frame** (a) Optically-measured mechanical spectrum of the (2,2) mode for fixed optical damping. Data are in open circles and solid lines are a numerical fit. As the electromechanical damping rate increases, the thermomechanical peak gets broader and shorter, indicating cooling. At the highest powers, the peak acquires a non-Lorentzian lineshape. (b) This peak is best fit by the sum of two Lorentzians with opposite phase, indicating that noise from the membrane mode (dashed) is interfering with that of the substrate mode (dotted). (c) COMSOL simulations of the membrane chip, with the color scale representing a normalized vertical displacement. The most membrane-like mode (left) has the motion confined to the membrane but nearby substrate modes (right) have significant motion in both the membrane and Si chip.

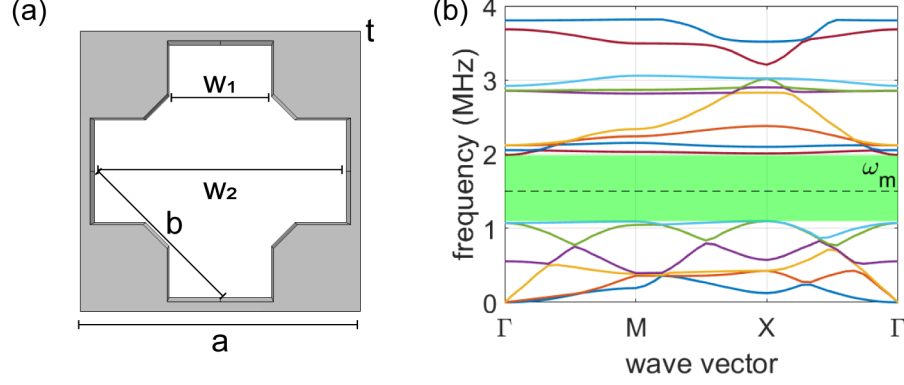


Figure 3.3: **PNC unit cell** (a) Dimensions of the optimized unit cell subject to fabrication constraints are $a = 1450 \mu\text{m}$, $w_1 = 550 \mu\text{m}$, $w_2 = 1350 \mu\text{m}$, $b = w_2\sqrt{2}$, $t = 380 \mu\text{m}$. The length b corresponds to the side length of a square cutout centered on the unit cell and rotated 45 degrees. (b) Band structure across the first Brillouin zone. The membrane mechanical mode will sit at the center of the band gap (green-shaded area)

Recalling the stepped impedance filter analogy, the mass contrast corresponds to an impedance mismatch between adjacent sections of the filter.

Constraints from the fabrication influence the optimal unit cell design (see [Appendix A](#)). The wafer from which the membrane chips are made is $380 \mu\text{m}$ thick. The minimum tether width that we can reliably fabricate is $100 \mu\text{m}$. Lastly, we could in principle add rather than remove material in order to create the necessary mass contrast. This has been successfully implemented for PNCs in the nitride membrane [[Høj et al. 2022](#); [Huang et al. 2024](#)], but it is difficult to achieve in Si PNCs due to the thickness of the wafer.¹ Within this design space, we are able to find a unit cell which has a ~ 1 MHz wide band gap centered at 1.5 MHz, as shown in [Fig. 3.3](#). The large bandgap ensures that small changes in membrane frequency or PNC fabrication do not ruin the isolation of the membrane.

3.1.2 Finite PNC effects

Due to the MHz-scale bandgap center frequency, the unit cell size is quite large, comparable to the side length of the membrane ($500 \mu\text{m}$). Only 1 to 1.5 unit cells physically fit into the membrane chip (see [Fig. 3.1 c](#)). The finite PNC cannot fully suppress all the spectator modes around the

¹One would need to deposit a high-density material in large amounts in order to create appreciable mass contrast.

membrane frequency, akin to a single-pole filter with limited extinction in stop band. This problem is exacerbated by uncertainty in the membrane and spectator mode frequencies due to variations in the fabrication ([Appendix A](#)). We cannot design the membrane and Si chip to deterministically avoid all the spectator modes.

To understand whether the finite PNC will effectively reduce the density of spectators modes, we develop a probabilistic figure of merit that accounts for a priori unknown membrane mode frequencies. This figure of merit will describe the likeliness of a frequency collision between a spectator mode and a membrane mode given some patterning of the Si chip and assuming a uniform probability distribution of the membrane frequencies within some range.

In COMSOL, we can simulate the eigenmodes of the Si chip around the membrane mode frequency. The coupling of each of these modes to the microwave circuit is encoded into an effective mass:

$$m_{\text{eff}} = \frac{\int_{\text{chip}} \rho(x) z^2(x) d^3x}{\langle z \rangle_{\text{pad}}^2} \quad (3.3)$$

where the integral in the numerator is over the whole membrane chip. $\rho(x)$ is the material density and $z(x)$ is the vertical displacement. We normalize by the squared-average vertical displacement of the capacitor pad $\langle z \rangle_{\text{pad}}^2$. Modes with significant (little) capacitor pad motion will have a small (large) effective mass and are easy (difficult) to cool with the radiation pressure force.

We can then calculate the electromechanical background thermal spectrum from these spectator modes. The thermomechanical position noise spectral density contribution from one spectator mode is [[Hauer, Maciejko, et al. 2015](#)]:

$$S_{\text{xx}}^{\text{th}}(\omega) = \frac{4k_{\text{B}}T\gamma_{\text{m,sub}}}{m_{\text{eff}}((\omega^2 - \omega_{\text{m,sub}}^2)^2 + (\gamma_{\text{m,sub}}\omega)^2)}. \quad (3.4)$$

From experimental observations, we find that the $Q_{\text{m,sub}} = 10^3 - 10^4$. We expect that a bath of low Q spectator modes is more likely to add significant noise to any given realization of the membrane frequency, so we assume that all spectator modes have $Q_{\text{m,sub}} = 10^3$.

We can use the classical noise spectral density for the spectator modes as they are not cooled

efficiently by the pumps and remain in the classical oscillator limit. To this background of classical oscillators, we add the contribution from our membrane mode when cooled to $n_m = 1$ at a total damping of $\Gamma_T = 1$ kHz, using the quantum position noise spectral density:

$$S_{xx}^{\text{quantum}}(\omega) = 2x_{\text{zpf}}^2 \frac{\Gamma_T(n_m + 1)}{\Gamma_T^2/4 + (\omega - \omega_m)^2} \quad (3.5)$$

where $x_{\text{zpf}}^2 = \frac{\hbar}{2m_{\text{eff}}\omega_m}$ is the zero-point fluctuations of the mechanical mode. This is one realization of the possible membrane and spectator mode spectrum. The spectator modes can change character if it hybridizes with the membrane. To capture that, we linearly sweep the membrane mode frequency in simulation between $0.9\omega_m \rightarrow 1.1\omega_m$ and generate many realizations of the total position spectrum. In each of these realizations, we integrate the motion in a $10\Gamma_T$ band centered on the membrane frequency to find how much of the motion is from spectator modes.

Our figure of merit is then the percentage of these realizations that result in more than 0.1 phonons of motion from the spectator modes (a tenth of the membrane occupation). This figure of merit is $\sim 40\%$ for the 2018 design, and reduced to $\sim 2\%$ for the optimized PNC design, more than an order of magnitude improvement. More details of this simulation procedure and application to other design challenges will be detailed in Kazemi Adachi's thesis.

3.1.3 PNC performance

With reasonable confidence in the design, membrane chips shielded by a PNC were fabricated. The process is akin to that used for the 2018 device [Burns 2019], with some changes to improve yield (Appendix A). The first round of devices with the PNC had sub-Hz intrinsic mechanical damping ($Q_m > 10^6$). In Fig. 3.4, we look more carefully at one of these devices, which was used for the work in [Brubaker et al. 2022] and [R. Delaney 2022].

We measure Q_m by performing mechanical ringdown measurements and fitting the exponential decay of the output amplitude to $e^{-\Gamma_T t/2}$. In the presence of a red-detuned pump, we create a large coherent state in the mechanical mode ("ring up") by applying an additional tone on-resonance with

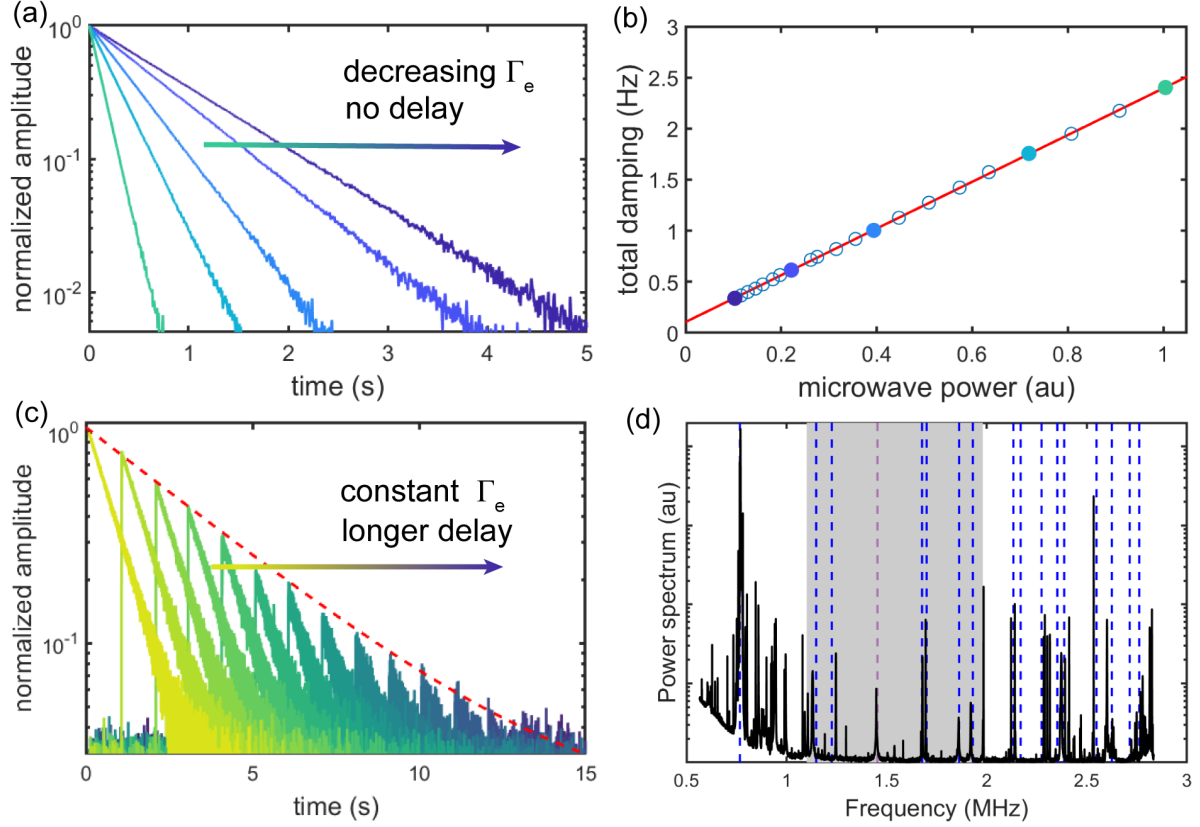


Figure 3.4: Improved Mechanical Properties (a) Mechanical Ringdown with no measurement delay. Each line corresponds to a different electromechanical damping rate. (b) Exponential decay rate extracted from mechanical ringdowns, with filled circles corresponding to data in the first panel. The y-intercept of the linear fit is the intrinsic mechanical damping rate γ_m . (c) Stroboscopic Mechanical Ringdown with constant damping. The delay between ring up and measurement is swept. Fitting the exponential decay envelope (red dashed line) of the amplitude measured at the beginning of the measurement also gives γ_m . (d) Electromechanical Spectrum. Many modes are visible in the spectrum (black line), including the (2,2) mode (gray dashed line), other membrane modes (blue dashed line), and some spectator modes. There are fewer spectator modes within the expected band gap (gray shaded region).

the cavity. The mechanical amplitude will decay at a rate Γ_T when the on-resonance tone is turned off at $t = 0$, as seen in Fig. 3.4a. Damping is linearly dependent on pump power (Eqn. 2.8), so the y-intercept of the fit to Fig. 3.4b will be γ_m . Note, this measurement is usually done purely electromechanically, as the optical measurement requires a lock tone which contributes its own pump power-independent damping rate. We find that $\gamma_m/2\pi = 113 \text{ mHz}$ ($Q_m = 1.28 \times 10^7$), two orders of magnitude larger than the previous designs. This is also approaching the limit $Q_{m,\text{max}} = 1.57 \times 10^7$ set by the bending loss and the radiation loss with ideal mounting ($\alpha = 1$), implying that the remaining mechanical loss is roughly equal parts bending and radiation.

As a sanity check, we implement another method of mechanical ringdown, shown in Fig. 3.4c. "Ringdown in the dark" or "Stroboscopic" readout turns off the red-detuned pump when the on-resonance tone is turned off. The mechanical oscillator then decays into its thermal bath at γ_m . After some delay, the electromechanical damping is turned back on, after which the oscillator amplitude will decay at Γ_T . The measured amplitude immediately after the damping is reapplied provides a snapshot in time of the intrinsic mechanical decay. By sweeping this delay, we can more directly access the intrinsic mechanical decay exponential. Using this technique, we find $\gamma_m/2\pi = 98 \text{ mHz}$ ($Q_m = 1.48 \times 10^7$), in slight tension with the previous measurement. For the subsequent analysis requiring this parameter, we use the first, more conservative measurement.

It is important to note that these measurements were made at the base temperature of our dilution refrigerator $T_{\text{bp}} = 40 \text{ mK}$. At this base plate temperature, the membrane has been observed to equilibrate to $T_{\text{eq}} \leq 100 \text{ mK}$ with the laser off, and at most $T_{\text{eq}} = 100 \text{ mK}$ with the laser on (see Fig. 2.5 or Fig. 8 of [Brubaker et al. 2022]). The mechanical quality factor can depend on temperature below $\leq 1 \text{ K}$, which has also been observed by [Yuan et al. 2015] and attributed to temperature-dependent surface losses [Villanueva and Schmid 2014]. In this device, we found that the mechanical loss increased linearly with temperature following $\gamma_m = a_\gamma T_{\text{bp}} + b_\gamma$ with $a_\gamma = 176 \text{ mHz/K}$ and $b_\gamma = 101 \text{ mHz}$. In a Sec. 5.6, we will explore how the mechanical dissipation at the milikelvin temperature could be due to the influence of two-level systems, the same model that explains the microwave circuit loss.

Lastly, we evaluate how the PNC affected the mechanical spectrum in [Fig. 3.4d](#). The mechanical spectrum is measured over a broadband that includes the designed mechanical bandgap, denoted by the gray shaded area. Outside of the PNC, there are many peaks not attributable to pure membrane modes (dashed blue lines), but within the bandgap, there are fewer of such non-membrane modes. In particular, there are relatively few substrate modes in the vicinity of the (2,2) mode (gray dashed line). We conclude that despite the relatively limited number of unit cells in the membrane chip, the PNC successfully reduces the prevalence of substrate modes in the bandgap and decreases the radiation loss of the membrane mode embedded in the structure.

3.2 Oxide-Oxide bonding

Construction of the transducers is a delicate process, prone to inconsistencies between each attempt. Thermal contraction during cooldown and unpredictable deformation of the cryogenic-compatible epoxy (Stycast 2850FT) used to affix components are two of the main sources of unreliability. The epoxy can change during thermal cycling, further exacerbating the issue. Control over the optical mode and the spacing of the electromechanical capacitor are two critical parameters that are challenging to optimize after the transducer has already been constructed. In principle, one could introduce piezoelectric actuators to control the mirror positioning and a DC voltage between the capacitor pads to control that spacing, but these additional tuning parameters introduce new avenues for vibrational noise to pollute the system.

Direct bonding of the constituent components can make the transducer robust to cryogenic thermal cycling by ensuring alignment of bonded chips and removing the need for epoxy. We use a hydrophilic oxide-oxide bonding chemistry to create intimate and robust contact between two silicon dioxide surfaces, which are either intentionally deposited or natively present on Si wafers [[Burns 2019](#); [Plöchl and Kräuter 1999](#)]. (more information on the bonding recipe can be found in [Appendix A](#)). As shown in [Fig. 3.5](#), we have explored this bonding for two applications: bonding the circuit chip to the mirror chip and/or bonding the membrane chip to the circuit chip. The first application ensures that the optical mode is always normally incident on both the membrane and flat mirror,




Figure 3.5: **Oxide-Oxide bonding for transducer construction.** Implementation of direct bonding to increase reliability of transducer construction. (a) Three-element cavity where each element (curved mirror, transducer chips, flat mirror) are epoxied (not shown) to separate mounting brackets and the membrane chip is epoxied (black ovals) to the circuit chip. (b) Two-element cavity where the flat mirror is bonded to circuit chip, reducing the likeliness of misalignment. (c) Fully bonded etalon, where the membrane chip is bonded to the circuit chip, removing epoxy from transducer construction.

which drastically reduces light scattering [Urmeý 2024]. The second removes the need for epoxy in construction of the electromechanical system.

The main challenge in realizing these applications is ensuring that the bonding does not disrupt the mechanical or electrical transducer components. For example, the capacitor gap spacing is very sensitive to deformation from the mismatched thermal contraction of SiO₂ and Si. If the circuit and membrane chips bow apart, the electromechanical coupling will decrease ($g_e \propto 1/d$) and the microwave circuit frequency could end up out of the band of our measurement chain ($\omega_e \propto \sqrt{d}$). We have also observed cases where the capacitor pads touch, either attracted by electrostatic forces, the Casimir effect, or simply bending of the two chips. This "collapsed" capacitor results in an unusable device, as there is no longer an electromechanical circuit. The issues can arise in both bonded and epoxied devices, but the bonding is a lithographically defined process with precise control. Design of the location and size of the oxide contact points presents an opportunity to control the capacitor spacing in the presence of deformation due to thermal contraction.

3.2.1 Circuit chip-mirror chip bonding

The flat mirror is a dielectric stack of tantala and silicon dioxide deposited wafer scale on Si by Five-Nine Optics [*FiveNine Optics – Boulder, Colorado / $R > 99.999\%$ 2024*], making a Bragg reflector with high reflectivity. The total stack is a few μm thick, with a 10 nm capping layer of SiO₂. In order to achieve the highest quality mirrors, the deposition occurs at several hundred Celsius. When cooled to room temperature, mismatched thermal expansion coefficients creates stress in the mirror coating, which will result in an appreciable wafer bow. The bonding process requires the two surfaces to have intimate contact, generally requiring flat substrates. The stress of the mirror coating is compensated by a 3 μm SiO₂ stress compensation layer also deposited at high temperature on the backside of the wafer, so that the mirror wafer is flat at room temperature, as shown in Fig. 3.6a.

The additional stress compensation layer effectively creates a symmetric material stack of oxide-silicon-oxide. Even though the materials are still mismatched, symmetry of the structure ensures that the wafer remains flat at all temperatures. After we bond the circuit wafer to the mirror

wafer, we break this symmetry. When the new bonded stack is cooled to cryogenic temperatures, the stack will again have a bow (Fig. 3.6b). This led to unpredictable capacitor gap spacings in the first devices with bonded mirrors.

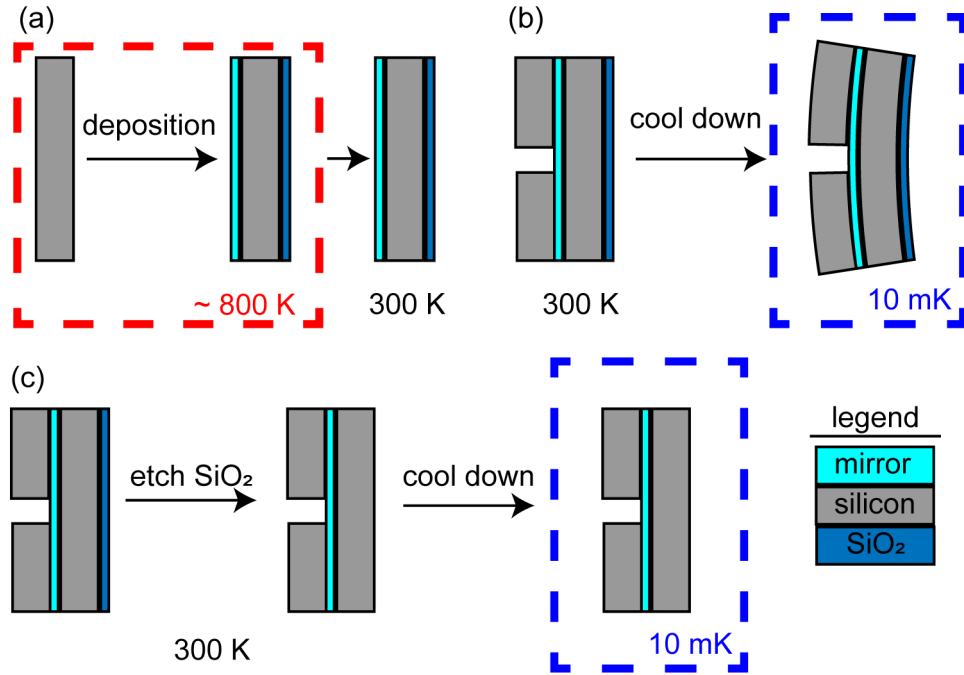


Figure 3.6: **Stress compensation of mirror coating** (a) A layer of SiO₂ (purple) deposited at high temperature counteracts the film stress of the mirror (cyan) coating at room temperature. (b) After bonding the circuit chip, with a hole in the center for optical access, the stack is no longer symmetric and bends when cooled to cryogenic temperatures. (c) Removing the stress compensation after bonding the circuit chip recreates a symmetric stack, which does not bend when cooled. Dimensions are not shown to scale.

We tried chip scale bonding the circuit chips to fused silica substrates, mirror chips of different side lengths, and even mirror chips diced from a 1 inch thick Si puck, hoping to better control the gap spacing by changing the materials and dimensions of the mirror substrates. Of these options, using 1-2 mm side length mirror chips (compared to the 9 mm side length circuit chip) proved most successful. The mirror chip is then a perturbation on the circuit chip, and the bowing is reduced. Smaller mirrors, however, have less bonding area and are thus less resistant to cryogenic cycling and routine handling. Additionally, chip scale bonding sacrifices the scalability of wafer scale processing. We ultimately abandoned these chip scale methods, instead developing a wafer scale method which

uses the same symmetry principles as the SiO₂ stress compensation layer.

If we can remake a symmetric material stack, we expect to minimize the deformation of the chip. In [Fig. 3.6c](#), this can be done by removing the SiO₂ stress compensation after bonding the wafer of circuit chips to the wafer of mirror chips. The material stack is now symmetric, and this etch can be done wafer scale with a one-sided anhydrous HF vapor for 90 minutes ([Appendix A](#)). This process minimizes the chip deformation from thermal contraction of the mirror coating without sacrificing wafer scale processing and robust bonding.

In addition to improving the reliability of the capacitor gap spacing, bonding the mirror chip and circuit chips reduced light scattering out of the optical mode, which improved $\kappa_{\text{int,o}}$, and reduced laser heating of the superconducting circuit ([Sec. 4.2](#)). The optical cavity could also be shortened due to the compactness of new cavity design, which improved the vacuum optomechanical coupling rate g_o by concentrating the light in the shorter section of the asymmetric cavity. These improvements are detailed in Maxwell Urmeý's thesis [[Urmeý 2024](#)]. Finally, there are also prospects to bond the curved mirror to the membrane chip using micro-fabricated mirrors [[Jin et al. 2022](#)], which could enable a fully-bonded, epoxy-free cavity construction.

3.2.2 Membrane chip-circuit chip bonding

Recently, we have attempted to extend the oxide process to construction of the electromechanical circuit. In the device used in [[R. Delaney 2022](#)] and [[Brubaker et al. 2022](#)], the gap spacing between the capacitor pads was determined by aluminum structural posts with the desired height on the circuit chip. We can replace each of these posts by a pair of posts, one on the circuit chip and one on the membrane chip, whose combined heights sum to the desired capacitor gap spacing. The chips are then held together by oxide bonding of those sets of posts.

An important design parameter is the size and shape of the bonding posts. Too large, and there is a high likeliness for dust or other contaminants to mar the bonding surface. Too small, and the bond between the chips may not survive cryogenic operation or routine handling. We found that 4 posts outside of the phononic crystal with a radius of 100-200 μm was a reliable design, achieving




Figure 3.7: **Bonded transducer etalons** (a) Design of the membrane chip. The oxide bonding posts are the circles, four outside the PNC and four on the membrane frame whose positions have been optimized with the process described in [Sec. 3.1.2](#). (b) Picture of two fully bonded transducer etalons, with the membrane, circuit and mirror chips all only connected via oxide bonding.

sufficient bond strength to survive multiple cryogenic cycles. The design of the membrane chip and a photo of two fully-bonded transducer etalons (every component except the curved mirror) are shown in [Fig. 3.7](#).

Initial tests of the fully bonded etalons were promising, yet inconsistent. We achieved gap spacings of 200 nm while maintaining $Q_m > 10^6$ with three devices. This was still relatively small subset of the total number of devices constructed, as most of the capacitor gaps ended up either too large to have a resonance in band of the microwave measurement chain or collapsed (10 such devices). COMSOL simulations of the thermal contraction of the structure suggest that the spacing should change by less than 1 nm between 300 K and 77 K, after which there should be negligible additional deformation. The capacitor gap is also influenced by the mounting of the device to the chip holder and any dust between the membrane and circuit chips. It is also possible that our simulations do not accurately capture the deformation from thermal contraction. We have some intuition of the deformation from thermal contraction of epoxied devices ([Appendix A](#)), but better understanding of the oxide bonding posts must be developed before it can become a more reliable construction method.

3.3 Pinned membrane

As discussed in the previous section, the capacitor gap separation is a difficult length to control with our flip chip architecture. Some of this arises from the length scales involved. In the device studied in [Brubaker et al. 2022] and [R. Delaney 2022], the structural posts which define the chip spacing are nearly 2.5 mm away from the capacitor gap which is designed to be 300 nm. Less than a milliradian of angular displacement is needed to explain the eventual 830 nm spacing. This distance is required for the posts to be outside of our limited PNC, as they could act like mechanical shorts to the circuit chip if placed inside the PNC. We find, however, posts placed on nodes of motion of the substrate or even the membrane can circumvent this limitation, as effectively the phonons do not "see" these posts if they are properly placed.

An early implementation of this can be seen in Fig. 3.7a, as the posts on the frame of the membrane are designed to line up with the nodes of the membrane motion. More recently, we believe that posts on the center of the PNC pads (large squares) or tethers (skinny lines) will be more effective. While these posts are closer to the capacitor than the posts outside of the PNC, the closest we can place a post is directly on the membrane. This idea was inspired by an optical measurement of $Q_m > 10^6$ of the modified membrane modes in a defective transducer where the capacitor pad was collapsed [Urmey 2024]. That measurement suggested that rigidly clamping part of the membrane was not necessarily catastrophic for the loss of the now modified mechanical modes, and the frequencies can be predicted by COMSOL simulations. In the remainder of this section, we will detail the design of the pinning post and measurements of the first iteration of pinned electromechanical devices. For clarity, we note that the transducers detailed in most other sections of this thesis do not have a pinned membrane, as it is a relatively new development²

²The annealed transducer discussed in Ch. 6 does have a pinned membrane. In this thesis, we only report the microwave properties of that device as full characterization is still underway.

3.3.1 Pinning post design

Intuitively, the location for the pinning post that is least perturbative to the mechanical mode is along a nodeline of motion of the unpinned mechanical mode shape. Additionally, the post should be as large as possible to create reliable contact between the membrane and circuit chips while minimally changing the mechanical mode coupling and dissipation. To confirm this intuition, we can use COMSOL to simulate the electro/optomechanical coupling rates relative to the unpinned mode, the bending loss limited Q , and (2,2) mode eigenfrequency. These simulations include only the membrane, capacitor pad, and pinning post with fully fixed boundary conditions along the perimeter of the membrane and the surface of the pinning post not contacting the membrane. We ignore any potential effects of the Si chip so we can quickly iterate on the design.

COMSOL will return eigenvalues and Q_{bend} is calculated using known material properties and the mechanical mode shape [Reetz et al. 2019]. The electro/optomechanical coupling rates are written:

$$g_e = G_e x_{\text{zpf},e} \mathcal{E}_{\text{mot}} \quad (3.6)$$

and

$$g_o = G_o x_{\text{zpf},o}. \quad (3.7)$$

G_e (G_o) have units of [frequency/distance], describing the susceptibility of the corresponding cavity frequency to motion of the membrane. The electrical case also includes a participation ratio \mathcal{E}_{mot} , which is the ratio of the mechanically-susceptible capacitance to the total capacitance of the circuit. These quantities will not be influenced by the design of the pinning post as they are determined by the design of the electromagnetic cavities. The pinning post will change the amount of motion at the capacitor pad or optical spot, which is represented by $x_{\text{zpf},e}$ and $x_{\text{zpf},o}$ respectively and can be calculated from the mechanical mode shape.

In our design process, we compared the coupling rates by comparing the zero-point motions of a given design to that of the unpinned mechanical mode shape. $x_{\text{zpf},e}$ is evaluated over the capacitor

pad using the effective mass defined by Eqn. 3.3. $x_{zpf,o}$ is evaluated similarly over the optical spot, which we choose as a $30\mu\text{m}$ radius circle centered at the unmetalized anti-node of the unpinned mechanical mode. This choice does not account for the Gaussian intensity distribution of the beam or movement of the optimal coupling spot in the pinned designs. The *relative* optomechanical coupling is insensitive to the choice of uniform beam profile, and we observe little shift of the optimal coupling spot compared to the beam waist in the designs considered here.

Fig. 3.8 reviews the simulation of these parameters to optimize the post location and radius. The post location sweep is performed with a $10\mu\text{m}$ radius post along the southwest to northeast diagonal, which maintains the reflection symmetry of the unpinned membrane design. The couplings are least affected when the post is between $-70\mu\text{m}$ and $-80\mu\text{m}$ from the center of the membrane which, as expected, is along the nodal line of the unpinned mode shape. This is also where the mechanical frequency is least perturbed. At all post positions, the post necessarily creates additional curvature in the mode shape and Q_{bend} is reduced from the unpinned case. Q_{bend} is most affected when the post approaches an anti-node of motion ($-100\mu\text{m}$), but the optimal post location avoids much of the additional loss.

The post radius is swept with the post defined in its optimal position along a nodal line of motion, as seen in Fig. 3.8d. Naturally, as the post size approaches zero, the design parameters asymptotically approach the unpinned values. It is unlikely that a $1\text{-}2\mu\text{m}$ radius post would reliably bond or that the chips could be aligned to that accuracy. In flip chip assembly, we find that the membrane and circuit chip can be reliably aligned to within $10\mu\text{m}$ with our non-computer assisted alignment (Appendix A). We thus chose a $10\mu\text{m}$ radius as compromise between perturbing the membrane mode and feasibility of construction.

There is a large design space unexplored in this thesis. For example, we could consider if the membrane could be pinned in such a way to increase the electromechanical coupling. In Fig. 3.8e, we see that for larger post radii, the electromechanical coupling increases at the expense of the optomechanical coupling. Clamping more of the membrane concentrates the motion at the capacitor pad because the post is closer to the capacitor pad than the optical spot. In the device

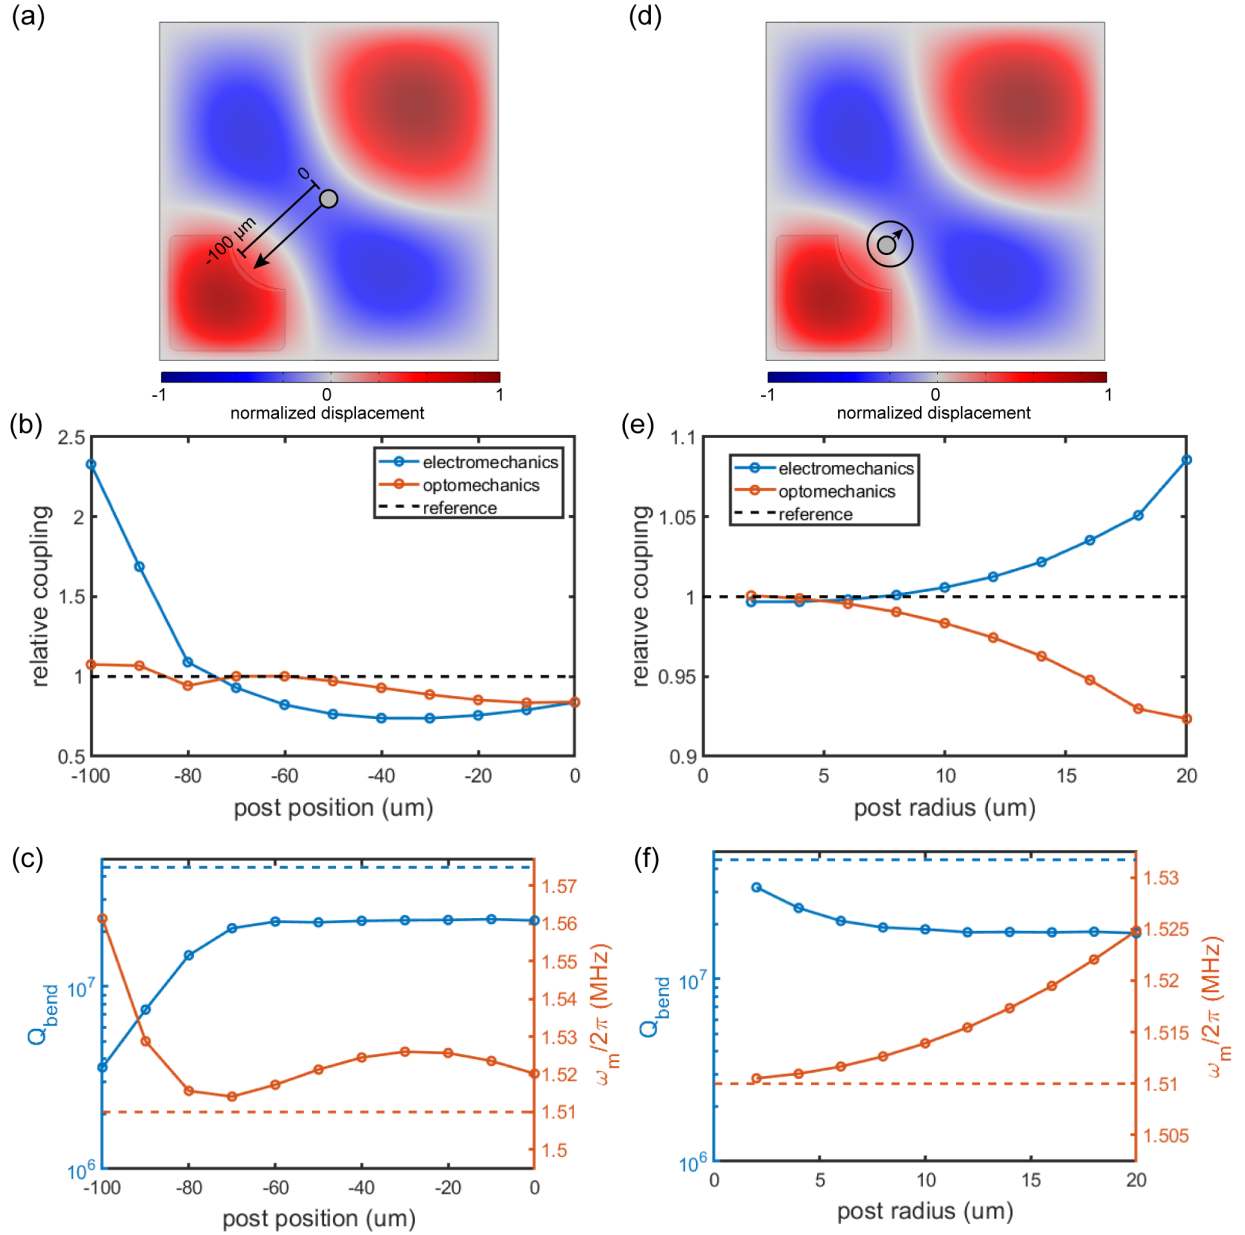


Figure 3.8: **Pinning post design** (a) Location sweep of pinning post with radius $10\ \mu\text{m}$, overlaid on the unpinned (2,2) mechanical mode shape. The post location is swept along the diagonal of the membrane, starting at the center and approaching the capacitor pad. (b) Relative electromechanical/optomechanical coupling as a function of post location. (c) Q_{bend} (blue) and ω_m (orange) as a function of post location, with dashed lines representing the simulated unpinned values. (d) Post radius sweep at optimized post location of $-75\ \mu\text{m}$, which lies along a nodal line of motion. (e-f) Relative coupling, Q_{bend} , and ω_m as a function of post radius, with dashed lines representing the simulated unpinned values.

measured in [Brubaker et al. 2022] and [R. Delaney 2022], $g_o \approx 40 \times g_e$. As those experiments were limited by microwave noise which scaled with microwave pump power, it could be advantageous to trade optomechanical coupling for electromechanical coupling. Additionally, one could consider more extensive pinning structures that could define exotic membrane mode shapes beyond that which is capable with our fabrication process. In this thesis, we restrict our experimental attempts to the more limited design space and fabricate devices with the design optimized in Fig. 3.8.

3.3.2 Experimental realization

Using our optimized design, we fabricate electromechanical etalons with the membrane, circuit, and mirror chip with the pinning post. In addition to the pinning post, these designs also had four posts outside the PC and four posts on the membrane frame optimized to be along the nodelines of the membrane motion, similar to the design in Fig. 3.7a. Some etalons were fully bonded and some still included epoxy between the membrane and circuit chips. A optical photograph looking through the membrane at the capacitor pad and pinning post is shown in Fig. 3.9. The pinning is the small dark dot within other concentric circles just above the capacitor pad. Using a white-light spectrophotometer (Appendix A), we measure the room-temperature capacitor gap spacing to be 200 nm, matching the height of the posts. We also infer that there is indeed contact between the membrane and circuit chip at the pinning post by the darkness of the spot.

Upon cooling down this device, we find a resonance frequency indicating that the capacitor gap spacing is still around 200 nm. Comparing the simulated mode frequencies to the measured mechanical mode spectrum in Table 3.1, the membrane is likely to pinned, especially considering the shifted fundamental frequency. The frequency of the fundamental mode is sensitive to pinning because the pinning post is not on node of motion for this mode. It is also generally easy to identify the fundamental mode in the spectrum as the lowest frequency mechanical mode. Assigning mode numbers to higher frequency modes is complicated by the fact that not all modes are well coupled to the microwave circuit.

Most importantly, in Fig. 3.9, a electromechanical damping sweep of the (2,2) mode extracts

Mode Number	Simulated (unpinned)	Simulated (pinned)	Measured
(1,1)	714 kHz	830 kHz	815 kHz
(1,2) or (2,1)	1.08, 1.15 MHz	1.15, 1.18 MHz	995, ? kHz
(2,2)	1.38 MHz	1.39 MHz	1.34 MHz

Table 3.1: **Pinned membrane mode frequencies.** Simulated and measured membrane mode frequencies for the four lowest order modes. The frequency of the fundamental mode is a good indicator of whether or not the membrane is pinned as it exhibits a large frequency shift upon pinning. As designed, the (2,2) mode is largely insensitive to pinning. One of the (1,2) or (2,1) mode has vanishingly small electromechanical coupling, so we do not observe it in the measured mechanical spectrum. The simulation parameters are a membrane side length of $550\text{ }\mu\text{m}$, a Si_3N_4 stress of 1 GPa, density of 3100 kg/m^3 , Young’s modulus 250 GPa, and Poisson’s ratio 0.27.

a vacuum electromechanical coupling of $g_e/2\pi = 4.98\text{ Hz}$ and an intrinsic mechanical damping rate of $\gamma_m/2\pi = 0.103\text{ Hz}$, both consistent with devices with similar spacings that are not pinned. This corresponds nearly an order of magnitude increase in the electromechanical cooperativity at all pump powers. Even though the membrane is held fixed at the point defined by the post, the coupling and loss of the mechanical membrane mode is largely unaffected. Optimizing g_e is especially important (Fig. 2.9) due to the large microwave pump-induced noise which has limited our transducers (Ch. 4).

This device represents one of the many electromechanical etalons assembled and measured. Many of the devices which did not end up pinned had a spacing indicating that the membrane and circuit chips had bowed apart. Work continues on understanding the thermal contraction of this transducer architecture. For example, we found that adding stycast between the membrane and circuit chips to already fully bonded devices was effective in decreasing the room-temperature and cryogenic capacitor gap spacing. To ensure pinning, we fabricated some devices where the pinning post was 10 nm taller than the other structural posts, hoping to slightly "tent" the membrane. Measurements of these devices found generally reduced $Q_m \approx 3.5 \times 10^5$, but it is unclear if the additional mechanical loss arises from the additional membrane curvature enforced by the taller pinning post.

In a small subset of the measured pinned devices, we observed a Duffing-type non-linearity in the electromechanical response [Hamel 1921; Korsch et al. 2008; Fong et al. 2012; Catalini et al.

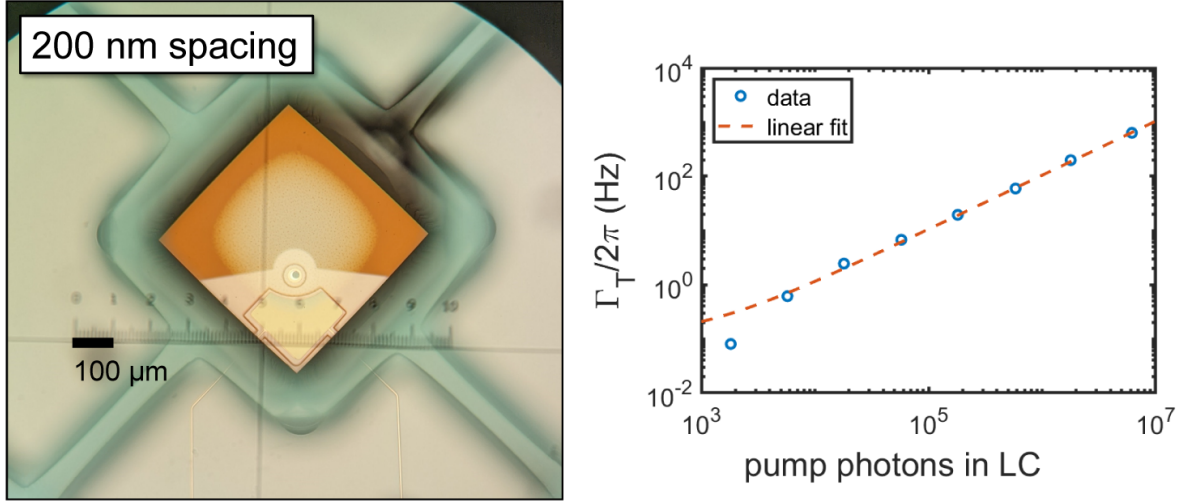


Figure 3.9: **Electromechanical properties of pinned membrane** On the left, a room-temperature optical micrograph looking from above the membrane chip, down through the transparent membrane onto the capacitor pads. The membrane is square in the center of the image, and came out as a $550\ \mu\text{m}$ square in this device. The apparent coloration of the capacitor pad and pinning post suggest that there is contact at the post (confirmed by a broad reflection spectrum measurement from optical to mid-IR frequencies). On the right, a electromechanical damping sweep. From a linear fit, we extract $g_e/2\pi = 4.98\ \text{Hz}$ and $\gamma_m/2\pi = 0.103\ \text{Hz}$.




Figure 3.10: **Duffing non-linearity of some pinned membranes.** Magnitude (top) and phase (bottom) response of electromechanically-induced transparency with the probe tone either swept upwards (blue) or downwards (red) in frequency. The non-Lorentzian, "shark-fin" type lineshape as well as the asymmetry between probe sweep directions is characteristic of a hardening Duffing non-linearity.

2021], shown in Fig. 3.10. In silicon nitride membranes, this behavior is due to nonlinear terms in the restoring force which generally become important when the amplitude of the membrane motion is on the order of the membrane thickness. In these measurements, however, the membrane motion is being damped to amplitudes in the picometer range, much smaller than the 100 nm membrane thickness. We suspect that this effect could arise from nonlinear restoring forces at the pinning post, especially if the membrane and post are not infinitely rigidly connected.

Despite these complications, we successfully fabricated and measured electromechanical etalons where the membrane is pinned on a node of motion of the (2,2) mode, with minimal degradation to the properties of the mechanical membrane. With additional exploration, the pinned membrane designs could be implemented in a full transducer, ideally resulting in reliable capacitor gap spacings. There is even hope to push to gap spacings of 100 nm or smaller ($g_e \propto 1/d$), which would further relax the large electromechanical pump powers needed in previous realizations of our transducer. One could also consider other bonding materials, such as superconducting metals which would create a galvanic connection between the membrane and circuit chips.

Chapter 4

Transducer Demonstrations and Limits

Anything you dream is fiction, and anything you accomplish is science, the whole history of mankind is nothing but science fiction.

- Ray Bradbury

At the beginning of 2020, we characterized a transducer device that incorporated several improvements from the transducer described in [Higginbotham et al. 2018]. While that this transducer did not achieve quantum performance, we demonstrated some prerequisites for quantum-enabled networking of superconducting qubits connected by optical fields. The largest roadblocks to $N_{\text{add,up}}, N_{\text{add,down}} < 1$ was also identified: microwave-pump power dependent loss and noise in the superconducting circuit.

In this chapter, I will describe these demonstrations, with focus on how those experiments were limited by the negative effects of the microwave pump. We demonstrated co-operation of a superconducting qubit and the transducer (Sec. 4.1) by optically reading out the qubit via upconversion of the classical microwave readout pulse with minimal backaction on the qubit (described thoroughly in Rob Delaney’s thesis [R. D. Delaney, Urney, et al. 2022; R. Delaney 2022]). This transducer also demonstrated the lowest added noise to-date for a mechanically-mediated, doubly-parametric architecture (Sec. 4.2) and ground-state cooling of the mechanical mode with the optomechanical interaction (described thoroughly in Max Urney’s thesis [Urney 2024] and [Brubaker et al. 2022]).

Microwave-pump induced noise and loss reduced the electromechanical cooperativity and transducer efficiency and increased the added noise, which limited the quantum efficiency of the

qubit readout and prevented quantum-enabled operation of the transducer (Fig. 2.10). In Sec. 4.3, we explore potential sources of the pump-induced microwave loss, with the expectation that the loss and noise arise from the same source. We find that the loss is ultimately due to electric field participation in the Si_3N_4 which composes the membrane.

The transducer used for the two demonstrations discussed in this chapter is shown schematically in Fig. 4.1, with insets showing more detailed depictions of the transducer components. The microwave resonator is a superconducting lumped element LC circuit made from NbTiN (blue), inductively coupled to a transmission line. One of the circuit’s capacitor pads is fabricated on the Si_3N_4 membrane, providing electromechanical coupling. The membrane (purple) is supported by a silicon frame which is patterned into a phononic crystal (see Sec. 3.1). The optical cavity is formed by a curved mirror on a silica substrate and a flat mirror on a silicon substrate, the latter of which is bonded to the circuit chip (see Sec. 3.2). The cavity mode (red) passes through the corner of the membrane opposite the capacitor pad, providing optomechanical coupling. The transducer is connected to the microwave measurement chain by wirebonds to a flexible circuit board, which is connectorized with an SMA. Free-space optics couple light into the transducer via a window in the side of the dilution refrigerator.

The microwave pump-power-dependent loss and noise were exacerbated by a relatively small electromechanical coupling rate $g_e/2\pi = 1.6$ Hz. This device had smaller coupling than other devices made around that time due to the large capacitor gap spacing of ≈ 830 nm, compared to the designed 300 nm separation. This particular transducer etalon was chosen from a selection of similar devices for having the least power-dependent loss in preliminary electromechanical measurements as well as large Q_m . See Table 4.1 for a list of all the parameters of this transducer.

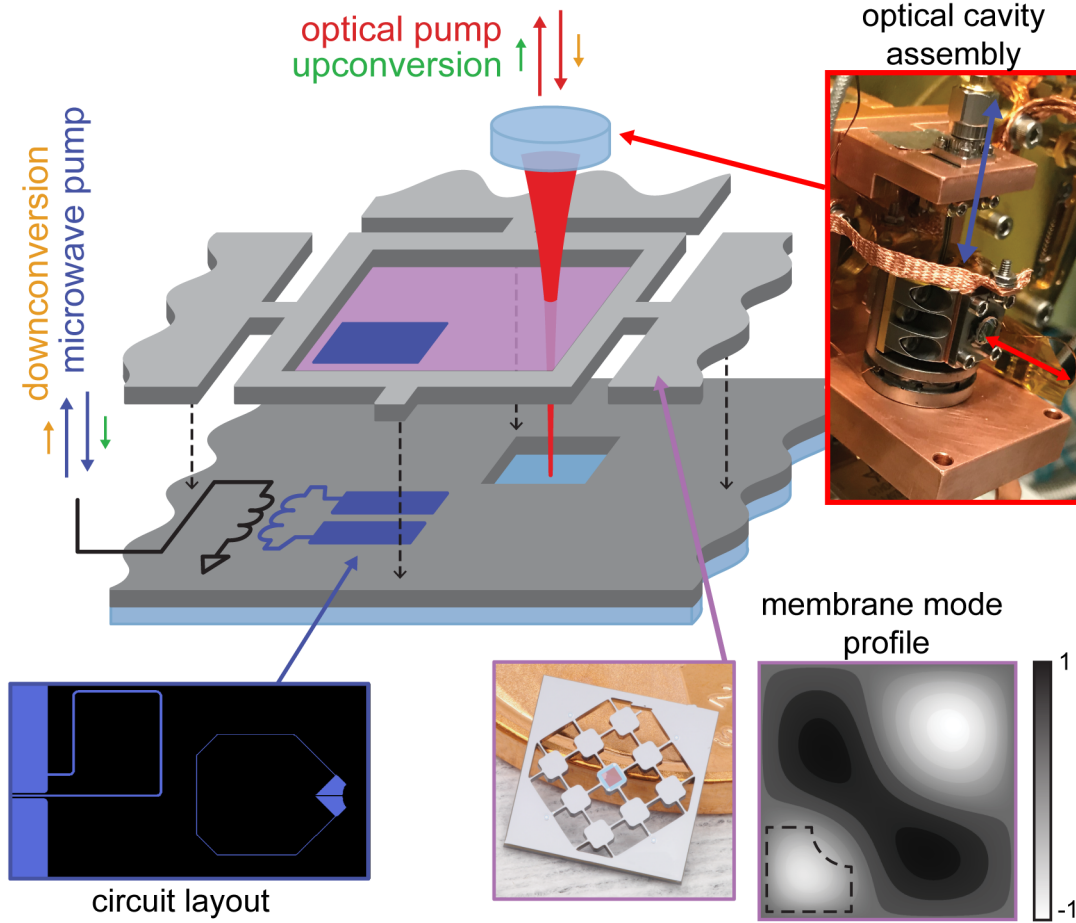


Figure 4.1: **Transducer Description.** The design of the transducer characterized in this chapter, with more detailed images in the insets. The superconducting circuit is made from NbTiN (blue) on a Si (dark grey) substrate (circuit chip). The Si_3N_4 membrane (purple) is held by a phononic crystal in membrane chip (light grey). The optical cavity is formed between a curved mirror on a silica substrate and a flat mirror (mirror chip) bonded to the circuit chip. Microwave signals reach the chip via SMA cables and optical signals are coupled via a free-space lens (blue and red arrows in the top right inset).

Parameter	Symbol	Value
Optical cavity frequency	ω_o	$\omega_o/2\pi = 277$ THz
Optical cavity external coupling	$\kappa_{\text{ext},o}$	$\kappa_{\text{ext},o}/2\pi = 2.12$ MHz
Optical cavity linewidth	κ_o	$\kappa_o/2\pi = 2.68$ MHz
Microwave circuit frequency	ω_e	$\omega_e/2\pi = 7.938$ GHz
Microwave circuit external coupling	$\kappa_{\text{ext},e}$	$\kappa_{\text{ext},e}/2\pi = 1.42$ MHz
Microwave circuit linewidth	κ_e	$\kappa_e/2\pi = 1.64 - 2.31$ MHz
Mechanical mode frequency	ω_m	$\omega_m/2\pi = 1.451$ MHz
Intrinsic mechanical dissipation rate	γ_m	$\gamma_m/2\pi = 113$ mHz
Vacuum optomechanical coupling	g_o	$g_o/2\pi = 60$ Hz
Vacuum electromechanical coupling	g_e	$g_e/2\pi = 1.6$ Hz
Modematching of pump beam to cavity mode	ϵ_{PC}	$\epsilon_{\text{PC}} = 0.86$ (0.80)
Modematching of cavity mode to LO beam	ϵ_{CL}	$\epsilon_{\text{CL}} = 0.91$ (0.79)
Bidirectional modematching	$\epsilon = \sqrt{\epsilon_{\text{PC}}\epsilon_{\text{CL}}}$	$\epsilon = 0.88$ (0.79)
Matched efficiency (peak efficiency)	η_{M} ($\eta_{\text{t,max}}$)	$\eta_{\text{M}} = 55\%$ ($\eta_{\text{t,max}} = 49\%$)
Optomechanical transducer gain	\mathcal{A}_o	$\mathcal{A}_o = 1.22$
Electromechanical transducer gain	\mathcal{A}_e	$\mathcal{A}_e = 1.08 - 1.16$
Lock backaction	$\gamma_{\text{lock}}n_{\text{lock}}$	$2\pi \times 40$ photons/Hz
Optical technical noise	$n_{\text{eff},o} = a_o\Gamma_o$	$a_o = 2.8 \times 10^{-6}$ photons/Hz
Microwave technical noise	$n_{\text{eff},e} = a_e\bar{n}_e + b_e$	$a_e = 2.93 \times 10^{-8}$ $b_e = 0.09$ photons

Table 4.1: **Electro-optic transducer parameters.** The range of values for the microwave cavity linewidth κ_e and electromechanical gain \mathcal{A}_e reflect the pump-power dependence of the microwave loss, taking the cavity from overcoupled to near critically coupled. The optical cavity modematchings and transducer efficiencies represent either operation with maximum efficiency (first number) or lowest added noise (second number).

4.1 Optically-detected superconducting qubit readout

In order to realize a quantum network architecture of superconducting qubits connected by optical fibers, quantum signals from the qubits need to be upconverted to THz frequencies without imparting backaction on the qubit. The first demonstration of quantum transduction suffered from optical heating of the qubit, destroying the state after measurement and necessitating pulsed operation of the optical pump [Mirhosseini et al. 2020]. While our transducer does not yet quite operate past the quantum threshold, we show that we can perform readout of the state of the qubit with minimal backaction by upconverting and optically detecting the classical microwave readout pulse without the need to pulse the conversion pumps. The ultimate quantum efficiency of the readout, however, is limited by the microwave-pump-power dependent loss and noise in the transducer.

The experimental setup is shown in Fig. 4.2. A superconducting transmon qubit is inserted into a high Q, high-purity aluminum superconducting cavity. The cavity is tuned into resonance with the transducer’s microwave circuit with a sapphire rod mounted to a piezoelectric actuator. Due to the strong dispersive coupling between the transmon and the cavity, the latter acquires a qubit-state-dependent frequency shift [Clerk et al. 2010]. After a short qubit preparation pulse, a readout pulse (green) at the cavity frequency ω_c injected via the weakly coupled port acquires a qubit-state-dependent phase shift, and comes out of the strongly coupled port towards the transducer. Directional isolation between the qubit and transducer is provided by circulators and an interferometric cancellation tone to prevent reflected microwave-pump power or optically-generated quasiparticles from dephasing the qubit.

The signal that enters the transducer is upconverted and detected using optical heterodyne. The phase acquired by the pulse during readout will result in a different detected voltage. In Fig. 4.2e, we plot a histogram of the detected voltage when the qubit was prepared in either the ground (teal) or excited (purple) state. We find two bi-modal Gaussian distributions with some overlap, a result of limited readout efficiency and finite state preparation error. After defining a voltage

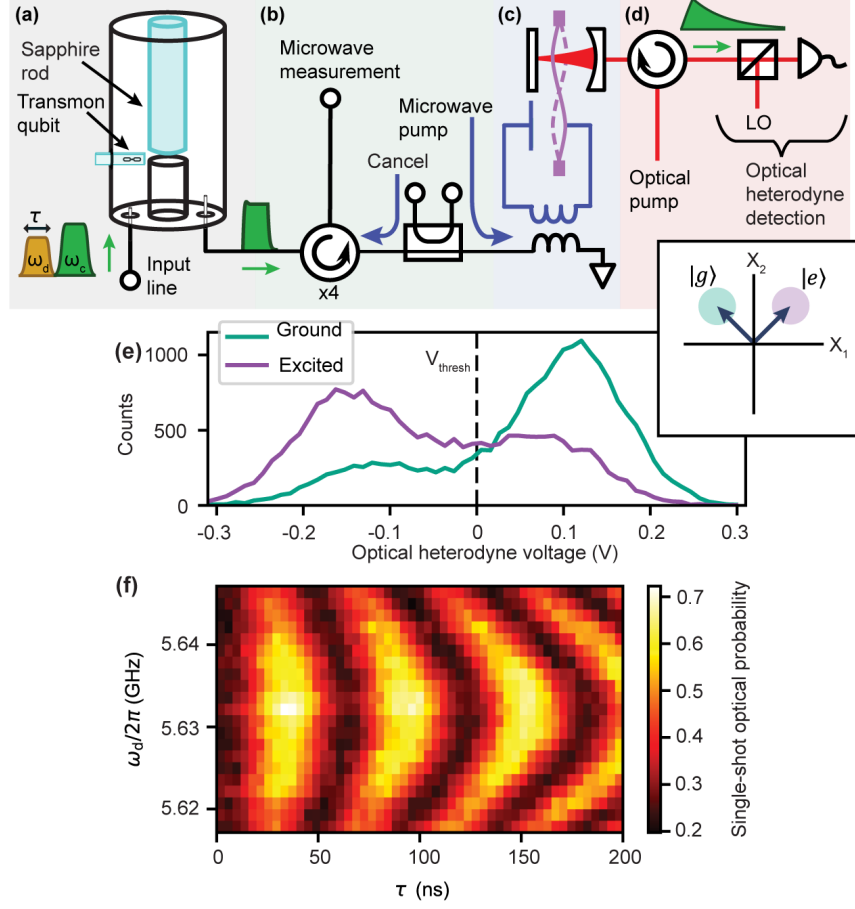


Figure 4.2: **Single-shot optically-detected qubit readout.** (a) The cavity quantum electrodynamic (cQED) system is composed of a transmon qubit coupled to a high Q superconducting aluminum cavity. The cavity is tuned into resonance with the transducer via a moveable sapphire rod. (b) Isolation between the cQED system is provided by 4 circulators and a tone to interferometrically cancel the microwave pump. The readout pulse (green) leaves the cQED, passes through the isolation, then upconverted by the transducer (c) and finally optically detected (d). The dispersive readout pulse has a phase which encodes the qubit state (see inset). (e) Histogram of the optical heterodyne voltage, depending on whether the qubit is prepared in the ground (teal curve) or excited (purple curve) state. V_{thresh} is the voltage threshold for single-shot readout. (f) Rabi oscillations measured optically as a function of the duration of the state preparation pulse (gold). This data was taken while the transducer was operated continuously with damping rates $(\Gamma_e, \Gamma_o)/2\pi = (0.5, 2.4)$ kHz.

threshold V_{thresh} to maximize the readout fidelity, we are able to optically detect Rabi oscillations of the superconducting qubit in Fig. 4.2f by varying the qubit drive frequency ω_d and duration τ (gold).

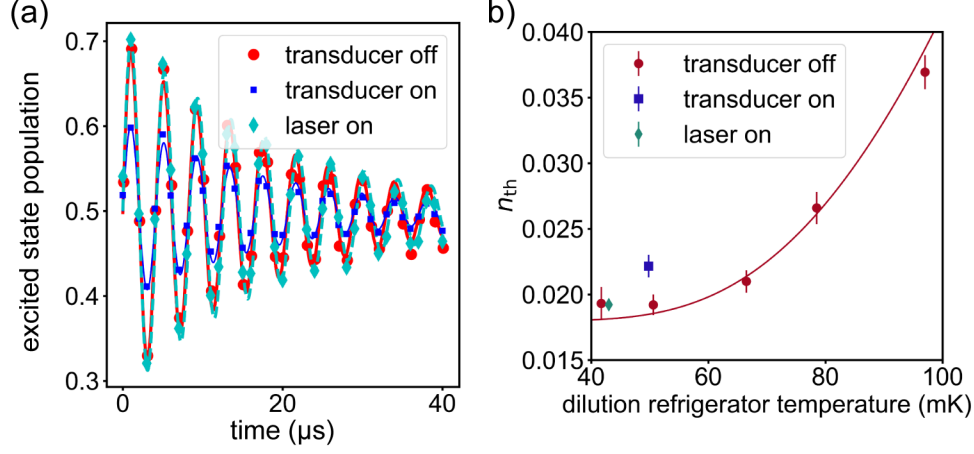


Figure 4.3: **Transducer backaction on qubit** (a) Ramsey fringes with both pumps off ($(\Gamma_e, \Gamma_o)/2\pi = (0, 0)$ kHz; red circles), only the laser on ($(\Gamma_e, \Gamma_o)/2\pi = (0, 5.0)$ kHz; cyan diamonds), and both the microwave and optical pumps on ($(\Gamma_e, \Gamma_o)/2\pi = (1.1, 5.0)$ kHz; blue squares). There is no change in T_2 when turning on the laser, and slight reduction when turning the microwave pump on. There is also a reduction in the contrast of the oscillations due to pump-power dependent loss of the microwave circuit. (b) Thermal population of the qubit extracted from the dephasing rate as a function of fridge temperature. Turning on the microwave pump appears to make the qubit hotter than the base plate.

To quantify the backaction on the qubit, we look at the dephasing time T_2 of the cQED system under different transducer parameters by performing a Ramsey measurement. This measurement is done in the microwave domain by detecting the fraction of the readout pulse which reflects from the transducer and is directed to the microwave readout chain. We fit the decay rate of the oscillations in Fig. 4.3, extracting $T_2 = 20.4 \pm 0.2 \mu\text{s}$ while the laser is on, compared to $T_2 = 20.7 \pm 0.2 \mu\text{s}$ when the laser is off. If there were any optically-generated quasiparticles in the LC circuit which found their way through the isolation to the qubit, we would expect to see a reduction in dephasing time. Remarkably, we cannot resolve any such degradation.

We do observe a small reduction when also applying the electromechanical pump. T_2 is reduced to $19 \pm 1 \mu\text{s}$, which corresponds to an excess backaction of only $\Delta n = (3 \pm 1) \times 10^{-3}$ photons

[Hatridge et al. 2013]. We believe that this excess backaction comes from local heating of microwave components by the large microwave pump powers required to achieve appreciable damping rates. Improving g_e would reduce the required pump powers, likely reducing the transducer’s backaction on the qubit. In addition to reduced T_2 , we also observe a reduction in the contrast of the Ramsey fringes due to microwave-pump power dependent loss of the LC circuit.

Lastly, we can use the qubit as a non-Gaussian resource by which to quantify the quantum efficiency of the whole optically-detected qubit readout chain $\eta_q = \eta_{\text{chain}}/N_{\text{det}}$. η_{chain} captures the various sources of loss between the qubit and optical detection, while N_{det} represents the noise measured in addition to the signal. The latter is dominated by noise emitted by the transducer’s LC circuit, while also having contributions from vacuum noise and added noise of ideal heterodyne detection. The maximum efficiency we measure is $\eta_q = 8 \times 10^{-4}$, in which losses account for $\eta_{\text{chain}} = 1.9 \times 10^{-3}$ (using independent measurement of N_{det}). A more complete breakdown of the inefficiencies can be found in Table 4.2.

Several of the sources of inefficiency are directly or indirectly related to the pump-power dependent microwave loss and noise in the manner described in Fig. 2.7 and Fig. 2.10. The electromechanical damping must be kept somewhat small to prevent overwhelming noise, meaning the total transducer bandwidth cannot match that of the qubit pulse and the transducer is operated far from its maximum efficiency. Loss of the LC circuit induced by the microwave pump and added noise generated in the circuit also contribute inefficiency to the readout. Thus, microwave-pump-power-dependent effects limited the quantum efficiency of the readout with the laser contributing negligible backaction, in contrast to other platforms where optical power is the dominant source of heating of the transducer and dephasing of the qubit.

4.2 Optomechanical ground-state cooling

While this transducer did not achieve quantum-enabled performance, we did reduce the added noise by an order of magnitude from [Higginbotham et al. 2018]. This improvement was enabled by phononic shielding of the membrane mechanical mode (Sec. 3.1), robust cavity construction

description of inefficiency	symbol/expression for efficiency	efficiency value
bandwidth mismatch	η_{bw}	0.15
microwave transmission loss	η_{MW}	0.34
cQED system cavity loss	η_{cav}	0.96
optical detection efficiency	ξ_{o}	0.28
transducer LC circuit loss	$\kappa_{\text{ext,e}}/\kappa_{\text{e}}$	0.53
optical cavity loss	$\kappa_{\text{ext,o}}/\kappa_{\text{o}}$	0.79
optical modematching	ϵ_{CL}	0.80
transducer damping rate mismatch	$\frac{4\Gamma_{\text{e}}\Gamma_{\text{o}}}{\Gamma_{\text{T}}^2}$	0.59
transducer gain	$\eta_{\text{G}} = \mathcal{A}_{\text{e}}\mathcal{A}_{\text{o}}$	1.5
transducer added noise	$1/N_{\text{det}}$	0.42

Table 4.2: **Contributions to quantum efficiency of qubit readout.** Microwave transmission loss accounts for insertion losses of the circulators and other components between the qubit and the transducer. Pump power-dependent loss of the microwave circuit contributes to the LC circuit loss, but also limits the transducer bandwidth and necessitates a large damping rate mismatch. The main contribution to the transducer’s added noise is pump-power dependent noise generated in the LC circuit.

via chemical bonding (Sec. 3.2 and [Urmeý 2024]), as well as changing the superconductor from Nb to NbTiN ([Burns 2019]). This version of the transducer needed less optical power to cool the membrane mode and was also better able to handle the deleterious effects of the laser, permitting ground-state cooling of the mechanical mode with the optical pump. Recalling Sec. 2.4, this is a necessary and sufficient condition for quantum downconversion in the absence of technical noise. We will highlight in this section that there is little technical noise induced by the optical pump, but quantum-enabled conversion is prevented by technical noise and circuit loss induced by the microwave pump.

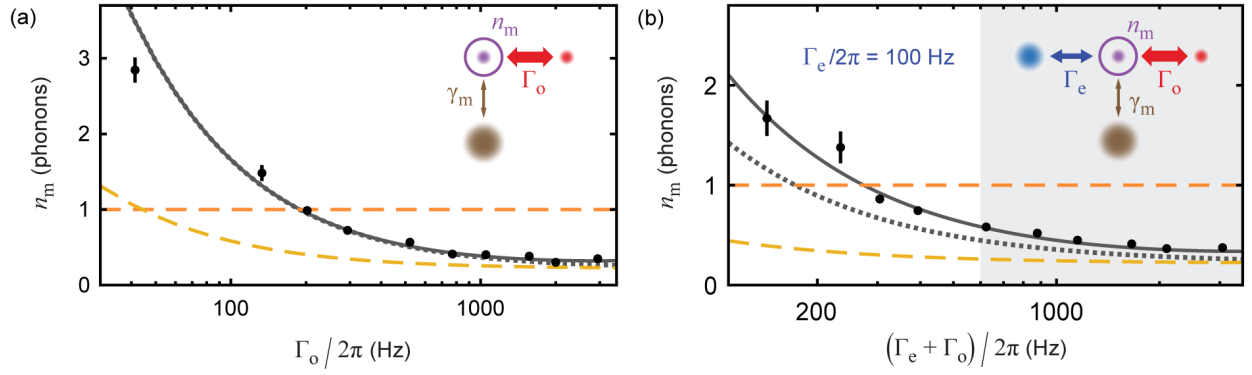


Figure 4.4: **Optomechanical ground state cooling.** (a) Pure optomechanical cooling of the membrane mode. The dots are data, with the solid gray line a fit, the dotted line a prediction in the absence of technical noise, and the yellow dashed line the backaction limit. There is little technical noise induced by the optical pump, and the mechanical mode is cooled to below 1 phonon of occupancy (orange dashed line). Inset illustrates the relative size of the occupations (blurred discs) and relative size of the coupling rates (arrows). (b) Optomechanical cooling with a fixed electromechanical damping rate. While the mechanical mode can still be cooled to the ground state, there is significant technical noise added by the microwave pump. Inset illustrates relative occupancies and coupling rates in the limit that $\Gamma_e \ll \Gamma_o$ (gray shaded region).

The occupancy of the mechanical mode in the presence of optomechanical and electromechanical damping is given by Eqn. 2.23

$$n_m = \frac{\gamma_m n_{th} + \Gamma_e (n_{min,e} + n_{eff,e}) + \Gamma_o (n_{min,o} + n_{eff,o})}{\Gamma_T} \quad (2.23)$$

where we have written the occupations of the microwave and optical cavities as the sum of the

backaction limit and pump-power dependent technical noise. In Fig. 4.4a, we see that the membrane mode can be purely optomechanically cooled to the ground state with negligible contribution from technical noise from the laser. With large Γ_o , the occupation approaches the backaction limit with a minimum mode occupancy of $n_m = 0.32$ phonons. Note, the backaction limit here (dashed yellow line) has contributions from imperfect sideband resolution and from the auxiliary lock beam.

To activate transduction, we turn on the microwave pump to a fixed damping rate of $\Gamma_e/2\pi = 100$ Hz in Fig. 4.4b. While the mode is still cooled to the ground state, there is a significant difference between the measured occupation and that predicted in the absence of technical noise. From the fit to this data, we extract a significant occupancy of the microwave circuit $n_{\text{eff},e} = 0.8 \pm 0.1$ photons which is independent of laser power. Thus, the microwave pump must induce this excess noise. This is shown schematically by the inset, where the occupation of the microwave bath is larger than that of the optical bath. Nevertheless, the mechanical mode is cooled to a final occupation of $n_m = 0.34$ phonons with sufficient Γ_o , again approaching the backaction limit.

The transducer efficiency and added noise in upconversion are given by Eqn. 2.18 and Eqn. 2.35:

$$\eta = \epsilon \frac{\kappa_{\text{ext},e}}{\kappa_e} \frac{\kappa_{\text{ext},o}}{\kappa_o} \frac{4\Gamma_e\Gamma_o}{\Gamma_T^2} \quad (2.18)$$

$$N_{\text{add,up}} = \frac{N_{\text{out},o,+} + \tilde{N}_{\text{out},o}}{\mathcal{A}_e\mathcal{A}_o\eta} = \frac{n_m}{\mathcal{A}_e \frac{\kappa_{\text{ext},e}}{\kappa_e}} \frac{\Gamma_T}{\Gamma_e} + \frac{\tilde{N}_{\text{out},o}}{\mathcal{A}_e\mathcal{A}_o\eta} \quad (2.35)$$

where we replaced the η_{up} with η for a more conservative estimate of the added noise. In order to minimize the added noise, we work with $\Gamma_o/2\pi = 85$ Hz such that less microwave power is required to achieve approximately matched damping rates. This reduces the excess loss and noise induced by the microwave pump. In Fig. 4.5, we observe a maximum efficiency $\eta_{\text{max}} = 47 \pm 1\%$ at $\Gamma_e/2\pi = 75$ Hz and a minimum added noise of $N_{\text{add,up}} = 3.2 \pm 0.1$ photons/[s · Hz] at $\Gamma_e/2\pi = 135$ Hz. This is an order of magnitude improvement in added noise compared to the 34 photons/[s · Hz] measured in [Higginbotham et al. 2018]. Further improvement to the added noise is limited by the excess noise

in the LC circuit induced by the microwave pump, which is why the added noise begins to increase for larger Γ_e .




Figure 4.5: **Minimum added noise sweep.** With the optomechanical damping rate fixed at $\Gamma_o/2\pi = 85$ Hz, we sweep the electromechanical damping rate and find the (a) transduction efficiency and (b) added noise in upconversion. The yellow shaded region indicates where $\Gamma_e \approx \Gamma_o$, corresponding to the maximum efficiency and minimum added noise. Black dots are data and the solid line is a fit.

A full breakdown of the contributions to the minimum added noise in upconversion is given in Table 4.3 by picking out the different terms in Eqn. 2.35 after substituting in the expression for the mechanical occupancy (Eqn. 2.23). Due to the various noise sources, the mechanical mode is not ground state cooled when the $N_{\text{add,up}}$ is minimized, with an occupancy of $n_m = 1.5$ phonons. The two largest contributors to the mechanical occupation are the technical noise in the microwave circuit and the low-frequency thermal bath.

To highlight that $n_{\text{eff,e}}$ is induced by the microwave pump and not the optical pump, we measure the properties of the microwave circuit at a range of circulating laser powers inside the optical cavity in Fig. 4.6. There is little difference between the linear power reflection measurement of the microwave circuit with just the lock beam at 0.3 mW and the optical pump at its maximum value of 74 mW. If scattered light was generating quasiparticles in the superconductor, we would expect the circuit loss to increase drastically. Instead, a only small increase in circuit loss is observed,

source of noise	expression	value (photons/[s · Hz])
residual thermal occupation	$\gamma_m n_{\text{th}} / (\Gamma_e \frac{\kappa_{\text{ext},e}}{\kappa_e} \mathcal{A}_e)$	1
microwave cavity backaction	$\Gamma_e n_{\text{min},e} / (\Gamma_e \frac{\kappa_{\text{ext},e}}{\kappa_e} \mathcal{A}_e)$	0.1
microwave technical noise	$\Gamma_e n_{\text{eff},e} / (\Gamma_e \frac{\kappa_{\text{ext},e}}{\kappa_e} \mathcal{A}_e)$	1.4
optical cavity backaction	$\Gamma_o n_{\text{min},o} / (\Gamma_e \frac{\kappa_{\text{ext},e}}{\kappa_e} \mathcal{A}_e)$	0.2
optical technical noise	$\Gamma_o n_{\text{eff},o} / (\Gamma_e \frac{\kappa_{\text{ext},e}}{\kappa_e} \mathcal{A}_e)$	$\ll 0.1$
lock backaction	$\gamma_{\text{lock}} n_{\text{lock}} / (\Gamma_e \frac{\kappa_{\text{ext},e}}{\kappa_e} \mathcal{A}_e)$	0.4
reflected output noise	$\tilde{N}_{\text{out},o} / (\mathcal{A}_e \mathcal{A}_o \eta)$	0.1

Table 4.3: **Contributions to minimum added noise in upconversion.** The first six contributions directly impact the membrane occupation n_m , while the last term is the noise on the laser which promptly reflects from the cavity which must be referred to the input.

possibly due to heating of the fridge. No increase in the microwave circuit occupancy is observed at the optical powers required to ground state cool the transducer, and remains small even at the largest optical powers. Combined with the opto/electromechanical cooling measurements, this confirms that the noise induced in the microwave circuit is induced by the microwave pump.

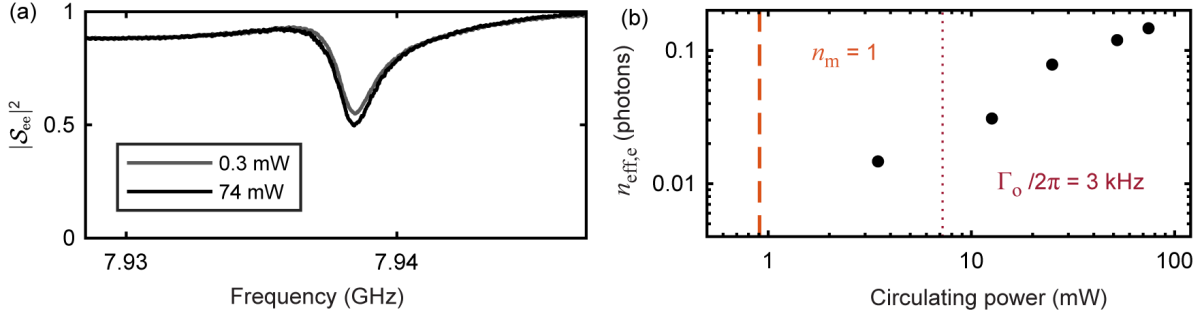


Figure 4.6: **Effects of laser on microwave circuit.** (a) Reflection measurements of microwave circuit with 0.3 mW (gray) and 74 mW (black) of circulating optical power. There is little change to the linear circuit parameters with an optical power 2 orders of magnitude larger than needed for ground-state cooling. (b) Microwave circuit noise induced by the laser. Any additional noise is too small to be resolved for powers needed for ground-state cooling, and remains negligible at large circulator powers.

4.3 Investigating microwave loss and noise

We now turn our attention to understanding the deleterious effects the strong microwave pump tone has on the properties of the LC circuit. The noise and loss in the circuit induced by the microwave pump (shown in Fig. 4.7) are the main limiting factors preventing quantum-enabled transduction. They limit the electromechanical cooperativity, transducer efficiency, minimum added noise, and bandwidth at which the noise is minimized, all key metrics for a mechanically-mediated microwave-to-optical quantum transducer (see Sec. 2.1).

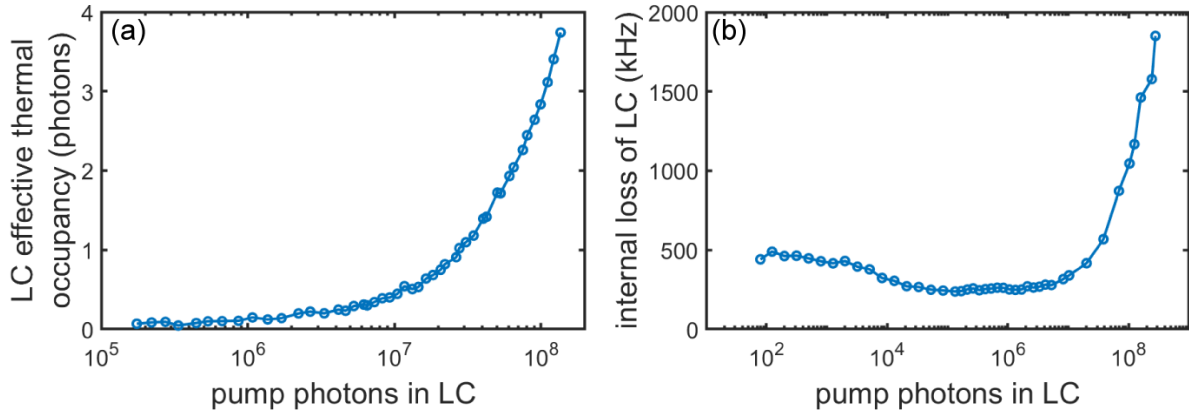


Figure 4.7: **Pump-induced microwave noise and loss.** (a) Occupancy of the LC circuit seen by the mechanics $n_{\text{eff},e}$ as a function of pump power, which has contributions from the phase noise of the generator and noise generated inside of the circuit. (b) Internal loss of the circuit $\kappa_{\text{int},e}/2\pi$ as a function of pump power. There is an initial decrease in loss due to saturation of the resonant TLS-circuit interaction and then an unexpected increase at large powers. $\Gamma_e/2\pi = 135$ Hz required $\bar{n}_e \sim 2 \times 10^7$ photons. Solid lines are a guide to the eye, and not a fit.

To quantify the noise in the LC circuit, we measure the microwave noise spectral density coming out of the LC circuit on resonance. This is referred to the output of the transducer by dividing by the microwave measurement chain efficiency, and then translated into an mode occupancy by dividing by $4\kappa_{\text{ext},e}/\kappa$. We parameterize the noise as linear with the number of pump photons in the circuit \bar{n}_e with a small zero-power offset:

$$n_{\text{eff},e} = a_e \bar{n}_e + b_e \quad (4.1)$$

with $a_e = 2.93 \times 10^{-8}$ photons/photon and $b_e = 0.09$ photons. We can also extract the microwave mode occupancy by fitting the converter added noise as a function of Γ_e , from which we extract parameters that agree with this direct measurement. A portion of this occupancy will be due to the phase noise of the generator (see [Sec. 2.4.1](#)). Our Rohde & Schwartz SMA100B generator provides the microwave pump (with ultralow phase noise afforded by option SMAB-B711) and is measured *in situ* to have a double-sideband phase noise spectral density of $10 \log_{10}[S_{\delta\phi,\delta\phi,e}(\omega_m)] = -149$ dBc/Hz when the carrier is at $\omega_e - \omega_m$, which accounts for about a quarter of the total measured noise.

The other, possibly related, effect of the microwave pump is excess circuit loss, especially at the largest microwave powers used in the last two sections. The behavior of the loss below $\bar{n}_e = 10^6$ is consistent with the power-dependent loss measured in other superconducting circuits, modeled by the resonant-type interaction with TLS defects [[Gao 2008](#); [Martinis et al. 2005](#); [Müller et al. 2019](#)]. The TLS model itself does not predict the increasing loss at $\bar{n}_e > 10^6$, and as a result, we spent a significant amount of time investigating many other potential sources of loss [[McRae et al. 2020](#)]. In particular, we were wary of the NbTiN, as this was a relatively new material in our fabrication process and in the general superconducting circuit community (when compared to Nb and Al). In the remainder of this section, we will review the results of these experiments. These included tests with trapped flux and critical temperature measurements of the superconducting films. We also performed surface and bulk characterization of the NbTiN. Ultimately, we found that Si_3N_4 , not the superconductor, was the source of this excess power-dependent loss and in the next [Ch. 5](#), we will see that this effect can be explained by the TLS model with a simple self-heating ansatz.

4.3.1 Trapped flux

When a thin-film superconductor is cooled through its superconducting transition temperature T_c in the presence of a magnetic field B_{cool} , there can be an incomplete expulsion of the magnetic field. This is especially true in type-II superconductors like NbTiN. The quantized flux trapped in the superconductor, known as vortices, can then behave as resistive charge carriers if allowed to propagate around the circuit [[Stan et al. 2004](#); [Song, Heitmann, et al. 2009](#)]. While the vortices can

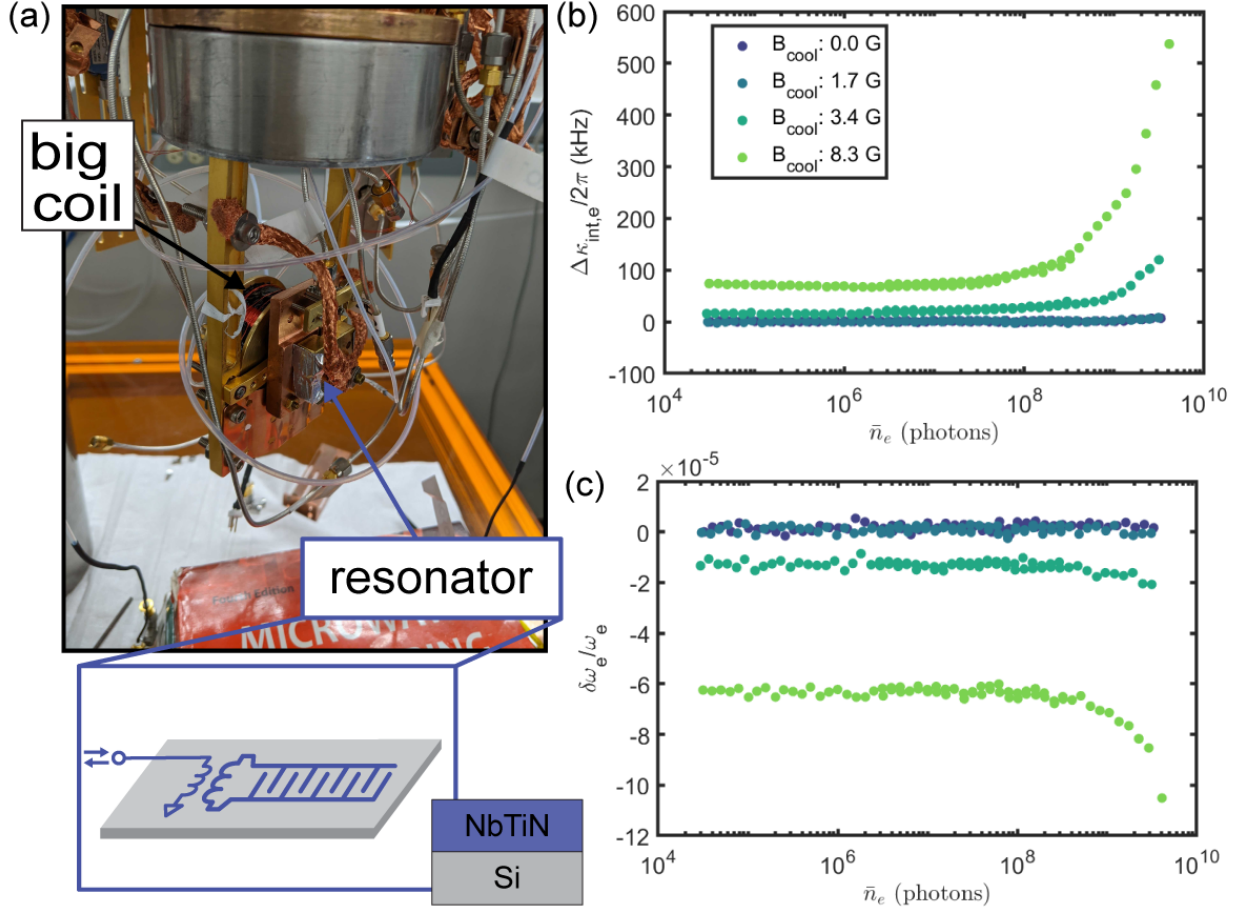


Figure 4.8: **Trapped flux in LC circuit.** (a) Experimental setup. The magnet coil is positioned approximately 1.5 cm away from the resonator such that the magnetic field is out-of-plane with the superconducting film. The coil is made from low-resistance copper wire to avoid heating the fridge excessively while applying tens of mA of current. The resonator is a planar lumped element LC circuit with a 5 μm trace width made with 200 nm NbTiN on Si (inset). To change the amount of trapped flux, we warm up the fridge to 20 K, and then cool back down while applying the magnetic field. (b-c) Internal loss and fractional frequency shift as a function of pump power and magnetic field relative to the value at low power $\bar{n}_e = 3 \times 10^4$ and $B_{\text{cool}} = 0$ G.

become pinned at defects in the superconducting trace, they become mobile and contribute loss in the presence of thermal energy or strong electromagnetic fields in the circuit. The magnitude of this additional loss will depend on the number of vortices and characteristic strength of the pinning potential.

Magnetic shielding of the fridge at room-temperature and of the device at the base plate are common techniques to mitigate this loss source [Kreikebaum et al. 2016; Nsanzineza 2016]. We can do neither effectively due to the need for free-space optical access to the membrane. Moreover, the mount for the transducer is made with Invar to prevent misalignment of the optical cavity during cooldown. Invar’s extraordinarily low coefficient of thermal expansion is intimately related to its magnetic properties [Lohaus et al. 2023]. Thus, there is possibly a sufficiently large out-of-plane magnetic field at the superconducting film due to the Invar, other magnetic components inside the fridge (ie circulators, stainless steel screws, etc.), and Earth’s magnetic field.

The width of the inductor in the LC circuit is designed with this constraint in mind. There is complete expulsion of the magnetic field below a critical field $B_m \sim \Phi_0/w^2$, where $\Phi_0 = h/2e$ is the superconducting flux quantum and w is the width of the trace. For the $5\,\mu\text{m}$ trace width in the transducer, the critical field is 0.65 - 3.3 G depending on whether the flux is trapped at a metastable or absolutely stable equilibrium [Stan et al. 2004]. We measured the magnetic field from the invar mounts to be 0.1-0.3 G, which in combination with stray magnetic fields could surpass the critical field.

To investigate this experimentally, we measure the power-dependent loss and frequency shift of a resonator which is cooled through T_c at a range of magnetic fields which are provided by a solenoid mounted near the resonator. The resonator is a planar, lumped-element LC circuit whose inductor is the same design as the transducer. The capacitor is an interdigitated finger capacitor, and the circuit is made from NbTiN on Si with a $5\,\mu\text{m}$ trace width. To reset the trapped flux, we warm up the fridge to 20 K and then apply a magnetic field while cooling back down to mK temperatures. The magnetic field is then turned off while characterizing the resonator. The magnet and resonator are housed in a magnetic shield at the base plate to shield ambient fields.

The experimental setup and measurements are detailed in [Fig. 4.8](#). At $\bar{n}_e = 3 \times 10^4$ and $B_{\text{cool}} = 0$ G, we find $\kappa_{\text{int,e}}/2\pi = 100.0$ kHz and $\omega_e/2\pi = 5.553$ GHz. We plot the loss and fractional frequency shift relative to this value to highlight the power- and magnetic-field dependence. The reported magnetic field is an estimate based on room-temperature measurements with a gaussmeter placed approximately at the location of the resonator. A more conservative measurement with the gaussmeter slightly farther away reduced the measured field by a factor of 2. We see that an increasing magnetic field contributes excess loss and shifts the frequency lower for $B_{\text{cool}} > 1.7$ G. The additional loss should scale linearly with B_{cool} above the critical field, which we use to extract a critical field of $B_m = 2.0$ G (1.0 G if we use the more conservative estimate of the field), consistent with that expected from theory.

We do observe a qualitatively similar increase in loss at large drive powers for $B_{\text{cool}} > B_m$ as seen in the transducer and the power-dependence of the loss also increases with increasing B_{cool} . However, the magnetic field required to reproduce the power-dependent loss in the transducer is unreasonably large. For the largest value of B_{cool} , the loss increases by 500 kHz from its low power value at $\bar{n}_e = 4 \times 10^9$ photons, while in the transducer this happens at two-order of magnitude lower power. Thus, we require a field much larger than 8 G (4 G conservatively) to replicate the scale of power-dependent loss in the transducer, which is inconsistent with the ≈ 1 G maximum fields we expect from independent measurements.

As a final note, loss from trapped flux can also be mitigated by the implementation of flux-pinning structures in the ground plane or inductor trace of the circuit [[Song, DeFeo, et al. 2009](#); [Bothner et al. 2011](#); [Chiaro et al. 2016](#)]. We implemented flux pinning holes in the capacitor pad of some transducer devices, as they are much wider than the inductor traces and thus the critical field for flux trapping is much smaller. No difference in the power-dependent loss was observed in those transducer devices, likely because the current density is negligible in the capacitor pads. There are some reasons to add strong magnets to the base plate of the fridge, such as to create eddy current dampers for mechanical vibrations [[Wijts 2013](#); [C. Lee et al. 2018](#)], in which case flux pinning geometries could be beneficial.

4.3.2 Quasiparticle loss

Quasiparticles, or unpaired charge carriers in a superconductor, can be generated from a number of sources. Generally, some external energy is added to the system which is sufficient to break Cooper pairs, which then have a finite recombination time. This energy can come from stray infrared photons, higher-order interactions from microwave photons, or thermal phonons if the device temperature is close enough to T_c . We have already seen that our circuit is robust to infrared light in the optical cavity (Fig. 4.6), and loss from this source is not expected to scale with microwave power. Multiphoton absorption of microwave signals can also generate pair-breaking radiation [Goldie and Withington 2012; de Visser et al. 2014], which should scale directly with microwave power. In the previous section, we found that we could put more than 10^9 photons into a planar resonator with the same inductor design as the transducer without inducing any excess loss (Fig. 4.8). As the inductor is the location of largest current density, we expect that loss from microwave pair-breaking radiation is negligible for the smaller powers used in the transducer experiments.

We do observe temperature-dependence to the loss in addition to the power-dependence which could be a result of thermally-generated quasiparticles. While we expect the T_c of our NbTiN film to be 13-17 K, the transition temperature is sensitive to film thicknesses and deposition conditions [Giachero et al. 2023]. If some step in our complicated fabrication process is drastically changing the NbTiN material properties (see Appendix A), we could inherit power-dependent loss from heating of the superconducting film which leads to thermally-activated quasiparticles.

With the help of Jordan Wheeler at NIST Boulder, we measured the T_c of a set of NbTiN films that experienced the same fabrication steps as those in our transducers. The metal on the circuit chips is 200 nm thick and deposited on bare Si, while the metal on the membrane chips is 30 nm thick and deposited on top of a 100 nm layer on Si_3N_4 on a Si substrate. We measure the T_c of four films in Fig. 4.9a, mapping out the two-dimensional parameter space of thick/thin NbTiN on Si/ Si_3N_4 on Si. For NbTiN on Si, we find $T_c = 15.5$ K, regardless of thickness. When the metal is deposited on Si_3N_4 , the T_c is reduced to 14 K for the thicker film and 13 K for the thinner film.

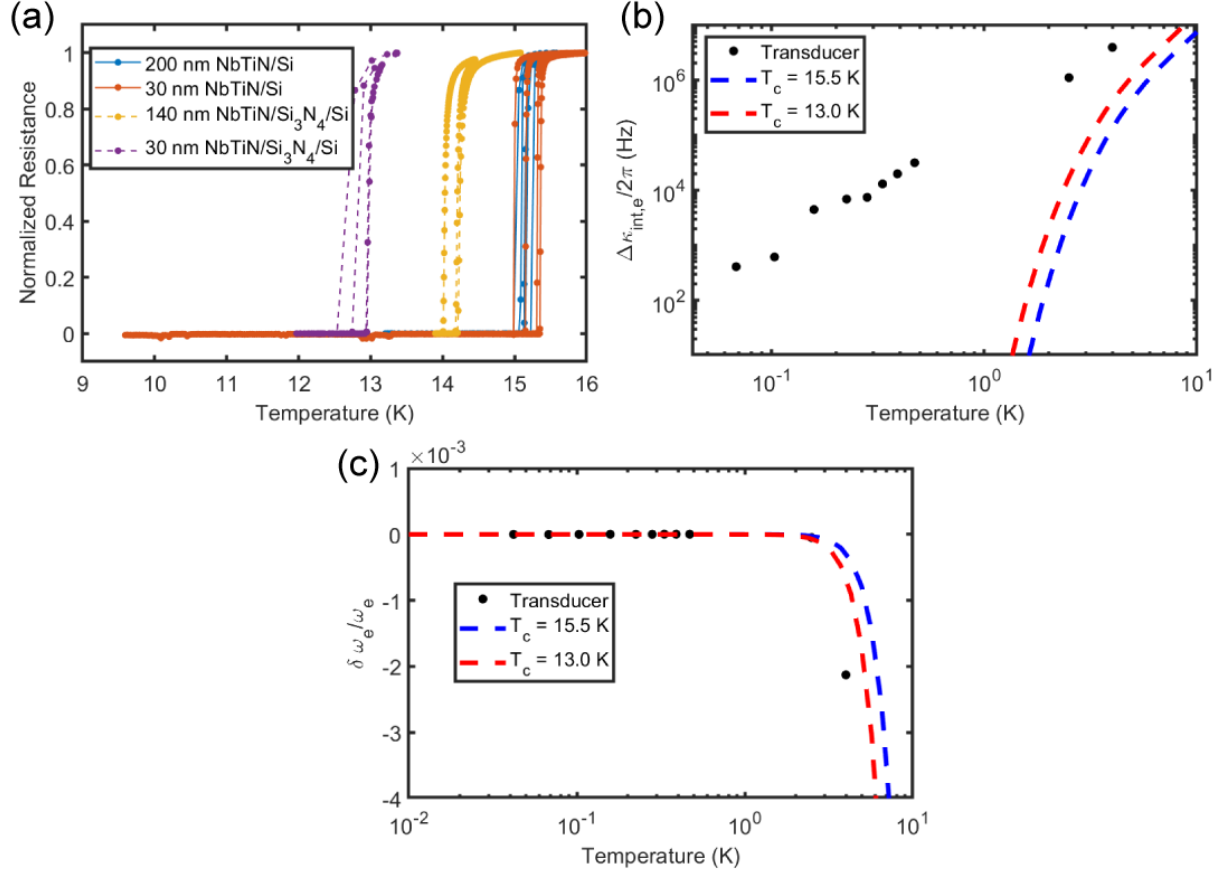


Figure 4.9: **NbTiN T_c and Thermal Quasiparticles.** (a) Temperature-dependent sheet resistance for a set of NbTiN films with different properties normalized to the normal state resistance. The resistance is measured while cooling down and warming up the adiabatic demagnetization refrigerator (ADR). Thermal lag between the thermometer and sample temperature likely explains the spread in T_c for the same film. (b) Circuit loss and (c) fractional frequency shift of the transducer relative to the value at $T = 40$ mK as a function of temperature. Dashed lines are predictions based on Mattis-Bardeen theory of surface impedance for the best (blue) and worst (red) case values of T_c .

This reduced value of T_c , however, predicts a temperature-dependent loss that is much smaller than observed dependence between 0.04 - 4 K. The thermally-excited quasiparticle induced loss and frequency shift of the circuit is given by [McRae et al. 2020]:

$$\frac{\kappa_{\text{int,e,qp}}}{\omega_e} + 2i\frac{\delta\omega_e}{\omega_e} = \frac{\alpha}{\omega_e\mu\lambda} (R_s + i\delta X_s). \quad (4.2)$$

α is the kinetic inductance fraction and λ is the London penetration depth for energies much smaller than the superconducting gap. The real and imaginary parts of the surface impedance R_s and X_s are calculated using the Mattis-Bardeen theory for the anomalous skin effect [Mattis and Bardeen 1958]. The kinetic inductance fraction is calculated to be $\alpha = 0.15$ by comparing the simulated and measured resonance frequency of planar resonators. The penetration depth depends on the normal state resistivity $\rho_n = 168 \times 10^{-8} \Omega$, which is measured by a room-temperature four-point probe, and the superconducting gap, which we approximate using T_c . As a final note, we truncate the energy integral for the complex conductivity in the Mattis-Bardeen equations at 10^5 times the superconducting gap to help convergence without loss of accuracy.

In Fig. 4.9b-c, we plot the measured temperature-dependent loss and fractional frequency shift of the transducer circuit relative to the values at base temperature, along with predictions given the best- (blue) and worst-case (red) values of T_c . While there is rough qualitative agreement with the fractional frequency shift, both values of T_c predict a temperature-dependent loss orders of magnitude smaller than observed. At the highest temperature of 4 K, thermally-excited quasiparticles can only account for 3-8% of the measured excess loss. Additionally, we would require a strong self-heating non-linearity to model the power-dependent loss. The superconducting film would need to be at 4-6 K to explain all of the power-dependent loss measured at $\bar{n}_e = 10^8$ photons, which we find unlikely given that the metal should thermalize well to the baseplate via the Si and Al wirebonds.

4.3.3 NbTiN surface and bulk analysis

The new superconductor in this transducer was introduced to reduce the power-dependent noise of the microwave circuit, which we often referred to as "parameter" noise. Measurements of the fractional frequency noise in resonators made with NbTiN and Nb found a factor of four reduction in the devices with NbTiN, which agrees with observations in the literature [Burns 2019; Barends et al. 2010]. This improvement could be attributed to the higher T_c and faster quasiparticle relaxation time of NbTiN, or the different oxide formation at the metal-air interface [Zhang et al. 2018; Nazib et al. 2023].

The NbTiN is sputtered using a process developed at NIST Boulder by Mike Vissers for Si substrates. Given our relative unfamiliarity with this material and the additional complication of sputtering onto the Si_3N_4 , we explored several material characterization techniques to probe the surface and bulk properties of the superconductor. While we found some differences between NbTiN on Si and NbTiN on Si_3N_4 on Si, finding an obvious connection to the power-dependent loss and noise proved difficult.

To determine the bulk material composition of the NbTiN, we employed secondary ion mass spectrometry with the help of QSpec Technologies [Lockyer et al. 2024; QSpec Technology, Inc. 2024]. In this technique, a sample under vacuum is bombarded with primary ions (in this case oxygen ions) which removes material at a previously calibrated rate. The secondary ions (in this case Nb and Ti ions) are ejected from the sample and pass into a mass spectrometer, generating an intensity vs time plot which can be converted into an elemental concentration vs depth profile.

For this analysis, we measured two films, one NbTiN on Si and the other NbTiN on 100 nm of Si_3N_4 on Si, both of which had experienced the usual fabrication steps. In Fig. 4.10, we have plotted the relative concentration of Nb x such that the chemical formula for the superconductor is more accurately given by $\text{Nb}_x\text{Ti}_{1-x}\text{N}$. For both films, Nb is dominant over Ti at the surface, and the bulk contains about two Nb atoms for every Ti atom, consistent with similar films deposited in comparable conditions [Burdastyh et al. 2020]. Curiously, the NbTiN on Si has composition

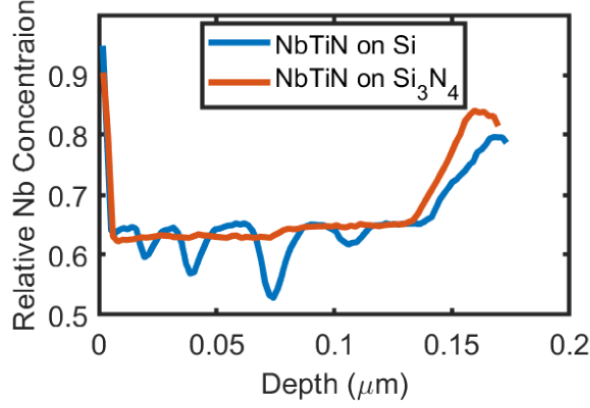


Figure 4.10: **Depth profile of $\text{Nb}_x\text{Ti}_{1-x}\text{N}$ composition.** The surface is predominantly Nb, while the bulk is about two-thirds Nb. The film on bare Si has some fluctuations in the composition, which could be due to fluctuations of the relative Nb and Ti beam target powers during sputtering. The y-axis here corresponds to x in $\text{Nb}_x\text{Ti}_{1-x}\text{N}$.

fluctuations on the order of ten percent as a function of thickness, which could arise from fluctuations in the Nb and Ti target powers during sputtering. Resonators made with this same film did not demonstrate any power-dependent loss. We conclude then that these fluctuations, while surprising, are unrelated to the power-dependent loss measured in the transducer. It also unknown if these fluctuations occur in every sputtering process.

Next, we turn to analyzing the surfaces, particularly that of the metal. The surface is a likely location of residues from the fabrication process as well as lossy dielectrics in the form of oxides. Surfaces are particularly important for the electromechanical capacitor, as the electric field must pass through the metal-air interface. We employ two techniques: X-ray photoelectron spectroscopy (XPS) and atomic force microscopy (AFM). The former measures the energy of electrons emitted from the photo-electric effect upon X-ray illumination and reveals the chemical composition and chemical bonding states of the surface. The latter performs a high-resolution two-dimensional scan of the surface in order to determine the surface roughness.

The XPS measurement and analysis was performed by Intertek [*X-Ray Photoelectron Spectroscopy (Allentown Lab) 2024*] over a $200\text{ }\mu\text{m}$ region with up to 0.1 eV resolution on two samples similar to those used for the SIMS analysis. The relative atomic percentages are tabulated in [Table](#)

Sample	C	F	N	Na	Nb	O	Si	Ti
NbTiN on Si	24.6	ND	4.5	ND	18.9	48.8	1.7	1.1
NbTiN on Si ₃ N ₄	41.1	1.7	1.2	0.5	11.9	41.0	0.7	2.1

Table 4.4: **Surface composition from XPS.** Relative atomic percentage after normalization and omission of hydrogen, which is undetectable by XPS. TiO₂, NbN, and Nb₂O₅ peaks were observed for both films. The most significant difference is the large atomic percentage of Carbon on the second sample, likely from organic residues. ND is short for not-detected.

4.4. For both samples, the Nb/Ti ratio at the surface is ≈ 0.9 , similar to that measured via SIMS. Analysis of the various peaks observed in the spectrum identified TiO₂, NbN, and Nb₂O₅ at the surface of both films. The main difference between the two samples is the larger concentration of carbon bonds on the second film, which from the spectrum are consistent with C-O bonds in organic residues. These residues should be removable by strong acids such as nanostrip. We observed little change in the power-dependent loss of planar lumped-element devices as a function of surface cleaning steps such as nanostrip, hydrofluoric acid, and oxygen ashing, leading to the conclusion that surface contaminants on the NbTiN are not the cause of the excess loss and noise.

The final materials characterization was AFM, performed with the help of John Biesecker at NIST Boulder. In principle, the various processing steps in the fabrication flow could negatively impact the surface roughness, which could lead to additional lossy oxides [Earnest et al. 2018; Lei et al. 2023]. In our samples, we do observe a degradation in the surface uniformity with additional processing, but the RMS roughness remains within a factor of 2 of the pristine Si wafer (Table 4.5). While increased surface roughness can lead to additional losses, this effect is significant when the surface roughness approaches tens of nanometers, much larger than the measured roughness of our films [Wisbey et al. 2010].

4.3.4 Identifying Si₃N₄-related dielectric loss

While the results of the last two section did not find any obvious reason for the NbTiN on Si₃N₄ to be defective, we confirm those conclusions with measurements of planar and non-mechanical

Sample	RMS Roughness (pm)
Bare Si	281
Si ₃ N ₄ on Si	265
30 nm NbTiN on Si	320
30 nm NbTiN in Si ₃ N ₄	530

Table 4.5: **RMS Roughness from AFM scans.** The AFM scans are taken over a 10 μm square. The samples seem to generally get rougher with more processing or material depositions, but exhibit a RMS roughness much smaller than the film thicknesses.

flip chip devices with Si₃N₄ incorporated into the material stack. By varying the geometry of these resonators, we are able to tease apart the impact of the various materials and their interfaces. In particular, interfaces can be host to two-dimensional electron gases [Koley and Spencer 2005; Park et al. 2010; Bristowe et al. 2014] and to piezoelectricity [H. Zhou et al. 2024], both of which can contribute excess dissipation. The metal-air interface (the surface of the metal) can also be host to oxides or organic residues which contribute dielectric loss. Ultimately, we identified the bulk Si₃N₄ as the source of the power- and temperature-dependent loss, not the superconductor.¹ In next chapter, we will quantitatively model this loss with the standard tunneling model of two-level systems. The remainder of this section will review the measurements which lead to the qualitative conclusion that the power- and temperature-dependent loss is related to the Si₃N₄ in the capacitive part of the circuit.

To understand the effect of Si₃N₄ in the microwave resonators, we design a set of planar lumped-element resonators with the Si₃N₄ in either the inductive or capacitive sections of the circuit (Fig. 4.11). The Si₃N₄ is only underneath the superconductor, as clarified by the inset. The inductive Si₃N₄ design will determine if the NbTiN on Si₃N₄ has reduced conductance or critical current, while the capacitive Si₃N₄ designs will probe the dielectric loss of the bulk Si₃N₄ as well as the

¹In previous work in this experiment, similar power-dependent loss was attributed to the superconductor and apparently cured by switching from Nb or faulty NbTiN to properly deposited NbTiN [Burns 2019]. It is possible that conclusion incorrectly identified the source of that loss, given the findings in this thesis. It also unclear what fraction of the power-dependent noise improvement with NbTiN is actually due to the changed superconductor, as opposed to reduced participation in the Si₃N₄

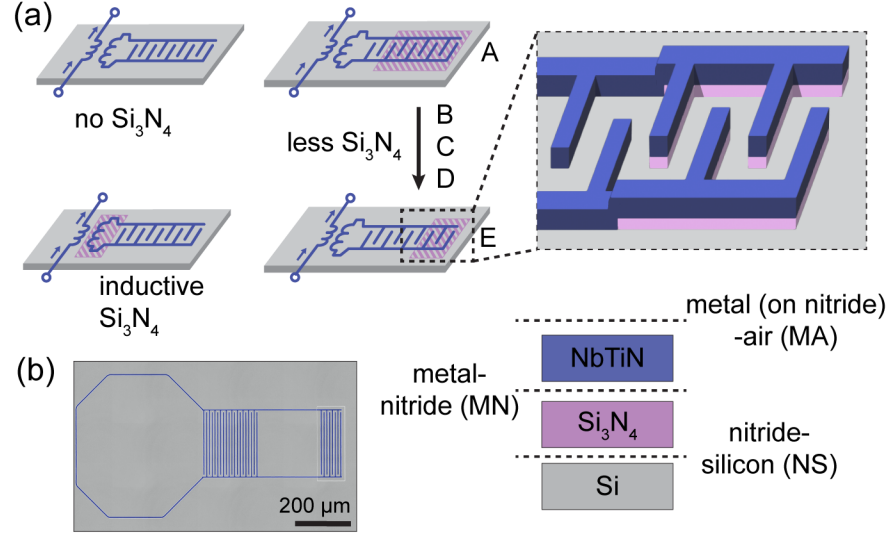


Figure 4.11: **Inductive and Capacitive Si_3N_4 .** (a) Schematics of planar, lumped element resonators with Si_3N_4 incorporated into the inductive or capacitive parts of the circuit. The cross-hatched pink regions represent that the Si_3N_4 is only directly underneath the NbTiN (blue). It is shown as a rectangle for clarity in the schematics, but the actual geometry is detailed in the inset. (b) False color optical micrograph of design D.

interfaces. The capacitor is designed as an interdigitated finger capacitor, and design A has the largest proportion of fingers on Si_3N_4 , while design E has the least. The resonators are fabricated on the same chip and coupled inductively to a common transmission line such that they can be measured simultaneously. Each resonator is designed at a slightly different frequency (all centered around 6 GHz), permitting selective addressing.

Fig. 4.12 presents the results of power- and temperature-dependent characterization of the microwave circuit parameters for these designs. Immediately, we identify that the inductive Si_3N_4 (brown) and no Si_3N_4 (light blue) designs both experience little change in the resonator parameters as a function of power and temperature, while all the capacitive designs do exhibit appreciable loss and frequency shift. The scale of these parameter changes also appears to correlate with the amount of Si_3N_4 that is incorporated into the interdigitated capacitors. These observations suggest that the power- and temperature-dependent effects arise from electric fields in some region that scales with the amount of Si_3N_4 in these designs. In addition to the bulk Si_3N_4 , the metal (on

nitride)-air, metal-nitride, and nitride-silicon all scale with the planar extent of the capacitive Si_3N_4 . In the transducer devices, the electric fields of the parallel plate could also extend into these regions. Disentangling the contributions from these various regions requires more quantitative analysis with finite-element simulations, which we detail in the next chapter.

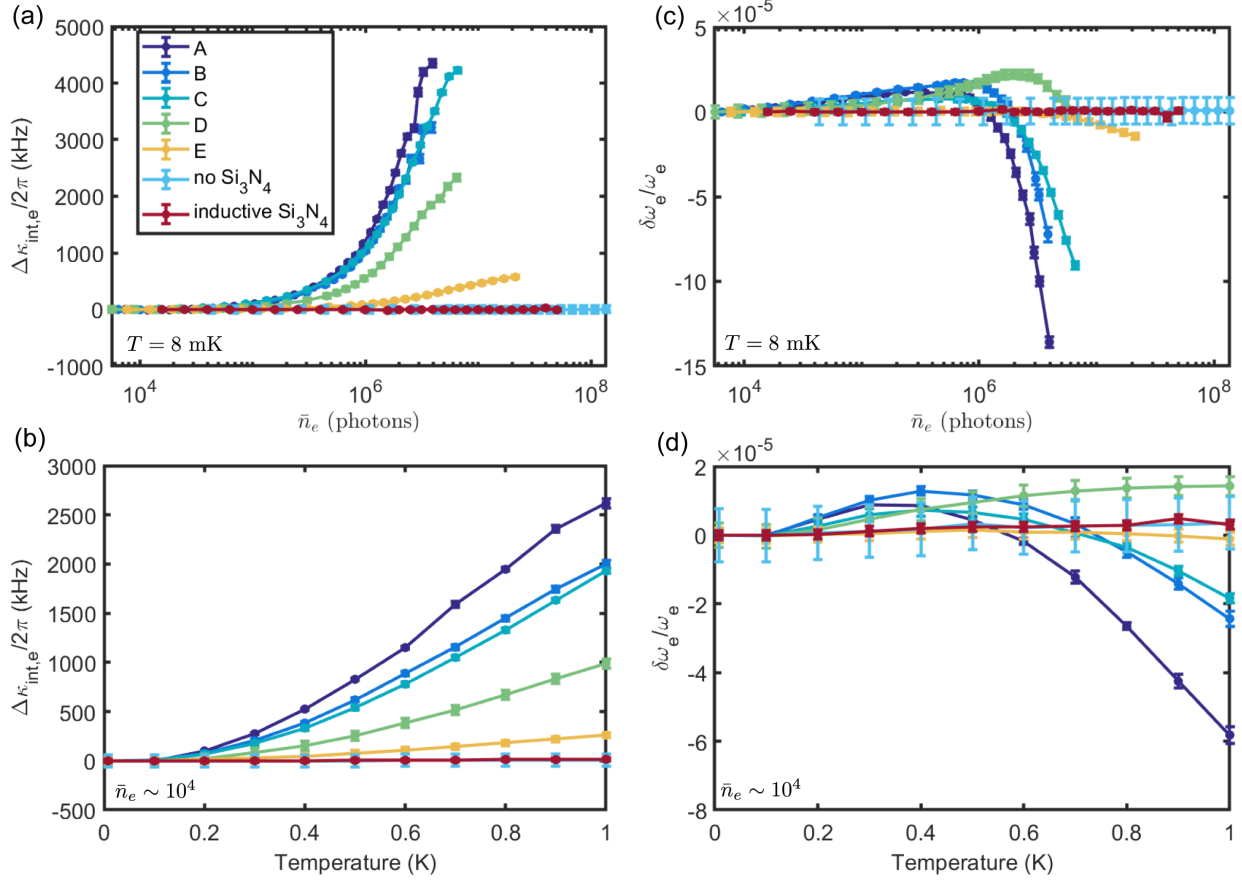


Figure 4.12: **Measurements of inductive and capacitive Si_3N_4** Power- and temperature-dependent (a-b) internal loss and (c-d) fractional frequency shift relative to values at $\bar{n}_e = 10^4$ and $T = 8$ mK. While the inductive Si_3N_4 resonator demonstrates similar behavior to the resonator with no Si_3N_4 , the excess power- and temperature-dependent loss of the capacitive Si_3N_4 resonators scales with increasing Si_3N_4 in the interdigitated capacitor. The frequency shift displays similar behavior, but the correlation is less clear. Solid lines are a guide to the eye, not a fit.

Chapter 5

Microwave properties of Si_3N_4

*You're one hundred percent positive that the ship which is crashed on the bottom
of this ocean is the ship which you said you were one hundred percent positive
could one hundred percent never crash?*

- Douglas Adams

At the end of the last chapter, we arrived at the conclusion that the power- and temperature-dependent microwave circuit loss is due to electric field energy in the bulk or interfacial dielectrics related to Si_3N_4 . In this chapter, we show that this is due to the bulk Si_3N_4 with the help of finite element simulation (Sec. 5.1). The conventional knowledge of dielectric loss in superconducting circuits is captured by the standard tunneling model of two-level systems [Gao 2008; Behunin et al. 2016; McRae et al. 2020; Phillips 1987; Hunklinger and Arnold 1976; Anderson et al. 1972], which in its most common form does not predict the excess loss or frequency shift we observe at high drive powers. Most of the modern research into the dielectric loss at cryogenic temperatures has been focused on the sub 10 mK, single-photon regime which is most relevant for superconducting qubits. This regime is well-described by the so-called "resonant"-type interaction, in which the microwave circuit and TLS directly exchange energy.

With help from Alec Emser and Cyril Metzger, we found the temperature-dependent loss and frequency shift we observe above 500 mK can be modeled by the so-called "relaxation"-type interaction of the TLS model (Sec. 5.2). This term is often ignored in modern studies as it is only relevant when the thermal energy $k_{\text{B}}T$ is similar to or greater than the TLS energy E , which usually occurs at temperatures unsuitable for qubit operation. We then introduce a simple self-heating ansatz

to phenomenologically model the power-dependent loss and frequency shift, achieving reasonable agreement and providing a possible source of the microwave circuit noise. We also explore the implications that the self-heating model has for the thermal conductance in our devices (Sec. 5.2.5).

To improve the dielectric properties of the Si_3N_4 , we took inspiration from techniques used in near-infrared nanophotonics experiments, with guidance from Scott Papp and David Carlson at NIST Boulder. A post-deposition, high-temperature thermal anneal of the Si_3N_4 improved the power- and temperature-dependent loss from the relaxation interaction by more than two-orders of magnitude (Sec. 5.3). Nuances in the modeling and the measured improvements to the TLS loss have implications for the model itself (Sec. 5.5). Finally, we attempt to use the TLS model to understand the temperature-dependent loss of the mechanical membrane mode (Sec. 5.6). Much of the discussion on the microwave properties of Si_3N_4 , the impact of annealing, and the self-heating ansatz was published in [Mittal et al. 2024].

5.1 Identifying bulk Si_3N_4

To disentangle the contributions to the loss rate from the various bulk dielectrics and interfaces, we will use a participation formalism, where the total microwave circuit loss $Q_i^{-1} = \kappa_{\text{int,e}}/\omega_e$ is given by a weighted sum of the material intrinsic loss rates q_j :

$$Q_i^{-1} = \sum_j F_j q_j^{-1} \quad (5.1)$$

where $F_j \equiv U_j/U_{\text{tot}}$ is the fraction of the total electromagnetic energy that is stored in the j -th material [Wenner et al. 2011; Read et al. 2023a]. To calculate this participation or "filling-fraction", we can employ finite-element simulation in Ansys HFSS eigenmode solver [Ansys HFSS / 3D High Frequency Simulation Software 2024]. We integrate the electric field energy stored in the j -th dielectric with volume V_j and dielectric constant ϵ_j and then divide by the total energy of the

geometry:¹

$$F_j = \frac{\int_{V_j} \epsilon_j |E|^2 dV}{\int_V \epsilon |E|^2 dV}. \quad (5.2)$$

For the interfaces, which are only represented as zero thickness sheets, we integrate the field over the surface and then multiply by the interface thickness $t_{\text{int}} = 3 \text{ nm}$. We employ the same thickness for all interfaces for ease of comparison. It is also the simplest assumption without detailed measurement of the interfaces. We also assume that the metal-nitride (MN) and nitride-silicon (NS) interfaces have the same dielectric constant as Si_3N_4 , while the metal (on nitride)-air (MA) interface is $\epsilon = 10$, in agreement with the techniques described in [Read et al. 2023a].

Design	Electric Energy Participation F_j			
	Si_3N_4	metal (on nitride)-air (MA)	metal-nitride (MN)	nitride-silicon (NS)
transducer	$[4.4 \times 10^{-4}, 1.7 \times 10^{-3}]$	0.141	1.2×10^{-5}	1.2×10^{-7}
A	0.0424	1.7×10^{-3}	1.3×10^{-3}	1.4×10^{-3}
B	0.0324	1.3×10^{-3}	1.1×10^{-3}	1.1×10^{-3}
C	0.0276	6×10^{-4}	5×10^{-4}	5×10^{-4}
D	0.0144	4×10^{-4}	4×10^{-4}	4×10^{-4}
E	2.5×10^{-3}	1×10^{-4}	2×10^{-4}	4×10^{-4}
screened Si_3N_4	1×10^{-7}	0.045	-	-

Table 5.1: **Electric field energy participation.** The Si_3N_4 participation of the transducer is provided as a range to account for misalignment (at most $20 \mu\text{m}$) of the capacitor pads. The transducer is simulated with a 200 nm gap spacing, and the screened Si_3N_4 is simulated with a 600 nm gap. The schematics for designs A-E are shown in Fig. 4.11.

From the table, we see that MA interface has the most participation in the transducer as the electric field largely exists between the capacitor pads. The second largest contribution in the table

¹For the magnetic field energy participation, we can integrate the surface current density over the relevant superconducting section and multiply by the film thickness. If the film thickness was larger than the penetration depth $\lambda_{\text{NbTiN}} = 300 \text{ nm}$, we would instead multiply by λ_{NbTiN} . This is then normalized to the total magnetic field energy.

is the bulk Si_3N_4 , as there are fringing fields of the parallel plate capacitor which extend into the membrane.² The planar resonators described in the last chapter (Fig. 4.11) are designed to sweep the Si_3N_4 participation F_{SiN} , and we design a new resonator that isolates the MA participation F_{MA} . This resonator has a flip geometry similar to the transducer, but the NbTiN completely encapsulates the Si_3N_4 which effectively screens the fringing fields. Thus, we will call it the "screened Si_3N_4 " resonator.

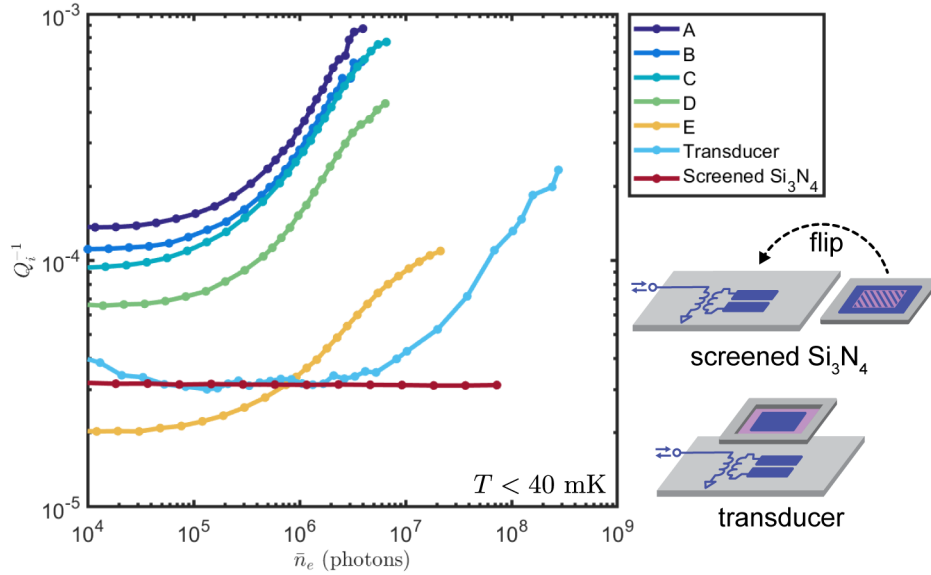


Figure 5.1: **Power-dependent loss comparison.** The internal loss of the circuit Q_i^{-1} as a function of power is plotted for the transducer, planar resonators A-E (Fig. 4.11), and the screened Si_3N_4 design (geometry described on by the schematic). Using the participation values in Table 5.1, the Si_3N_4 can explain the loss values with a intrinsic material loss of 10^{-3} .

We compare the power-dependent loss of each design in Fig. 5.1, all measured at the base temperature of a dilution refrigerator. In contrast to the previous chapter, we plot the absolute value of the loss Q_i^{-1} normalized by the resonator frequency rather than the relative power-dependent value. We can then use Eqn. 5.1 in combination with the calculated participations to infer intrinsic material loss. While the loss is the sum of the contributions from each dielectric, we will start with the simpler assumption that one source dominates over the others when calculating the intrinsic

²Participation in the Si substrate is also substantial, especially for the planar resonators A-E. We can safely ignore the Si participation as the resonator purely on Si had no power- or temperature-dependent loss (see Fig. 4.12)

material dissipation.

Immediately, we observe that loss of the screened Si_3N_4 resonator is not power-dependent. This eliminates the MA interface from the list of potential sources of the excess loss, as this geometry has vanishing participation in the other dielectrics. Using the calculated participation for the MN and NS interfaces, the loss tangents of those interfaces would need to approach one to model the resonators at the highest pump powers, at which point the interface should behave more like a conductor than an insulator. This would require a process such as dielectric breakdown [Padovani et al. 2024; Sze 1967; Apodaca 2005], which occurs at field strengths of $\sim 10 \text{ MV/cm}$ (assuming the properties of the interfaces are similar to that of the Si_3N_4). This is many orders of magnitude larger than achieved in our devices, where $\bar{n}_e = 10^8$ corresponds to field strengths in the 1 kV/cm range. Additionally, dielectric breakdown is usually a destructive process, yet we observe little to no hysteresis when measuring the power-dependent loss.

This leaves the bulk Si_3N_4 as the last potential candidate. To explain the loss at the highest powers, the Si_3N_4 film would need an intrinsic dissipation of 10^{-2} , which is much larger than the values reported for the low-temperature, single photon regime loss [Paik and Osborn 2010]. Later in this chapter, we will see how the loss of the Si_3N_4 can increase rapidly with temperature and power when self-heating is considered. The loss of the planar, lumped element resonators A-E does scale linearly with F_{SiN} across a range of base plate temperatures T_{bp} and pump powers, shown in Fig. 5.2. Up to misalignment of the capacitor pads, the power-dependent loss of the transducer also follows this general trend. With all of these observations, we conclude with high confidence that the power-dependent loss is a result of electric field participation in the Si_3N_4 .

5.2 Generalized model of circuit loss and frequency shift

The prevailing model for the dielectric properties of amorphous materials at cryogenic temperatures is the standard tunneling model (STM) of two-level systems (TLS) [Phillips 1987]. Amorphous materials can host defect states wherein an electron, atom, or group of atoms within the solid can tunnel between two nearly degenerate configurations, forming a TLS [Müller et al. 2019]. The electric




Figure 5.2: **Loss versus F_{SiN} .** The internal loss of resonators A-E scales linearly with the electric field participation ratio in the Si_3N_4 . This relationship holds at different pump strengths \bar{n}_e and temperatures T_{bp} . Error bars indicate $\pm 1\sigma$ uncertainty and dashed lines are linear fits.

and elastic dipole moments of these TLS couple to external electric or strain fields. In our resonators, the coupling between the electric dipole moments of a bath of TLS and the electric field of the circuit manifests as an additional dissipation channel, as the TLS quickly decay via phonons. The model is derived in detail in a number of modern works, so here we review the architecture of the model, highlighting the assumptions and discussing the main conclusions. [Behunin et al. 2016; Wollack et al. 2021; MacCabe et al. 2020; Emser et al. 2024].

The STM of tunneling states (TS) (a slightly more constrained model) assumes that we can model any defect with an asymmetric double-well potential with an asymmetry energy Δ between the wells and a tunneling energy Δ_0 (see Fig. 5.3). The Hamiltonian of this system can be written using the Pauli operators σ_j in the localised (non-diagonal) basis as:

$$H_{TS} = -\frac{\Delta}{2}\sigma_z + \frac{\Delta_0}{2}\sigma_x. \quad (5.3)$$

The presence of a weak electric field ξ modifies the asymmetry energy between the wells. For TLS that contribute to thermal properties below 1 K, microwave fields at frequencies ≤ 20 GHz, corresponding to wavelengths of order 10 mm. This is much larger than the characteristic size of TLS, which was recently imaged at the $10 \mu\text{m}$ scale [Hegedüs et al. 2024]. Thus, we can treat the TLS as




Figure 5.3: **TLS Hamiltonian** (a) An asymmetric double-well potential for a generalized defect. The ground Ψ_- and excited Ψ_+ states are superpositions of population in either well and separated by the TLS energy E . (b) The electric field of the circuit couples to the ensemble of TLS via their average dipole moment \mathbf{p} , and the TLS can interact with each other with strength J . The predominant method of TLS decay is to a bath of phonon modes $u_{\mathbf{q}}$ and governed by their elastic dipole moment γ . This illustration is adapted from Fig. 3 of [Behunin et al. 2016].

a point-like defect with an electric dipole moment $p_0 \equiv -\frac{1}{2} \frac{\partial \Delta}{\partial \xi}$. Expanding the TLS Hamiltonian to first order in the electric field and moving to the basis that is diagonal in the absence of the electric field, the electric field-TS Hamiltonian is:

$$H_{TS-el} = -\frac{E}{2} \sigma_z + \hbar \omega_e (\hat{b}^\dagger \hat{b} + 1/2) + p_0 \xi_{\text{vac}} \left(\frac{\Delta}{E} \sigma_z + \frac{\Delta_0}{E} \sigma_x \right) (\hat{b}^\dagger + \hat{b}) \quad (5.4)$$

where the TS energy $E = \sqrt{\Delta^2 + \Delta_0^2}$ and we have quantized the resonator's electric field $\delta \xi = \xi_{\text{vac}} (\hat{b}^\dagger + \hat{b})$. Note, the scale of the longitudinal and transverse coupling strengths are related by the structure of the double well potential. The more general two-level systems (TLS) model generalizes the TS model by introducing independent couplings M and D for the transverse and longitudinal interactions:

$$H_{TLS-el} = (D \sigma_z + M \sigma_x) \xi_{\text{vac}} (\hat{b}^\dagger + \hat{b}). \quad (5.5)$$

Throughout this chapter, we will point out where we are making the TLS generalization. Lastly, the defect energies can also be modified by strain fields, in which case the derivation is the exact same, except we use the elastic dipole moment γ and replace the coupling rates by their elastic counterparts. In a real system, the electric field of the resonator is coupled to a large ensemble of TLS, which then decay into a collection of phonon modes $u_{\mathbf{q}}$. The TLS can also interact with each

other with strength J_{ij} . In the following subsections, we will derive the effect of the microwave field on a single TLS, and then integrate over the whole ensemble ignoring TLS-TLS interactions. TLS-TLS interactions become relevant when understanding the thermalization of TLS bath. Our derivation follows the formalism developed in [Emser et al. 2024].

5.2.1 Susceptibility of a single TLS

Using linear response theory, we will derive a generalized electrical susceptibility $\chi(\omega)$, describing the interaction between a single TLS and an AC electric field. The frequency shift $\delta\omega_e$ and damping rate $\kappa_{\text{int,e}}$ experienced by the microwave mode due to the TLS is given by:

$$\delta\omega_e = -\frac{\xi_{\text{vac}}^2}{\hbar} \Re[\chi(\omega_e)] \quad (5.6)$$

$$\kappa_{\text{int,e}} = \frac{2\xi_{\text{vac}}^2}{\hbar} \Im[\chi(\omega_e)] \quad (5.7)$$

where the vacuum electric field amplitude is $\xi_{\text{vac}} = \sqrt{\hbar\omega_e/(2\epsilon_0\epsilon_r V)}$ with ϵ_r the relative dielectric constant ($\epsilon_r = 7.6$ for Si_3N_4) and V the material volume.

To accomplish this, we employ a formula developed by [Trif et al. 2018] for computing the susceptibility of a Hamiltonian in the presence of DC and AC perturbations. The first two terms of Eqn. 5.4 constitute the static Hamiltonian H_0 while the third term contains the time-dependence $V(t) = -\hat{p}\delta\xi(t) = \frac{\partial H_0}{\partial \xi}\delta\xi(t)$. By inspection, we see that

$$\hat{p} = -p_0 \left(\frac{\Delta}{E} \sigma_z + \frac{\Delta_0}{E} \sigma_x \right). \quad (5.8)$$

The dynamics of the density operator $\rho(t)$ of the TLS obey the Heisenberg equations of motion under $H(t) = H_0 + V(t)$ with an additional term to account for the relaxation of the TLS to its environment. The susceptibility can be found by evaluating $\chi(\omega) = \delta\langle\hat{p}\rangle/\delta\xi(t)$ where $\langle\hat{p}\rangle = \text{Tr}[\hat{p}(t)\rho(t)]$:

$$\begin{aligned}
\chi(\omega) &= \sum_n \chi^{(n)}(\omega), \\
\chi^{(n)}(\omega) &= \underbrace{-f_n \frac{\partial^2 \epsilon_n}{\partial \xi^2}}_{\chi_C^{(n)}} - \underbrace{\frac{i\Gamma_{nn}}{\omega + i\Gamma_{nn}} \left(\frac{\partial \epsilon_n}{\partial \xi} \right)^2 \frac{\partial f_n}{\partial \epsilon_n}}_{\chi_D^{(n)}} \\
&\quad - \underbrace{\hbar\omega \sum_{m \neq n} \frac{|\langle m | \frac{\partial H_0}{\partial \xi} | n \rangle|^2}{\epsilon_n - \epsilon_m} \frac{f_n - f_m}{\epsilon_n - \epsilon_m - \hbar\omega - i\hbar\Gamma_{nm}}}_{\chi_{ND}^{(n)}},
\end{aligned} \tag{5.9}$$

in which $|n\rangle$ and ϵ_n are the eigenstates and eigenenergies of the bare Hamiltonian. f_n is the equilibrium occupation of the n -th state and $\hat{\Gamma}$ is the TLS relaxation matrix which describes the rates at which the states decay to the environment. We assume that $\Gamma_1 \equiv \Gamma_{11} = \Gamma_{22}$ and $\Gamma_2 \equiv \Gamma_{12} = \Gamma_{21}$. The sum is over the two TLS states Ψ_+ and Ψ_- with energies $\pm E/2$.

The susceptibility has been suggestively organized into three contributions: the curvature χ_C , diagonal χ_D and non-diagonal χ_{ND} components. The curvature contribution is an adiabatic shift that arises from the second-order perturbation theory expansion. The diagonal contribution is responsible for ‘relaxation’-type damping from the longitudinal coupling, while the non-diagonal term is responsible from ‘resonant’-type damping from the transverse coupling. Evaluating these expressions gives:

$$\chi_C(\omega) = \frac{1}{2} \tanh\left(\frac{E}{2k_B T}\right) \left(\frac{\Delta}{E} \frac{\partial^2 \Delta}{\partial \xi^2} + 4 \frac{\hbar^2 g_x^2}{E \xi_{\text{vac}}^2} \right) \tag{5.10}$$

$$\chi_D(\omega) = \left(\frac{\hbar g_z}{\xi_{\text{vac}}} \right)^2 \frac{1}{k_B T} \frac{\text{sech}^2 \frac{E}{2k_B T}}{1 - i\omega \Gamma_1^{-1}} \tag{5.11}$$

$$\chi_{ND}(\omega) = - \left(\frac{\hbar g_x}{\xi_{\text{vac}}} \right)^2 \tanh\left(\frac{E}{2k_B T}\right) \left[\frac{2}{E} - \frac{1}{E - (\hbar\omega + i\hbar\Gamma_2)} - \frac{1}{E + (\hbar\omega + i\hbar\Gamma_2)} \right] \tag{5.12}$$

where we introduced the coupling rates $g_x = \frac{p_0}{\hbar} \frac{\Delta_0}{E} \xi_{\text{vac}}$ and $g_z = \frac{p_0}{\hbar} \frac{\Delta}{E} \xi_{\text{vac}}$.

5.2.2 Resonant interaction

In any real device, the microwave circuit is coupled to a bath of TLS with some distribution in frequency. This corresponds to integrating over all possible values Δ and Δ_0 with some distribution function in the continuum limit $f(\Delta, \Delta_0)d\Delta d\Delta_0$. Following [Phillips 1987], this distribution function is assumed to be independent of Δ such that $f(\Delta, \Delta_0) = P/\Delta_0$ where P is the density of states.³ For the resonant interaction, we can accurately model the data by ignoring the detailed origin of the TLS and simply integrate E as our only variable over a TLS local density of states $P(E)$, $\sum_{\text{TLS}} \longleftrightarrow \int dV \int dE_{\text{TLS}} P(E)$, which effectively assumes the TLS model as opposed to the TS model.⁴

A key assumption of the STM is a uniform density of states $P(E) = P$ and that the TLS are uniformly distributed in the volume. In amorphous materials, there is likely large inhomogeneous broadening that leads to the TLS density states being roughly uniform within some energy ranges. We will explore the consequences and deviations from the uniform distribution assumption in Sec. 5.5. Taking this approximation, as well as assuming $\partial^2 \Delta / \partial \xi^2 \approx 0$, so that the first term of the χ_{ND} cancels with the χ_C , the non-diagonal susceptibility of the TLS ensemble sampled by the resonator's electric field is:

$$\chi_{ND}(\omega) = \frac{\hbar^2 P V_h}{\xi_{\text{vac}}^2} \int_0^{\omega_{\text{max}}} d\omega_{\text{TLS}} g_x^2 \tanh\left(\frac{E}{2k_B T}\right) \left[\frac{1}{\omega_{\text{TLS}} - (\omega + i\Gamma_2)} + \frac{1}{\omega_{\text{TLS}} + (\omega + i\Gamma_2)} \right] \quad (5.13)$$

where the integral is cutoff at ω_{max} , the maximum transition frequency of the TLS ensemble. The choice of ω_{max} is somewhat arbitrary, and we will define quantities that are independent of this choice shortly. V_h represents the volume sampled by the electric field. Assuming the maximum TLS energy is much larger than the thermal energy, we can use the integral representation of the

³ P has units of inverse energy inverse d -dimensional volume

⁴Integrating the non-diagonal susceptibility over Δ and Δ_0 actually leads to an unbounded value when the TLS energy distribution is assumed to be uniform unless a lower-cutoff value for Δ_0 is introduced.

digamma function Ψ to evaluate the integral:

$$\chi_{ND}(\omega) \approx -2PV_h\overline{M}^2 \left[\Psi \left(\frac{1}{2} + \frac{\hbar\Gamma_2}{2\pi k_B T} + \frac{\hbar\omega_e}{2\pi i k_B T} \right) - \ln \left(\frac{\hbar\omega_{\max}}{2\pi k_B T} \right) \right]. \quad (5.14)$$

We have introduced the orientation- and polarization averaged transverse coupling potential $\overline{M} = \hbar\overline{g}_x/\xi_{\text{vac}}$. Dephasing rates for TLS in silica at 20 mK have been reported to be $\leq 2\pi \times 20$ kHz [Behunin et al. 2016], so the second argument of the digamma function can be safely ignored as $\hbar\Gamma_2/2\pi k_B T \ll 1/2$. We can also remove the dependence on the cutoff frequency by expressing the frequency shift relative to the resonant frequency at zero temperature:

$$\frac{\Delta\omega_e}{\omega_e} \equiv \frac{\omega_e(T) - \omega_e(0)}{\omega_e} = \frac{\delta\omega_e(T)}{\omega_e} - \frac{\delta\omega_e(0)}{\omega_e} \quad (5.15)$$

where $\omega_e = \omega_e(T = \infty)$ represents the resonator frequency in the absence of coupling to TLS. Now, we can evaluate Eqn. 5.6 and Eqn. 5.7 for the resonant interaction:

$$\left(\frac{\Delta\omega_e}{\omega_e} \right)_{\text{res}} = \frac{F \tan \delta_{0,\text{res}}}{\pi} \left[\Re \left\{ \Psi \left(\frac{1}{2} + \frac{\hbar\omega_e}{2\pi i k_B T} \right) \right\} - \ln \left(\frac{\hbar\omega_e}{2\pi k_B T} \right) \right] \quad (5.16)$$

$$Q_{\text{res}}^{-1}(T) = F \tan \delta_{0,\text{res}} \tanh \left(\frac{\hbar\omega_e}{2k_B T} \right). \quad (5.17)$$

The filling fraction $F = V_h/V$ is the same as the participation introduced in Eqn. 5.2 and the loss tangent $\tan \delta_{0,\text{res}} = \pi P \overline{M}^2 / (\epsilon_0 \epsilon_r)$ encodes the intrinsic material parameters from the resonant interaction.

So far, we have worked in the limit of a weak driving field. As the microwave circuit is driven harder, the TLS population can be coherently driven between the two states. This leads to Rabi oscillations of the TLS population with a Rabi frequency $\Omega_R = 2\overline{M}\xi_0/\hbar$ where ξ_0 is the steady state amplitude of the driving field. As the number of photons in the resonator $\bar{n}_e = \frac{2\epsilon_0\epsilon_r\xi_0^2 V}{\hbar\omega_e}$ increases, an individual TLS approaches equal population in either state and can no longer absorb energy from the resonator. Solving the steady-state solutions to the Bloch equations, we find that the loss

inherited by the resonator is modified:

$$Q_{\text{res}}^{-1}(T) = \frac{F \tan \delta_{0,\text{res}}}{\sqrt{1 + \bar{n}_e/n_c}} \tanh\left(\frac{\hbar\omega_e}{2k_B T}\right). \quad (5.18)$$

The "saturation" of the resonant damping begins at the critical photon number $n_c = \frac{\hbar\epsilon_0\epsilon_r V}{2\omega_e M^2 T_1 T_2}$ where $T_1 = 1/\Gamma_1$ ($T_2 = 1/\Gamma_2$) is the TLS relaxation (decoherence) time.

Having derived these expressions, we now take a step back and review the physics of the resonant interaction. As the temperature is increased, each TLS approaches a thermal state with equal population in Ψ_- and Ψ_+ , which prevents the TLS from absorbing energy from the resonator and yields the characteristic "saturation" of Q_{res}^{-1} . For the frequency shift ($\Delta\omega_e/\omega_e$), TLS with energy less than the resonator energy (positive detuning $\omega_e - \omega_{\text{TLS}} > 0$) lead to a blue shift while TLS with energy greater than the resonator energies (negative detuning $\omega_e - \omega_{\text{TLS}} < 0$) lead to a red shift (as expected from the Rabi Hamiltonian). In the zero-temperature limit, when all TLS are weighted equally, coupling to the TLS bath leads to a overall red shift ($\delta\omega_e(T=0)/\omega_e < 0$, but by definition $\Delta\omega_e/\omega_e = 0$). Now as the temperature is increased, the low-frequency TLS become saturated and contribute less strongly to the susceptibility, leading to a net red shift of the resonator frequency from the zero temperature limit ($\Delta\omega_e/\omega_e < 0$). This effect is captured by the first term in [Eqn. 5.16](#). When $2\pi k_B T > \hbar\omega_e$, TLS with energy greater than the resonator become saturated. The resonator eventually experiences a blue shift from the zero temperature limit ($\Delta\omega_e/\omega_e > 0$) as it approaches the limit of being uncoupled from the TLS ($T \rightarrow \infty$) which is captured by the second term in [Eqn. 5.16](#).

By looking at the structure of the integrand in [Eqn. 5.13](#), we can also extract which sections of the TLS energy distribution determine the scale of the dissipation and frequency shift. In [Fig. 5.4](#), we see that the frequency shift samples a much wider distribution of TLS energies than does the loss. This is because the real part of the non-diagonal susceptibility scales as $1/(\omega_{\text{TLS}} - \omega_e)$ while the imaginary part scales as $1/(\omega_{\text{TLS}} - \omega_e)^2$. The loss tangent $\tan \delta_{0,\text{res}}$ extracted from the frequency shift and loss are then only equivalent under the assumption that the TLS energy distribution




Figure 5.4: **TLS energies sampled by resonant interaction.** The real (blue) and imaginary (red) parts of the non-diagonal susceptibility χ_{ND} is plotted as a function of TLS frequency ω_{TLS} in the zero-temperature limit. Each curve is normalized to its own maximum positive value. Note that we plot the $-\Re[\chi_{ND}]$, which is a proxy for the frequency shift (Eqn. 5.6).

is uniform. The energy dependence of the imaginary part also helps explain why the frequency shift is not power-dependent. At large driving fields, there will only be an appreciable electric field at frequencies within the linewidth of the resonator. The frequency shift, however, samples a distribution much wider, so saturation of the TLS closest in energy to the resonator will have little effect on the resonator frequency.⁵

5.2.3 Relaxation interaction

For the relaxation interaction, we cannot simply integrate over the TLS energy. Instead, we perform a change of variables from Δ and Δ_0 to the TLS energy E and the energy relaxation time T_1 , which will be a much more convenient set of variables. We assume that TLS relaxation is dominated by single-phonon processes, even in the presence of TLS-TLS interactions [Burnett et al. 2014; Faoro and Ioffe 2015a]. The relaxation time can be calculated from Fermi's golden rule for the probability $p_{g \rightarrow e}$ that a thermal phonon excites the TLS:

$$T_1^{-1} = p_{g \rightarrow e}(1 + e^{E/k_B T}) \quad (5.19)$$

⁵This is true in the continuum limit, with many weakly coupled TLS. In the limit of a few strongly coupled TLS, saturation of the TLS closest to ω_e does impact the frequency shift.

$$p_{g \rightarrow e} = \sum_{\alpha} \frac{2\pi}{\hbar} |\langle g | H_{TLS-ph} | e \rangle|^2 \frac{g_d(\omega_{TLS})}{\hbar} N(\omega_{TLS}) \quad (5.20)$$

where is the TLS-phonon interaction Hamiltonian H_{TLS-ph} is the same as H_{TLS-el} except with a strain field and elastic dipole moments. g_d is the Debye density of states for a d -dimensional phonon bath and $N(\omega_{TLS})$ is the Bose-Einstein factor describing the thermal phonon distribution at ω_{TLS} . The sum is over all phonon polarizations α . The matrix element is given by the position- and orientation-averaged elastic transverse coupling \overline{M}_{γ} times the vacuum strain fluctuations. In total, the relaxation rate to due to one-phonon processes in d dimensions after averaging over polarizations and orientations is given:

$$T_1^{-1}(E) = \frac{\overline{M}_{\gamma}^2}{\overline{v}^{d+2}} \frac{\pi S_{d-1}}{(2\pi)^d} \frac{E^{d-2} \Delta_0^2}{\hbar^{d+1} \rho_d} \coth\left(\frac{E}{2k_B T}\right) \quad (5.21)$$

where \overline{v} is the polarization-averaged acoustic velocity, S_d is the d -dimensional unit-hypersurface area, and ρ_d is the d -dimensional density. This is the same as Eq. (7) in [Behunin et al. 2016]. The relaxation time is minimized when $E = \Delta_0$, or when the double well potential is symmetric, giving $T_{1,\min}^{-1}(E) = A_d E^d \coth(E/2k_B T)$. All of the unitful prefactors have been lumped in A_d for succinctness. The dimensionality of the phonon bath is set by the temperature and geometry of the dielectric. In particular, phonon modes can become frozen out at sufficiently low temperatures, which can reduce the effective dimensionality.

With the expression for the relaxation time in hand, we can use the Jacobian transformation to perform a change of variables:

$$\frac{\partial E}{\partial \Delta} = \frac{\Delta}{E}, \quad \frac{\partial T_1}{\partial \Delta_0} = -\frac{2T_1}{\Delta_0} \quad (5.22)$$

$$g(E, T_1) dE dT_1 = f(\Delta, \Delta_0) \frac{\Delta_0 E}{2\Delta T_1} d\Delta d\Delta_0 \quad (5.23)$$

$$g(E, T_1) = -\frac{P}{2T_1 \sqrt{1 - T_{1,\min}(E)/T_1}} \quad (5.24)$$

where $g(E, T_1)$ is the distribution function in energy and relaxation time. With these new integration

variables and after assuming uniformity over the TLS volume V_h , the diagonal susceptibility becomes:

$$\chi_D(\omega) = \frac{\hbar^2}{\xi_{\text{vac}}^2 k_B T} \int_0^\infty dE \int_{T_{1,\min}(E)}^\infty dT_1 g_z^2 \frac{\text{sech}^2 \frac{E}{2k_B T}}{1 - i\omega\Gamma_1^{-1}} g(E, T_1). \quad (5.25)$$

Unlike the resonant interaction, we will not neglect the energy dependence of the coupling term (ie, we use the TS model).

$$g_z^2 = \frac{p_0^2 \xi_{\text{vac}}^2}{\hbar^2} \frac{\Delta^2}{E^2} = \frac{p_0^2 \xi_{\text{vac}}^2}{\hbar^2} \left(1 - \frac{T_{1,\min}(E)}{T_1} \right) \quad (5.26)$$

Collecting terms, the loss and fractional frequency shift due to the relaxation interaction is given by plugging the diagonal susceptibility into [Eqn. 5.7](#) and [Eqn. 5.6](#):

$$Q_{\text{rel}}^{-1} = \frac{FP\bar{p}_0^2 \omega_e}{2\epsilon_0 \epsilon_r k_B T} \int_0^\infty dE \int_{T_{1,\min}(E)}^\infty dT_1 \text{sech}^2 \left(\frac{E}{2k_B T} \right) \sqrt{1 - \frac{T_{1,\min}(E)}{T_1}} \frac{1}{1 + \omega_e^2 T_1^2} \quad (5.27)$$

$$\left(\frac{\Delta\omega_e}{\omega_e} \right)_{\text{rel}} = \frac{-FP\bar{p}_0^2}{4\epsilon_0 \epsilon_r k_B T} \int_0^\infty dE \int_{T_{1,\min}(E)}^\infty dT_1 \text{sech}^2 \left(\frac{E}{2k_B T} \right) \frac{1}{T_1} \sqrt{1 - \frac{T_{1,\min}(E)}{T_1}} \frac{1}{1 + \omega_e^2 T_1^2}. \quad (5.28)$$

We again define the $\Delta\omega_e/\omega_e$ as the frequency shift relative to the resonator frequency at zero temperature. This coincides with $\delta\omega_e/\omega_e$ in this case as the frequency shift from relaxation damping at zero temperature is zero ($\delta\omega_e(T=0)/\omega_e = 0$). When we fit the circuit loss and frequency shift from this effect, we evaluate this expression in entirety. To gain intuition, the relaxation interaction is often considered in limit either the $\omega_e T_1 \gg 1$ or $\omega_e T_1 \ll 1$ corresponding to long and short-lifetime TLS.

At temperatures $k_B T \lesssim \hbar\omega_e$, TLS decay is suppressed as there are few thermal phonons which can induce relaxation of the defects (captured by the $\coth(1/T)$ dependence of [Eqn. 5.21](#)). The TLS decay time can thus grow large at low large at low temperatures ($\omega_e T_1 \gg 1$) and we can approximate $1/(1 + \omega_e^2 T_1^2) \approx 1/(\omega_e^2 T_1^2)$. Then, the integrals over T_1 can be evaluated using the following identity:

$$\int_a^\infty dx \frac{1}{x^d} \sqrt{1 - \frac{a}{x}} = \frac{\sqrt{\pi}}{2} \frac{a^{1-d} \Gamma(d-1)}{\Gamma(d + \frac{1}{2})} \quad (5.29)$$

where $\Gamma(d)$ is the gamma function. Using the expression for T_1 and performing a change of variables $x = E/(2k_B T)$:

$$\lim_{\omega_e T_1 \gg 1} Q_{\text{rel}}^{-1} = \frac{FP\bar{p}_0^2 \omega_e}{3\epsilon_0 \epsilon_r} \frac{\bar{M}_\gamma^2}{\bar{v}^{d+2}} \frac{S_{d-1}}{(2\pi)^{d-1} \hbar^{d+1} \rho_d} (2k_B T)^d \int_0^\infty dx x^d \text{sech}^2(x) \coth(x) \quad (5.30)$$

$$\lim_{\omega_e T_1 \gg 1} \left(\frac{\Delta \omega_e}{\omega_e} \right)_{\text{rel}} = \frac{-2FP\bar{p}_0^2}{15\epsilon_0 \epsilon_r} \left(\frac{\bar{M}_\gamma^2}{\bar{v}^{d+2}} \frac{\pi S_{d-1}}{(2\pi)^{2d}} \frac{1}{\hbar^{d+1} \rho_d} \right)^2 (2k_B T)^{2d} \int_0^\infty dx x^{2d} \text{csch}^2(x) \quad (5.31)$$

where the remaining integrals just evaluate to a number. The expressions contain both the electric and elastic dipole moments because the TLS absorbs electrical energy from the resonator and subsequently decays to phonons. We can pick out that the loss obeys a T^d power-law, while the frequency shift follows a T^{2d} power-law. For typical values of the material parameters, the frequency shift from the relaxation damping in this limit is negligible compared to the that of the resonant damping. The microwave circuit loss in this limit is more conveniently written by choosing some reference temperature T_0 :

$$\lim_{\omega_e T_1 \gg 1} Q_{\text{rel}}^{-1} = F \tan \delta_{0,\text{rel}} \left(\frac{T}{T_0} \right)^d. \quad (5.32)$$

Note, by choosing a reference temperature, we are effectively selecting an energy scale at which the material loss tangent from this interaction $\tan \delta_{0,\text{rel}}$ is defined. All of the pre-factors in [Eqn. 5.30](#) get wrapped into $\tan \delta_{0,\text{rel}}/T_0^d$.

As the temperature of the system increases, the decay time of the TLS decreases as more phonons become thermally activated. There is also the potential for higher order phonon processes to further enhance the decay rate. In the limit $\omega_e T_1 \ll 1$, we approximate $T_{1,\text{min}} \rightarrow 0$. This greatly simplifies the evaluation of the integrals and we find that the loss reduces to a temperature- and dimension-independent value while the frequency shift behaves logarithmically with temperature and linearly with phonon bath dimension:

$$\lim_{\omega_e T_1 \ll 1} Q_{\text{rel}}^{-1} = \frac{FP\bar{p}_0^2 \pi}{2\epsilon_0 \epsilon_r} \quad (5.33)$$

$$\lim_{\omega_e T_1 \ll 1} \left(\frac{\Delta\omega_e}{\omega_e} \right)_{\text{rel}} = -\frac{FP\bar{p}_0^2}{\epsilon_0\epsilon_r} \frac{d}{2} \ln(T/T'_0) \quad (5.34)$$

with T'_0 a reference temperature for the frequency shift. In practice, we will choose $T_0 = T'_0$ for simplicity. The rate of the frequency shift with temperature is set by the dimensionality of the phonon bath, and the resonator experiences a red shift that will eventually overwhelm the blue shift from the resonant term.

The crossover between these two regimes happens at $\omega_e T_{1,\min} \sim 1$. To get a sense of the temperature at which this crossover occurs, we point out that integral over energy strongly weights $E = 2k_B T$ due to the $\text{sech}^2(E/2k_B T)$. Thus at the crossover temperature T_{cross} , we can evaluate the expression for $T_{1,\min}$:

$$\omega_e \sim T_{1,\min}^{-1}(T_{\text{cross}}) = A_d E^d \coth(E/2k_B T) = A_d (2k_B T_{\text{cross}})^d \rightarrow k_B T_{\text{cross}} \approx \frac{1}{2} (\omega_e / A_d)^{1/d} \quad (5.35)$$

which suggests that the crossover temperature is a sublinear function of the resonator frequency and acoustic material parameters for $d > 1$. Thus, resonators at significantly different frequencies using dielectrics with different material parameters will have similar crossover temperatures. This could help explain why the temperature scale at which TLS damping is observed is $\lesssim 1$ K across a wide range of resonator frequencies and amorphous materials [Phillips 1987].

Let's again step back and understand the physics modeled by the derived expressions. Unlike the resonant interaction, relaxation damping does not involve any direct energy exchange with the microwave resonator. Instead, the defects are driven out of thermal equilibrium with the phonon bath by the electric field of the resonator. They then absorb a phonon from the bath and decay in a random direction, leading to an irreversible dissipation event. Relaxation damping is also not saturable in the same sense as the resonant damping (the loss is a monotonically increasing function of temperature). As more TLS become thermally populated, the loss increases as more defects can be pushed away from equilibrium by the resonator field. The loss eventually "saturates" (now in the sense that it stops increasing with the power law behavior) at sufficiently high temperatures to the value given by Eqn. 5.33.

To understand which TLS are sampled by the relaxation interaction, we plot the integrand of the non-diagonal susceptibility after numerically evaluating the integral over T_1 for several temperatures in Fig. 5.5 assuming a $d = 2$ dimensional phonon bath. It's clear that relaxation damping samples a much wider distribution of TLS energies than resonant damping. The distribution sampled by the imaginary contribution has a center and characteristic width that is proportional to $2k_B T$, so the distribution moves to higher energy and broadens with increasing temperature.⁶

The width of these distributions also explains why the relaxation term is not power-dependent in a manner similar to the frequency shift from resonant damping. Only the defects with energies within the linewidth of the microwave resonator will experience an appreciable electric field, which are a small subset of the TLS energies that influence relaxation damping. Additionally, the rate of energy loss from the resonator is governed by the decay rate of the TLS, which is mediated by phonon interactions that should not depend on the microwave field. In Sec. 5.2.5, we show that relaxation damping can still influence power-dependent loss if there is significant self-heating.

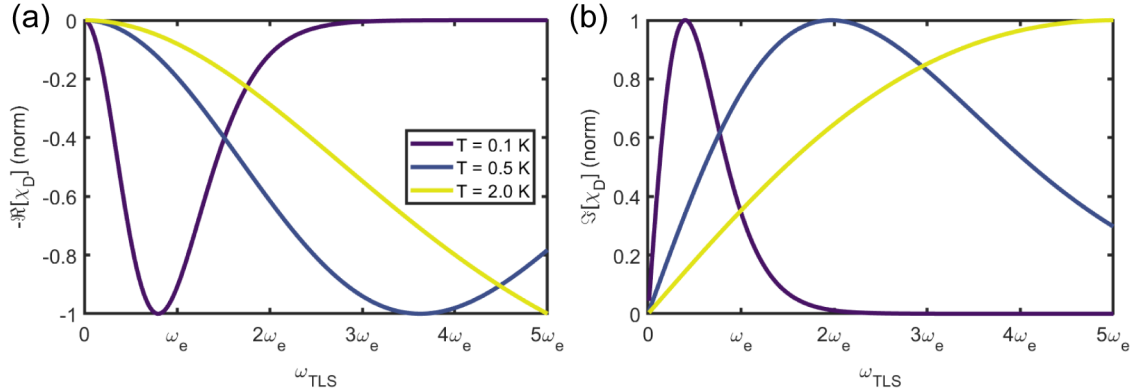


Figure 5.5: **TLS energies sampled by relaxation interaction.** Real (a) and Imaginary (b) parts of the diagonal susceptibility χ_D plotted as a function of TLS frequency ω_{TLS} for several temperature values, assuming a $d = 2$ dimensional phonon bath. Each curve is normalized to its own maximum value. Note that we plot the $-\Re[\chi_D]$, which is a proxy for the frequency shift (Eqn. 5.6).

⁶A similar effect is seen in three-dimensions, but in one-dimension, the distribution is always centered on zero energy and just the width increases with temperature

5.2.4 Extracting material properties

Including the TLS damping mechanisms, we will describe the loss and fractional frequency shift as the sum of several contributions:

$$Q_i^{-1}(\bar{n}_e, T) = Q_{\text{bg}}^{-1} + Q_{\text{res}}^{-1}(\bar{n}_e, T) + Q_{\text{rel}}^{-1}(T) + Q_{\text{qp}}^{-1}(T) \quad (5.36)$$

and

$$\left(\frac{\Delta\omega_e}{\omega_e}\right)(T) = \left(\frac{\Delta\omega_e}{\omega_e}\right)_{\text{res}}(T) + \left(\frac{\Delta\omega_e}{\omega_e}\right)_{\text{rel}}(T) + \left(\frac{\Delta\omega_e}{\omega_e}\right)_{\text{qp}}(T) \quad (5.37)$$

where Q_{bg}^{-1} is a constant value of background loss and the thermal quasiparticle contributions are parameter-free estimates using independently measured material properties of superconductor (see [Sec. 4.3.2](#)).

To study the power- and temperature-dependent loss and frequency shift, we excite the resonant circuits with a variable-strength microwave drive tone detuned by 1.5 MHz, while measuring the electrical resonator response using a network analyzer emitting a probe tone with at most 1% of the power of the drive. This mimics the operation of the electromechanical resonators. Each design is inductively coupled to a common on-chip transmission line, and the device is mounted to the base plate of a dilution refrigerator and cooled to a base temperature of $T_{\text{bp}} = 10$ mK. The drive and probe tones access the device through an array of attenuators and cables in the fridge whose microwave transmission response is measured prior to cooling down so that we may calibrate the intracavity photon number \bar{n}_e . To measure the temperature-dependent loss, we hold the drive power fixed and vary the base plate temperature. We fit the frequency-dependent transmission response to extract the resonator's internal loss Q_i^{-1} , external coupling Q_e^{-1} , and center frequency ω_e [[Khalil et al. 2012](#); [Probst et al. 2015](#)] (see [Appendix B](#)).

The power- and temperature-dependent loss and frequency shift measured for the capacitive Si_3N_4 resonators (see [Fig. 4.11](#) for schematics) over a broad range of power and temperatures is plotted in [Fig. 5.6](#) along with fits to the TLS-only model. Empirically, the temperature-dependent




Figure 5.6: **Power- and temperature-dependent properties of Si_3N_4 .** (a-b) Power- and temperature-dependent internal loss of designs A-E. (c-d) Power- and temperature-dependent fractional frequency shift. The TLS model captures the temperature-dependent effects very well, but fails to model the power-dependent effects at $\bar{n}_e > 10^5$. Error bars indicate $\pm 1\sigma$ uncertainty.

loss above $T_{\text{bp}} = 500 \text{ mK}$ was well modeled by a quadratic, so we fix $d = 2$ in the model. We fit a global model to the loss, with the background loss of each resonator, the resonant loss tangent $\tan \delta_{0,\text{res}}$, the relaxation loss tangent $\tan \delta_{0,\text{rel}}$, the critical photon number n_c , and the pre-factor in the expression for the minimum TLS decay time A_2 as free-parameters. The fractional frequency shift is fit independently of the loss.

The temperature-dependent loss is well captured by our combined model (Eqn. 5.36). There is the characteristic saturation of the resonant damping, and then the power-law dependence from relaxation damping. The highest temperature points $T_{\text{bp}} > 1 \text{ K}$ also begin to deviate from the quadratic dependence as we approach the high-temperature limit of the relaxation damping. There is also reasonable agreement between the model and the temperature-dependent frequency shift. In addition, the measured power-dependent loss and frequency shift is predicted well by the TLS model for low drive powers. We extract $\tan \delta_{0,\text{res}} = (1.4 \pm 0.1) \times 10^{-3}$ and $\tan \delta_{0,\text{rel}} = (3.4 \pm 0.1) \times 10^{-3}$ with $T_0 = 0.5 \text{ K}$ for the material loss tangents from resonant and relaxation damping, respectively.⁷ For reasons still unclear to us, this resonant loss tangent is 1-2 orders of magnitude larger than that reported previously for high-stress, stoichiometric Si_3N_4 thin films deposited via low-pressure chemical vapor deposition [Paik and Osborn 2010].

5.2.5 Self-heating

The TLS-only model, however, fails to capture the power-dependent effects for $\bar{n}_e > 10^5$. The loss begins to increase rapidly, as measured in the transducer (Fig. 4.7). The frequency shift experiences a small blue shift, followed by a rapid red shift. Interestingly, these power-dependent behaviors look qualitatively similar to the temperature-dependent effects predicted by relaxation damping. In addition, we notice that the power-dependent loss of a single resonator eventually collapses onto a single curve at each temperature in Fig. 5.7. This suggests that both the base plate temperature and power dissipated by the resonator can lead to heating of the relevant degrees of

⁷The relaxation loss tangent is defined by fitting the full expression Eqn. 5.27 and then taking the low temperature limit and choosing a reference temperature as in Eqn. 5.32.

freedom and that their might be an effective temperature which captures both effects.




Figure 5.7: **Evidence of self-heating.** Power-dependent loss of design A as a function of temperature. For each temperature, the power-dependent loss falls along the same curve for sufficiently strong pumping. Error bars indicate $\pm 1\sigma$ uncertainty.

To that end, we can define an effective temperature that is elevated above the baseplate temperature by an amount T_{heat} due to power dissipated inside the microwave resonator:

$$T_{\text{eff}}(\bar{n}_e) = T_{\text{bp}} + T_{\text{heat}}(\bar{n}_e). \quad (5.38)$$

Given a measured value of $Q_i^{-1}(\bar{n}_e, T_{\text{bp}} = 10 \text{ mK})$, we can find the temperature T_{eff} that would result in the same loss rate using the fits to the temperature-dependent loss with the TLS-only model. Carrying this out for each resonator design, we see that inferred temperatures seem to collapse onto a single trend line in Fig. 5.8, lending further support for the heating model. Designs with more Si_3N_4 have larger microwave loss, and thus dissipate more power into the phonon bath. The power dissipated per unit volume is constant across designs as the Si_3N_4 participation F_{SiN} scales proportionately with the Si_3N_4 volume. Fitting all of the data to a single power-law $T_{\text{eff}}(\bar{n}_e) = A\bar{n}_e^\beta + T_{\text{bp}}$, we found empirically that $\beta \approx 0.5$ and subsequently fixed this exponent in the fitting, extracting $A = 7.97 \times 10^{-4}$. For a self-consistency check, we can supplement the predictions of the TLS model with the self-heating $Q_i^{-1}(\bar{n}_e, T_{\text{eff}}(\bar{n}_e))$ and $\left(\frac{\Delta\omega_e}{\omega_e}\right)(T_{\text{eff}}(\bar{n}_e))$. The TLS and self-heating model reasonably predicts the power-dependent effects for strong drive strengths. Even though the self-

heating is parameterized by the circuit loss, it provides an approximate prediction of the frequency shift as well.

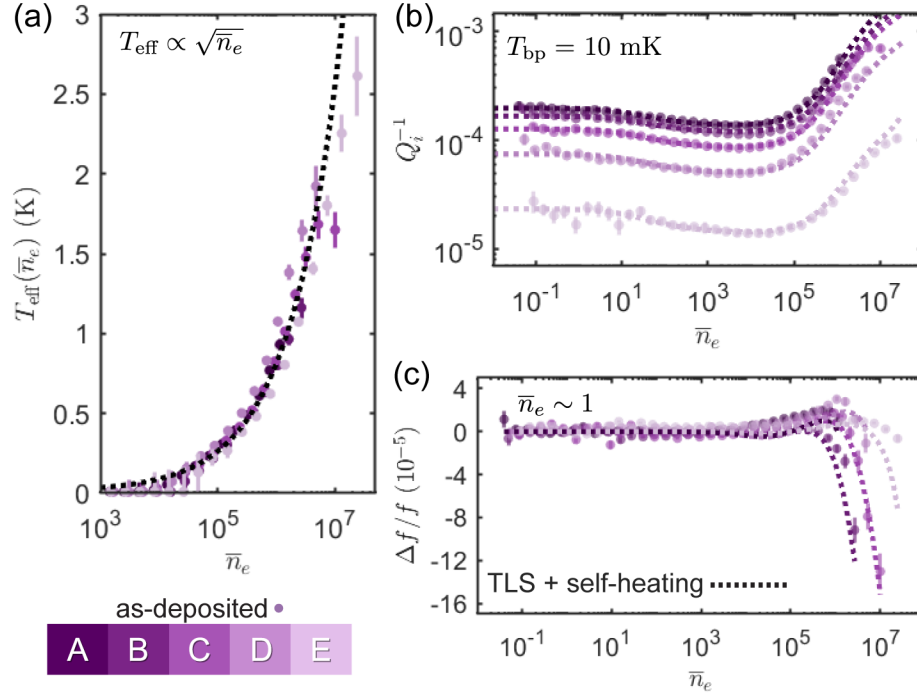


Figure 5.8: **TLS and self-heating model.** (a) T_{eff} for each design collapses to a single trend, which we find approximately obeys $\sqrt{\bar{n}_e}$. (b-c) Self-consistency check of heating model that adapts the TLS model with a power-dependent temperature. The self-heating is calculated with the measured dissipation, but provides reasonable agreement with the frequency shift as well. Error bars indicate $\pm 1\sigma$ uncertainty.

The fits to the TLS and self-heating model provide valuable insights into features of the TLS bath. We noticed that the relaxation damping scaled as T^2 , which is indicative of a two-dimensional phonon bath. This could be plausible for the transducer, where the TLS live in a two-dimensional, suspended membrane but these planar test devices sit on a $380\text{ }\mu\text{m}$ thick Si chip. The Si_3N_4 , however, is patterned in the same geometry as the interdigitated fingers ($5\text{ }\mu\text{m}$ wide and $220\text{ }\mu\text{m}$ long rectangles), and is only 100 nm thick. At low temperatures, phonon modes can become frozen out, which could adjust the dimensionality of the phonon bath. The thermal phonon wavelength can be found by $\lambda_{ph} \approx \frac{h v_{ph}}{2 k_B T}$, where $v_{ph} = 8000\text{ m/s}$ is the phonon velocity in Si_3N_4 .⁸ The phonon bath could look two-dimensional for temperatures $0.05\text{ K} \lesssim T \lesssim 2\text{ K}$ as thermal phonons with wavelength less than 100 nm are frozen out, but not those with wavelength less than $5\text{ }\mu\text{m}$. At higher temperatures, the bath would look three-dimensional and at lower temperatures, one-dimensional. We cannot resolve these crossovers in the resonator loss data as relaxation damping is negligible compared to resonant damping the low-temperature limit ($T \ll T_0$), and independent of bath dimensionality in the high-temperature limit (Eq. 5.33). Similar observations of a mismatch between the expected and measured phonon bath dimensionality of the relaxation damping was also reported by [Wollack et al. 2021] and is explored further in [Behunin et al. 2016].

The introduction of the self-heating effect raises the question: What is getting hot? Our analysis suggests significant heating of the relevant degrees of freedom, up to more than 2 K with $\bar{n}_e = 10^7$. In principle, $T_{\text{eff}}(\bar{n}_e)$ could reflect the temperature of either the whole Si chip, the Si_3N_4 film, or the bath of TLS. In these measurements, the resonators are all on the same chip, coupled inductively to a common transmission line. We observe no power-induced loss when driving one circuit and measuring the loss of another circuit on the same chip, indicating the Si chip remains cold.

The Si_3N_4 film could heat independently of the chip if there is an appreciable boundary (Kapitza) thermal resistance at the Si- Si_3N_4 interface [Swartz and Pohl 1989; S.-M. Lee and Cahill

⁸ $v_{ph} = \sqrt{K/\rho}$ where K is the bulk modulus. K is related to Young's modulus E and Poisson's ratio ν by $K = \frac{E}{2(1-2\nu)}$ for isotropic materials. Here, we used $E = 250\text{ GPa}$, $\nu = 0.27$, and $\rho = 3100\text{ kg/m}^3$.

1997], as thermal conduction of the NbTiN is suppressed in the superconducting state. Given a temperature discontinuity ΔT across an interface with area A , the thermal boundary resistance is given

$$R_{\text{Bd}} = \frac{\Delta T}{P_{\text{diss}}/A} \quad (5.39)$$

where $P_{\text{diss}} = \hbar\omega_e^2 \bar{n}_e Q_{\text{TLS}}^{-1}$ is the power dissipated in the Si_3N_4 by the interaction with TLS ($Q_{\text{TLS}}^{-1} = Q_i^{-1} - Q_{\text{bg}}^{-1} - Q_{\text{qp}}^{-1}$). In order for this effect to explain all the heating, we find that $R_{\text{Bd}}(T = 1 \text{ K}) \sim 2 \times 10^6 \text{ K cm}^2/\text{W}$, which is much larger than direct measurements of this property. At $T = 300 \text{ K}$, $R_{\text{Bd}} = 2 \times 10^{-4} \text{ K cm}^2/\text{W}$ has been reported for the $\text{Si}/\text{Si}_3\text{N}_4$ interface [Bai et al. 2009]. The dimensionality of the phonon bath d is theorized to determine the dependence of the $R_{\text{Bd}} \sim T^{-d}$. To compare the room-temperature value with our measurement, we will assume that the phonon bath is three-dimensional between $1 \text{ K} \leq T \leq 300 \text{ K}$. The room-temperature value thus predicts $R_{\text{Bd}}(T = 1 \text{ K}) = 5 \times 10^3 \text{ K cm}^2/\text{W}$, still three orders-of-magnitude smaller than required. Additionally, calculations using the acoustic or diffuse mismatch model predict $R_{\text{Bd}}(T = 1 \text{ K}) = 10^2 - 10^3 \text{ K cm}^2/\text{W}$ [Swartz and Pohl 1989], consistent with the values measured in the literature but not with the magnitude of heating observed in our resonators.

This leaves the possibility that the TLS bath itself is at an elevated temperature, which can occur if the TLS-TLS interaction are stronger than the TLS-phonon interaction ($J_{ij} > \gamma$) [Black and Halperin 1977]. The energy of one TLS will be modified by the energy of every other defect nearby due to their electric dipole coupling:

$$H_{\text{TLS-TLS}} = \sum_j \sum_{i \neq j} \frac{1}{2} J_{ij} \sigma_{z,i} \sigma_{z,j}. \quad (5.40)$$

The interaction rate J_{ij} depends on the distance between defects r_{ij} , scaling as $J_{ij} \propto r_{ij}^{-3}$ for three-dimensional systems.⁹ Thus, if there is a large density of defects, interactions between TLS can become appreciable, which is plausible considering the large density of hydrogen defects measured in Sec. 5.4. Additionally, [Würger and Bodea 2004] found that the thermal conductivity due to TLS

⁹We expect our system to be three-dimensional for photons, as those are not frozen out at mK temperatures

interactions scales as $T^{4/3}$. The rate at which heat is removed from the TLS bath can be expressed:

$$P_{\text{out}} = \int_{T_0}^{T_{\text{eff}}} \kappa_{th}(T) dT, \quad (5.41)$$

so if $\kappa_{th}(T) \propto T^{1.3}$, then the rate of heating leaving the system should scale as $P_{\text{out}} = T_{\text{eff}}^{1+1.3}$. Assuming equilibrium, the microwave power dissipated in the TLS bath $P_{\text{diss}} \propto \bar{n}_e$ is equal to rate at which leaves the system, implying that $T_{\text{eff}} \propto \bar{n}_e^{1/2.3} \sim \bar{n}_e^{1/2}$. Note, this is also possible if the phonon bath to which the TLS are coupled is one-dimensional, but this is not supported by the data.

We believe this self-heating effect could manifest in other material platforms. Similar power-dependent loss increases have been observed in amorphous silicon [Shu et al. 2021], crystalline silicon [Checchin et al. 2022], and sapphire [Read et al. 2023b]. The heating non-linearity is of particular importance in systems where pathways for thermalization are limited, such as the suspended membrane of the transducer. In the next chapter, we will explore how the heating of the TLS bath could contribute to noise in the microwave circuit of our transducer.

5.3 Annealing reduces microwave-frequency dielectric loss

Having identified the source of the power- and temperature-dependent loss and the physics underpinning these effects, we now look to improve the dielectric properties of our Si_3N_4 . Previous work found that hydrogen defects left over from the low-pressure chemical vapor deposition (LPCVD) gasses (dichlorosilane and ammonia) were correlated with increased microwave loss in the single-photon, low temperature regime where the resonant interaction dominates [Paik and Osborn 2010]. In discussions with collaborators Scott Papp and David Carlson at NIST Boulder, we learned that hydrogen defects are significant source of loss in Si_3N_4 optical waveguides in the near-infrared [Luke et al. 2013; Ji et al. 2021]. At these wavelengths, the N-H and Si-H bonds have vibrational modes that directly absorb optical energy. High-temperature, post-deposition annealing of the Si_3N_4 is a common technique in that field to remove the hydrogen impurities, and has enabled infrared optical resonators with Q surpassing one million.




Figure 5.9: **Power- and temperature-dependent properties of annealed Si_3N_4 .** (a-b) Power- and temperature-dependent internal loss of designs A-E, now fabricated with annealed Si_3N_4 . (c-d) Power- and temperature-dependent fractional frequency shift. The loss at all powers and temperatures is improved, and there is no longer any self-heating evident, as all of the data is well-described by the TLS model. Error bars indicate $\pm 1\sigma$ uncertainty.

At first glance, it is not clear if annealing will improve the excess microwave power- and temperature-dependent loss that we observe. The N-H and Si-H vibrational modes have energies approaching 100 THz, meaning they are unlikely to be the TLS determining the dielectric properties at 6 GHz. Thus, it is astonishing that when we fabricate the same resonator designs, except now performing a four-hour, 1100 C anneal in an oxygen atmosphere,¹⁰ the power- and temperature-dependent loss of the microwave circuit is dramatically improved. In Fig. 5.9, we also see that there is no longer any need for any self-heating assumption. The dielectric properties follow the predictions from the TLS-only model, including continued saturation of the resonant damping with power and temperature and no power-dependent frequency shift.

For the material loss tangents, we find $\tan \delta_{0,\text{res}} = (4.8 \pm 0.4) \times 10^{-4}$ and $\tan \delta_{0,\text{rel}} = (8 \pm 8) \times 10^{-6}$ with reference temperature $T_0 = 0.5$ K for the annealed Si_3N_4 . This represents about a factor of three improvement to the resonant damping, and essentially a complete removal of the relaxation damping. Optimistically, the loss tangent we extract is consistent with zero, and conservatively, annealing improves the relaxation damping by 60 times. A full list of the fitted resonator parameters is in Appendix B.2.

Perhaps most importantly, this improvement is also observed in the transducer geometry (Fig. 5.10). Our TLS and self-heating model can also fully capture the power- and temperature-dependent loss measured in the transducer used in [Brubaker et al. 2022] and [R. Delaney 2022]. The effective temperature now follows $T_{\text{eff}}(\bar{n}_e) = 1.45 \times 10^{-4} \sqrt{\bar{n}_e} + T_{\text{bp}}$, so the self-heating effect is 5.5 times weaker in the transducer geometry. This is due to the different electric field distributions. Given the same \bar{n}_e , there is a smaller electric field in the Si_3N_4 of the transducer geometry and thus less energy dissipated per unit volume.

In addition to annealing the Si_3N_4 , we design the motional capacitor pad to be larger than the capacitor pads on the circuit chip, partially screening the Si_3N_4 from the electric field of the parallel plate capacitor. Averaging over potential misalignment, we expect F_{SiN} to be reduced by a factor of four. We also deposit the Si_3N_4 with a different ratio (7:1) of precursor gasses in order to maintain

¹⁰Oxygen atmosphere is recommended by the furnace manufacturer, Tystar, to better drive out hydrogen impurities

tensile stress in the membrane (Sec. 5.3.1). Infrared spectroscopy measurements of films deposited with this ratio of gases reveal the same hydrogen defect density as films deposited with the standard 3:1 recipe (Sec. 5.4). We thus expect annealing, not deposition gas ratios, to be the key step in optimizing the dielectric properties of high-stress LPCVD Si_3N_4 . The combination of annealing and screening solve one of the two effects that limited the transducer demonstrations in Ch. 4.

5.3.1 Tensile stress of annealed Si_3N_4

Residual Stress (GPa)	3:1 NH_3 :DCS	7:1 NH_3 :DCS
as-deposited	1.04 ± 0.01	1.12 ± 0.06
annealed	-	0.84 ± 0.02

Table 5.2: **Annealing reduces tensile stress of Si_3N_4 .** We observe a $\sim 25\%$ reduction in the tensile stress of Si_3N_4 deposited with a 7:1 ratio of the precursor gases. While the stress of the annealed, 3:1 film was not measured, we expect a similar reduction from the stress of the as-deposited 3:1 film, which would correspond to a stress of 750 MPa.

Annealing is known to change the structural properties of Si_3N_4 [Kurita et al. 2022; Parkhomenko et al. 2023], which could negatively impact the mechanical properties of the membrane. In particular, the membrane frequency and intrinsic dissipation depend on a large tensile stress. The structural properties of the nitride can also change based on the ratio of the precursor gases, ammonia (NH_3) and dichlorosilane (DCS) [Tönnberg 2006]. Our collaborator at NIST suggested we try depositing the film with a 7:1 ratio of NH_3 :DCS, as opposed to our usual 3:1 ratio.

To measure the tensile stress of the film, we deposit a 100 nm layer onto both sides of a bare Si wafer. After taking a reference scan of the wafer with a k-space MOS Ultra-scan tool [kSA MOS UltraScan/ ThermalScan 2024] which measures deformation of the wafer, we etch the Si_3N_4 from one side of the wafer. The now asymmetric material stack will result in an appreciable wafer bow, which is again measured by the Ultra-scan tool to extract the tensile stress. Table 5.2 summarizes the result of these measurements. Annealing reduces the stress of the as-deposited Si_3N_4 by about 25%, which is consistent with measurements from experiments using rapid thermal annealing [W. Jiang,

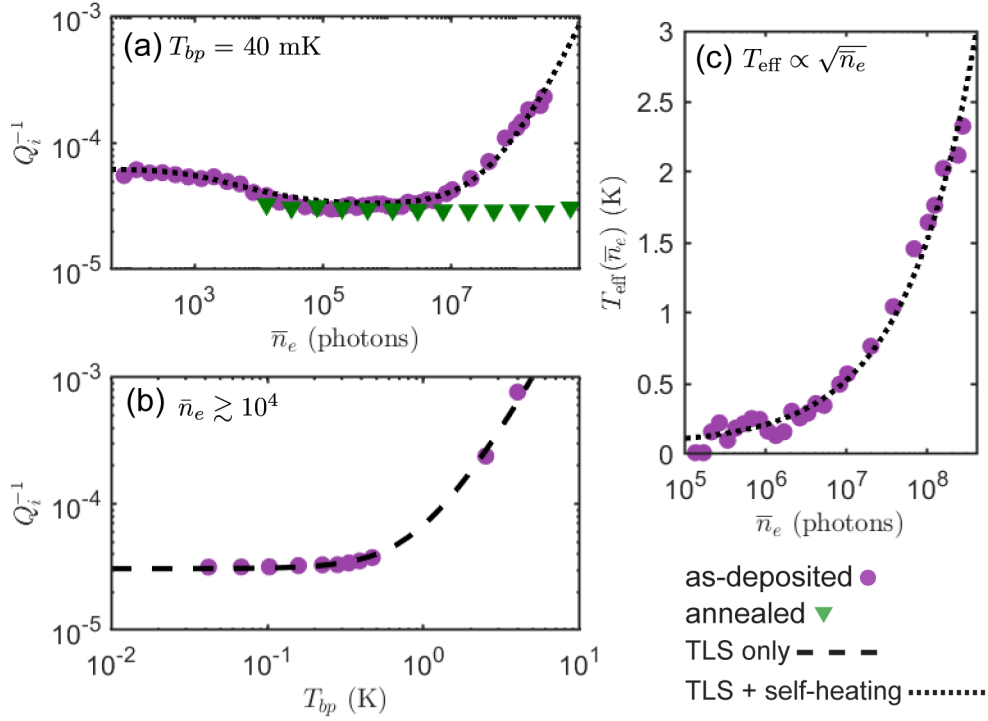


Figure 5.10: **TLS and self-heating in transducer improved by annealing.** (a) Power-dependent loss of a transducer with as-deposited (blue circles) and annealed (green triangles) Si_3N_4 , with a TLS and self-heating fit to the former. The transducer with annealed Si_3N_4 (Aspen 1) no longer has power-dependent loss. The motional capacitor pad is also designed to screen the Si_3N_4 from the electric field and reduce F_{SiN} . (b) Temperature-dependent loss in a transducer with as-deposited Si_3N_4 , along with a fit the TLS model. (c) Self-heating map extracted from the transducer loss data. Similar to the planar designs, the T_{eff} scales as the square root of the intracavity photon number.

D. Xu, et al. 2016]. This will correspond to a 25% decrease in the membrane mode frequencies and a 25% decrease in the maximum achievable Q_m , which both must be accounted for in the design of the membrane. To solve this issue, we could use rapid thermal annealing as opposed to furnace annealing in order to more accurately control the annealing time, which could provide control of Si_3N_4 stress (and perhaps even increase it). A reduced annealing time, however, could possibly grant only a partial improvement of the microwave loss so both microwave loss and mechanical stress must be co-optimized.

5.4 Microscopic origin of TLS

To better understand the impact of annealing on the dielectric loss of Si_3N_4 , we use Fourier-transform infrared spectroscopy (FT-IR) to attempt to identify the microscopic source of the TLS. For our Si_3N_4 film, we expect to primarily find molecular vibrational modes of the Si-N bond and Si-H or N-H impurities [Paik and Osborn 2010; Herth et al. 2010].

To prepare samples for the FT-IR, we deposit a 400 nm thick layer of Si_3N_4 onto both sides of a Si wafer. This film is much thicker than that used for the resonators because a thicker film will increase the absorption signal. We dice the wafer into four samples, and then etch the film on one side of each of the samples to obtain different thicknesses. This results in a set of samples with Si_3N_4 thickness ranging from 800 nm (no Si_3N_4 etched) to 500 nm (300 nm of Si_3N_4 etched on one side of the wafer). A similar procedure is used to prepare the annealed Si_3N_4 samples, except the anneal is performed directly after the deposition.

Our FT-IR spectrometer model is Thermo Scientific Nicolet iS50, operated in transmission mode [Boyle et al. 2022]. The transmission mode of operation is less sensitive than attenuated total reflection techniques, but nonetheless reveals the film impurities given proper sample preparation and data analysis. After purging the sample chamber with nitrogen gas for 30 minutes, we take a background spectrum with a bare Si wafer. For the spectra of each Si_3N_4 film, we subtract the background spectra and then perform an atmospheric correction to remove absorption features from residual water and CO_2 in the sample chamber. A polynomial fit is used to remove the baseline

structure [Mazet et al. 2005]. The areas under the absorption peaks are extracted from Gaussian fits to the spectra.

The spectra in Fig. 5.11 reveal that the as-deposited Si_3N_4 is host to hydrogen impurities, whose atomic concentration is substantially reduced after annealing. In addition to the Si-N stretching mode at 890 cm^{-1} (not shown), we observe peaks in the absorption spectrum at 3330 cm^{-1} and 2210 cm^{-1} for the as-deposited films corresponding to the stretching modes of the N-H and Si-H bond, respectively. The height of these peaks decreases with decreasing film thickness. Notably, the peaks of the spectrum for the annealed films are nearly indistinguishable from the background.

Using well-established calibrations of the infrared absorption cross-sections of these bonds, we can calculate the atomic hydrogen percentage in the film using the Lanford and Rand method [Lanford and Rand 1978], described by

$$\text{atomic \%H} = \frac{[\text{Si-H}] + [\text{N-H}]}{[\text{Si}] + [\text{N}] + [\text{Si-H}] + [\text{N-H}]}, \quad (5.42)$$

where $[\text{Si}]$, $[\text{N}]$, etc. denote the densities in atoms per unit volume. We have made the assumption that all of the -H in the film are bonded to either N or Si. $[\text{Si-H}]$ can be determined by $[\text{Si-H}] = A_{\text{Si-H}}/(\sigma_{\text{Si-H}} \times t)$, where $A_{\text{Si-H}}$ is the area under the Si-H absorption peak, $\sigma_{\text{Si-H}}$ is the IR absorption cross-section, and t is the thickness of the film ($[\text{N-H}]$ is found in an analogous manner). The sum $[\text{Si}] + [\text{N}]$ is the atomic density of Si_3N_4 in units of atoms per unit volume, which can be determined from the density of the film.

We find that annealing reduces the hydrogen content of the Si_3N_4 by at least an order of magnitude. The hydrogen percentage for the annealed films is an upper-bound, as we cannot resolve any features in the spectrum above the background for the annealed film. In terms of atomic concentration, there are about 2×10^{21} hydrogen atoms per cubic centimeter, or one defect every 1.26 nm assuming an isotropic distribution. This density of impurities could support the conclusion that the TLS bath itself is getting hot (Sec. 5.2.5). Additionally, the hydrogen content for both films is independent of thickness, implying the impurities are uniformly distributed throughout the bulk.




Figure 5.11: **Infrared spectroscopy of Si_3N_4 .** The N-H and Si-H bond stretching vibrational modes are in the IR spectrum between 60 and 100 THz at wavenumbers of 3330 cm^{-1} and 2210 cm^{-1} , respectively. We measure Si_3N_4 films of different thicknesses with Fourier-transform infrared spectroscopy (FT-IR) before and after annealing, noticing a significant reduction in the absorbance at these wavenumbers. (b) The atomic hydrogen percentage in the film is calculated from the spectrum measurements. The annealed film has more than an order of magnitude less hydrogen content than the as-deposited film. Error bars represent $\pm 1\sigma$.

As a final note, we observed atomic hydrogen percentages of 20-30% for unannealed Si_3N_4 deposited with a 7:1 ratio of precursor gases and LPCVD Si_3N_4 obtained from Rogue Valley Microdevices, a foundry from which we sourced Si_3N_4 before the furnace at NIST was operational.

The near-IR modes identified by the FT-IR are likely too high in frequency to be the source of loss in the Si_3N_4 , but they do serve as markers of hydrogen impurities that can host other TLS. Hydrogen impurities in alumina (Al_2O_3) are strong candidates for GHz-frequency TLS [Müller et al. 2019; Holder et al. 2013], where an interstitial hydrogen atom tunnels between adjacent oxygen atoms [Gordon et al. 2014]. Analogous defects in the 20-50 GHz range have been found in niobium films, where interstitial hydrogen atoms tunnel between adjacent nitrogen impurities [Morr et al. 1989; Steinbinder et al. 1991; Cannelli et al. 1998]. Similar defects could exist in our Si_3N_4 films. While the FTIR data presented do not conclusively identify a specific physical TLS realization, they do suggest that the source of the TLS is intimately related to the presence of hydrogen impurities, which we have shown are removed by thermal annealing.

5.5 TLS energy distribution

Interestingly, annealing realizes only a modest improvement to the resonant-type loss, but vastly reduces the relaxation-type loss. As noted in Sec. 5.2.2, the resonant interaction couples the electrical resonator to TLS with energy close to ω_0 , while in Sec. 5.2.3, we saw that the relaxation interaction mainly couples to off-resonant TLS. Similarly, the frequency shift of each interaction samples a different subset of TLS energies than does the loss, which could explain the discrepancy between the self-heating predictions for the frequency shift as that model is extracted from the loss data.

The relaxation contribution starts to dominate over the resonant contribution above $T_{\text{bp}} \sim 500$ mK, suggesting that the impurities removed by annealing were host to TLS with energies predominantly above $500 \text{ mK} \times 2k_B/h \simeq 20$ GHz. This observation motivates looking for defects around that energy and suggests that the annealing changes the frequency-dependent structure of the TLS density of states. The anomalous temperature dependence of the heat capacity of glasses is

expected to be due to a weak energy dependence $P(E) \propto E^\mu$ of the TLS energy density [Behunin et al. 2016]. In the generalized tunneling model, the non-uniform energy distribution is thought to be a result of TLS-TLS interactions, which could be consistent with the heating of the TLS bath [Faoro and Ioffe 2015b]. Specific heat measurements in glasses at cryogenic temperatures from the 1980's support this weak energy dependence, with $\mu = 0.2 - 0.3$ [Hunklinger and Raychaudhuri 1986]. To our knowledge, there has been one attempt to directly measure the energy distribution of defects, in which three lumped element resonators between 1-5 GHz were measured [Skacel et al. 2015]. By extracting the loss tangent at each frequency, they observed a weak energy dependence for amorphous silicon dioxide, consistent with the measurements from glasses.

We designed the following experiment in order to extend the results of [Skacel et al. 2015] and perhaps identify the 20-50 GHz TLS defect which causes relaxation damping in our devices. Instead of distinct lumped element resonators, we design a single transmission line resonator in a coplanar waveguide (CPW) geometry [Göppl et al. 2008]. The free-spectral range of such a structure decreases proportionally to the length of the transmission line ($\text{FSR} \sim 1/L$), so a long transmission line will have a dense spectrum of harmonics at microwave frequencies. By measuring the power- and temperature-dependent properties of the each resonator, we can extract the material parameters and look for structure in the loss tangent versus resonator frequency.

The device design is described in Fig. 5.12. The transmission line is 12.105 cm long, with interdigitated finger capacitors providing the coupling to the microwave readout chain. Similar to the planar, lumped element designs studied earlier in this chapter, the Si_3N_4 is only underneath the metal. Finite-element simulations and conformal mapping techniques can be used to find the effective dielectric constant, free-spectral range, and F_{SiN} [Gevorgian et al. 1995; Murray et al. 2018]. While we expect little dispersion in the CPW below 100 GHz, the coupling capacitors could change the electromagnetic mode shape, and thus F_{SiN} , at different frequencies, which still needs to be understood [Garg et al. 1979].

Measurements have begun on this resonator design, revealing an FSR of 455 MHz which gives 18 modes within the 4-12 GHz band of our measurement chain. In addition to furthering the

type of study performed in [Skacel et al. 2015], we hope to perform some "spectral-hole burning" experiments [Capelle et al. 2020]. By applying a microwave pump to one mode, we can saturate the TLS of that energy and then measure the microwave properties of the other modes. This could help understand which TLS are contributing most strongly to the relaxation interaction, which mainly samples off-resonant TLS. If these measurements prove successful, one could imagine extending this technique to other dielectrics, such as aluminum oxide, silicon dioxide, and silicon or working with an even longer transmission line to increase the mode density.

5.6 TLS description of $\gamma_m(T)$

As a final note on dissipation from the TLS model, we speculate that the temperature-dependent loss of the mechanical membrane described in Sec. 3.1.3 could be due to relaxation damping. For low-frequency resonators such as the mechanical membrane $\omega_m \sim 2\pi \times 1$ MHz, the TLS that would contribute to resonant damping are thermally saturated even at $T_{bp} = 10$ mK. That leaves relaxation damping, which has successfully described the temperature-dependent loss of nanobeams and two-dimensional membranes with significant patterning that could reduce the effective dimensionality [Hauer, Kim, et al. 2018; X. Zhou et al. 2019; Seis et al. 2022]. For a uniform two-dimensional membrane, we might expect the relaxation damping to scale with T^2 , yet we observe a linear dependence. While the effective dimensionality can be reduced if the relevant thermal phonon wavelengths exceed the system dimensions, this is not obviously the case for our $500 \mu\text{m}$ side length square membranes [Behunin et al. 2016]. Further study of the temperature-dependent mechanical intrinsic loss could enlighten ways to further improve γ_m .

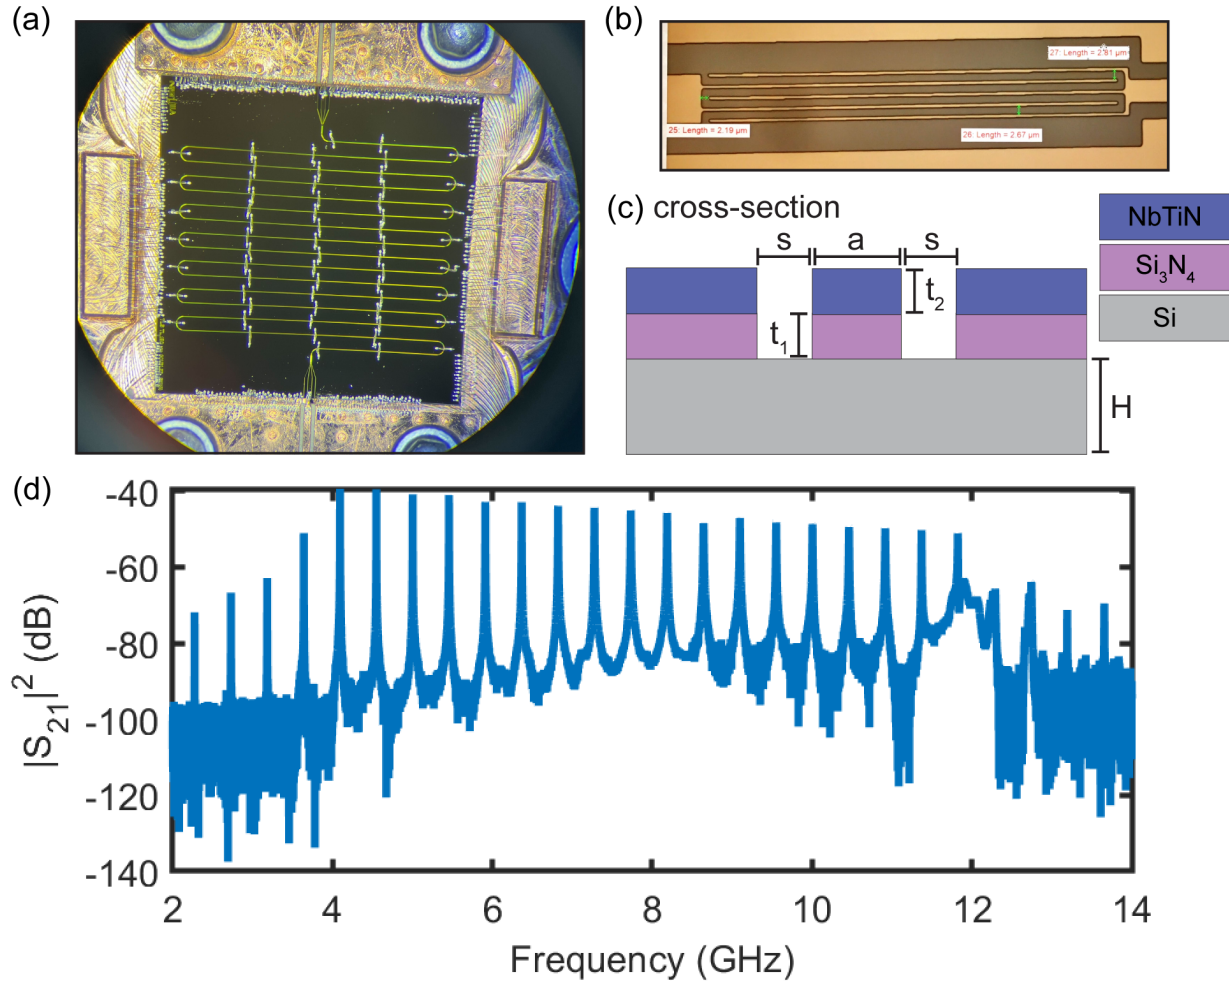


Figure 5.12: **Multimode CPW resonator.** (a) Optical micrograph of the fabricated device, with wirebonds across the ground planes. (b) Optical micrograph of the interdigitated finger capacitors (c) Cross-sectional slice of the material stack, with a center trace width $a = 9 \mu\text{m}$, slot width $s = 5 \mu\text{m}$, dielectric thickness $t_1 = 100 \text{ nm}$, superconductor thickness $t_2 = 140 \text{ nm}$, and substrate thickness $H = 380 \mu\text{m}$. (d) Transmission measurements, showing modes spaced by a 455 MHz between 4 and 12 GHz.

Chapter 6

Characterization of microwave noise

Stumbling across the truth isn't the same as making things up.
- Octavia Butler

We have operated under the assumption that improving the excess power-dependent loss in the transducer would also reduce the power-dependent noise observed in the transducers. Indeed, fluctuation-dissipation theorem argues that the dissipation experienced by a system under external drive is directly related to the fluctuations that system experiences from the environment [Kubo 1966]. We pursued this route because measuring the microwave loss is a much faster and simpler experiment when compared to measuring microwave noise. The latter generally requires careful calibration of the measurement chain. Having remedied the power-dependent loss, however, we now turn to directly measuring the microwave noise, and in particular, the impact of annealing on the noise from the Si_3N_4 dielectric.

Before diving into characterization of the noise, it is useful to compare the total microwave circuit occupation observed in our transducer with other electromechanical experiments. In Fig. 6.1, we plot the occupation of the microwave circuit as a function of the electromechanical pump strength for several experiments with either aluminum-based mechanical resonators (square) or Si_3N_4 nanobeams or membranes (circles). It is immediately clear that the total circuit occupation in our transducer with as-deposited Si_3N_4 is on par or better with that measured in other systems. The fact that many of these other systems achieve ground state cooling (except for [Suh 2011]) is due to either larger Q_m which reduces the coupling to the environment, larger ω_m which reduces the




Figure 6.1: **Total circuit occupation in electromechanical experiments.** Comparison of the microwave cavity occupancy as a function of pump strength for several electromechanical experiments. Squares indicate experiments with aluminum mechanical resonators [Teufel et al. 2011; Massel et al. 2012; Wollman 2015; Youssefi et al. 2023] while circles label experiments with Si_3N_4 nanobeams with Al [Suh 2011; Fink et al. 2016], Si_3N_4 membranes with Al [Liu et al. 2022], Si_3N_4 membranes with Al/NbTiN [Seis et al. 2022]. or Si_3N_4 membranes with NbTiN [Brubaker et al. 2022].

occupancy of the thermal bath, or larger g_e which limits the effect of internal cavity noise from the strong pump. In our system, we use a larger mass and lower frequency oscillator than many of these other experiments, specifically to achieve appreciable parametric coupling to both the electrical and optical resonators.

It is also curious that the total circuit occupation in many of these systems seem to follow a similar trend, despite different mechanical and microwave resonator frequencies and material platforms. This could suggest some underlying noise source that is common to all the systems, despite the different geometries and materials. One constant is the use of a Si substrate, and TLS in the native SiO_2 are present in all of the experiments. Except for [Suh 2011] and [Brubaker et al. 2022], these experiments have minimal impact from phase noise of the microwave generator, either due to filtering or a large mechanical frequency, so we do not believe this combined trend is due to some limit of the generator technology. Regardless of the source of this pump-power dependent noise, this observation further motivates our work to increase the electromechanical coupling rate (Sec. 3.3) so that the experiment can be operated with fewer pump photons, and thus less noise.

We begin by describing details of the noise measurement in Sec. 6.1. In order to accurately determine the total circuit occupation and pump-induced internal cavity noise, we must understand the impact of the technical noise of the microwave generator and how it can be modified by the measurement chain. These insights are used in Sec. 6.2 to calibrate the internal cavity noise in the same designs with either annealed or as-deposited Si_3N_4 that were used to understand the microwave loss. We find that annealing does reduce the noise of the microwave cavity, both in the planar devices and the transducer. The self-heating effect introduced in the last chapter provides a potential model for the internal cavity noise, which we explore in Sec. 6.3 in addition to the noise predicted by the TLS model. Lastly, we discuss two projects in Sec. 6.4 that are a work in progress: a two-mode circuit design that isolates signals from the microwave pump and a near-quantum limited amplifier supplied by NIST to improve the microwave measurement efficiency.

6.1 Noise calibrations

As introduced in [Sec. 2.4](#), the total microwave circuit occupation has two sources: technical noise (predominantly phase noise) from the microwave generator n_ϕ and internally generated noise n_c (or internal cavity noise). Both effects generate real photons in the circuit which can drive the mechanical mode, so extracting the total circuit occupation from electromechanical measurements measures $n_\phi + n_c$. One can also directly measure the thermal spectrum of the microwave circuit. The generator phase noise leaves the microwave circuit with phase given by resonator susceptibility so there is interference between the promptly reflected noise and that which enters the circuit ([Eqn. 2.34](#)). This is further complicated by use of a interferometric cancellation tone to prevent saturation of the measurement chain (see [Fig. 4.2](#)). This tone carries the same phase noise as the electromechanical pump, leading to additional interference. In this section, we will calculate these effects such that we can accurately determine n_c .

6.1.1 Generator phase noise

Phase noise on the electromechanical pump arises from fluctuations of the generator frequency. Generally, it scales proportionally with the carrier frequency and inversely with the detuning from the carrier so we could reduce the influence of phase noise by working at lower microwave frequencies or large mechanical frequencies. As we are only concerned with the phase noise in narrow bandwidth $\Gamma_T/2\pi \sim 10$ kHz, we can achieve a more significant improvement filtering the noise in that small bandwidth. The Q of this filter must be at least $\omega_e/\omega_m > 6000$, achievable at room temperature using oxygen-free high-conductivity copper (OHFC) filter cavities [[Joshi et al. 2021](#)]. The noise measurements of the transducer described in [[Brubaker et al. 2022](#)] and the capacitive Si_3N_4 resonators did not include such a filter, but we later borrowed a filter from John Teufel at NIST that provides >30 dB of suppression in a 50 kHz bandwidth with only 3 dB of insertion loss at the carrier frequency ([Fig. 6.2](#)). This was used for noise measurements of a transducer with annealed Si_3N_4 , described in [Sec. 6.2](#).




Figure 6.2: **Microwave filter cavity.** Filter response, with the cavity tuned to cancel the $\omega_m/2\pi = 1.5$ MHz component of the generator technical noise. More than 35 dB of suppression is achieved with only 3 dB of insertion loss at the pump frequency (red dashed line). The cavity is made of OHFC copper. The total length of the cavity and the insertion depth of the sapphire rod provide two tuning knobs to optimize noise suppression. The cavity is loaned to us from John Teufel.

6.1.2 Coherent pump cancellation

Due to the relatively small electromechanical coupling rate ($g_e \ll \omega_m, \kappa_e$), large electromechanical pump tones are required to enhance the interaction. In our experiment, this often results in about -50 dBm of power incident on the transducer. This could lead to saturation of the microwave measurement chain, particularly of the HEMT amplifier, which would throw off the measurement efficiency calibrations described in Fig. 2.5. We circumvent this problem by interferometrically canceling the strong microwave pump at the input of the HEMT. The microwave components and setup for this pump cancellation for two resonator configurations is depicted in Fig. 6.3.

The transducers are measured in reflection in order to maximize the bidirectional transduction efficiency. This configuration, however, requires a cryogenic microwave switch in order to measure multiple devices. The capacitive Si_3N_4 designs are measured in a "hanger" configuration, where multiple resonators are simultaneously coupled to a single transmission line.¹ This configuration naturally accommodates frequency-multiplexing, but signals generated in the circuit are split between leftward and rightward propagating waves in the transmission line.

¹Akin to clothes hangers on a closet rod, where each hanger is a different resonator and the transmission line is the shared rod.

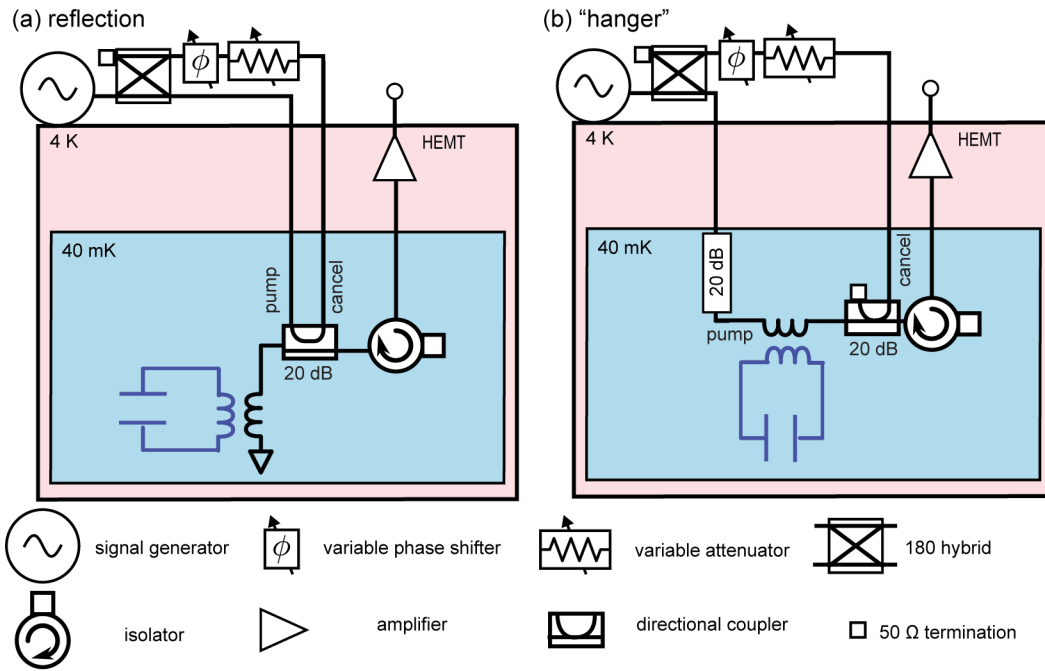


Figure 6.3: **Pump cancellation schematic.** Microwave measurement chain to interferometrically cancel the electromechanical pump at the input of the HEMT with the resonator measured in (a) reflection or (b) "hanger" style. Attenuation on the pump and cancel lines at the higher temperature plates are equivalent but not shown. Also not shown is an attenuator on the pump line at room temperature to match the insertion loss of the variable phase shifter and variable attenuator. List of symbols used in the schematic (bottom).

In both cases, the cancellation tone is tuned to cancel the power of the pump tone at the output of the directional coupler. The technical noise of the pump will reflect off of the resonator and acquire a phase shift relative to the technical noise of the cancellation tone, leading to either constructive or destructive interference. Critically, this effect does not change the noise sent to the resonator or the occupation of the cavity seen by the mechanical mode. It only changes the noise we measure with our microwave measurements.

To calculate the scale of this correction, we find the sum of the phasor of the noise on the pump and cancel tones at the input to the HEMT. The susceptibility of the resonator in the reflection and hanger configurations is:

$$S_{11}(\omega) = \frac{(\kappa_{\text{int,e}} - \kappa_{\text{ext,e}}) - 2i(\omega - \omega_e)}{(\kappa_{\text{int,e}} + \kappa_{\text{ext,e}}) - 2i(\omega - \omega_e)} \quad (6.1)$$

$$S_{21}(\omega) = \frac{\kappa_{\text{int,e}} - 2i(\omega - \omega_e)}{(\kappa_{\text{int,e}} + \kappa_{\text{ext,e}}) - 2i(\omega - \omega_e)}. \quad (6.2)$$

The transmission in hanger style can be written $S_{21} = \frac{1}{2}(1 + S_{11})$, which highlights the interference between the transmitted power and the power that interacts with the cavity. The technical noise of the generator reflects off of the circuit on resonance. Meanwhile, the cancellation tone has an equivalent magnitude and opposite phase to the coherent pump tone, which reflects off of the circuit at a detuning $\omega_e - \omega_m$. If we assume that there is a negligible phase shift and equivalent attenuation at ω_e and $\omega_e - \omega_m$ between the resonator and input of the HEMT, the technical noise seen by the measurement has the phasor:

$$\zeta_{\text{corr}} = \zeta |S_{11}(\omega_e)| e^{i \arg S_{11}(\omega_e)} + \zeta |S_{11}(\omega_e - \omega_m)| e^{i \arg S_{11}(\omega_e - \omega_m) + \pi} \quad (6.3)$$

where ζ is the magnitude of the phase noise at the input of the circuit. The expression for the hanger-style measurement is equivalent, except with $S_{11}(\omega)$ replaced with $S_{21}(\omega)$.

To gain insight, we plot the value of Eqn. 6.3 in Fig. 6.4 with $\zeta = 1$ assuming $\omega_m/2\pi = 10$ MHz and $\kappa_{\text{ext,e}}/2\pi = 0.5$ MHz for both the reflection (blue) and hanger (red) configurations. For

an over-coupled resonator $\kappa_{\text{ext,e}} > \kappa_{\text{int,e}}$ measured in reflection, the phase noise of the pump tone experiences a π phase shift ($\arg S_{11}(\omega_e) = \pi$), leading to constructive interference and measured phase noise larger than that seen by the resonator. At critical coupling (black dashed line), the on-resonance reflection goes to zero ($|S_{11}(\omega_e)| = 0$), so the phase noise on the pump tone is fully attenuated and only the phase noise of the cancellation tone is measured. When the loss exceeds the coupling, the noise of the pump acquires no phase shift ($\arg S_{11}(\omega_e) = 0$), and there is destructive interference of the noise due to the cancellation.




Figure 6.4: **Measured phase noise correction.** Interference of technical noise of the generator on the pump and cancel tones leads to modification of the output spectrum for resonators measured in reflection (blue) or hanger (red) configurations. The vertical dashed line corresponds to critical coupling $\kappa_{\text{int,e}} = \kappa_{\text{ext,e}}$.

In the hanger configuration, the noise of the pump tone always acquires zero phase shift ($\arg S_{21}(\omega_e) = 0$), so there is destructive interference of the noise for all values of the overcoupling ratio. When the resonator is perfectly overcoupled $\kappa_{\text{int,e}} = 0$, the on-resonance transmission goes to zero ($|S_{21}(\omega_e)| = 0$) as all of the noise is reflected and again, the measurement only sees the noise of the cancellation tone. For our resonators, the pump is not infinitely detuned from the resonator (imperfect sideband resolution), so the pump acquires a non-zero phase. This reduces (increases) the measured phase noise relative to the infinitely detuned case for overcoupled (undercoupled) resonators as the relative phase between the pump and its noise is not exactly π (zero).

Most of the experiments in Fig. 6.1 employed significant filtering of the generator phase noise

using room-temperature microwave filter cavities (Sec. 6.1.1). Phase noise was thus a negligible contribution to the total circuit occupation. In [Brubaker et al. 2022], we did not include any filtering and moreover, the power-dependent loss means the correction factor is also power-dependent. We will see that, overall, this effect is a small correction (<50%) to the measured noise, but could be an important part of identifying the physics of the internal cavity noise.

6.2 Impact of annealing on microwave noise

To measure the total circuit occupation induced by the electromechanical pump, we measure noise power at the output of the microwave circuit in a small bandwidth centered on ω_e as a function of pump power. Due to limitations of our measurement setup (namely bandwidth of the lock-in amplifier and added noise of the HEMT), it is difficult to resolve the Lorentzian lineshape of the thermal noise of the cavity, although this has been previously observed in our transducers [Burns 2019]. Instead, by measuring on-resonance, we extract the height of the Lorentzian, which is proportional to the cavity occupancy when expressed in photon number units above the measurement noise.²

The power-dependent frequency shift of the planar devices with as-deposited Si_3N_4 means we must dynamically tune the frequency of the drive tone to maintain a constant detuning relative to the resonator frequency. It is difficult to set the detuning appropriately for a given drive strength, because the drive detuning sets the intracavity photon number which may in turn shift the resonator frequency. We converge on the correct detuning by iteratively setting the pump detuning relative to an assumed resonance frequency and then fitting the measured resonator response. We repeat this loop until the predicted detuning differs from the measured detuning by less than 0.01%. This process is fairly time-consuming, so we only measure the microwave noise of designs A and E for the resonators with as-deposited Si_3N_4 .

The spectrum then contains three contributions: the measurement noise, the phase noise of the generator after accounting for the interference effects described in the last section, and the cavity

²In voltage spectral density units, the area under the Lorentzian is proportional to the occupancy.

thermal noise (internal cavity noise):

$$N_{\text{meas}} = 1/\eta_{\text{meas}} + n_{\phi} + R \frac{\kappa_{\text{ext,e}}}{\kappa_{\text{e}}} n_{\text{c}}. \quad (6.4)$$

The measurement efficiency is calibrated in one of two ways. If the microwave resonator is coupled to a mechanical mode, we can use the membrane as a calibrated noise source, as described in [Fig. 2.5](#). For the planar resonators with no mechanical mode, we instead inject a probe tone of known power, which can be used to calibrated the noise floor given that the attenuation on the input line is known. The attenuation is measured while the fridge is warm, so we expect at worst a factor of 2 overestimate of the measurement efficiency calibrated this way.

The contribution from generator phase noise is:

$$n_{\phi} = \zeta_{\text{corr}} \dot{N}_{\text{e}} S_{\delta\phi, \delta\phi, \text{e}} / 2 \quad (6.5)$$

where the factor of two turns the one-sided double-sideband spectral density $S_{\delta\phi, \delta\phi, \text{e}}$ into a two-sided double-sideband spectral density. We measure the $S_{\delta\phi, \delta\phi, \text{e}}$ using a Rohde & Schwarz FSWP phase noise analyzer for each pump frequency. ζ_{corr} accounts for the interference that the phase noise experiences from interacting with the cavity and the pump cancellation. The value of the pump photon flux \dot{N}_{e} is also sensitive to changes in the input line attenuation when cooling down. For the total circuit occupation, the pre-factor $R = 4$ for measuring in reflection, and $R = 2$ for the hanger configuration.

The power-dependent internal cavity noise n_{c} for designs A and E with as-deposited Si_3N_4 and designs A-E with annealed Si_3N_4 is shown in [Fig. 6.5](#) after correctly subtracting away the phase noise contribution to the measured noise power. There are only two points of comparison for as-deposited Si_3N_4 , but the data suggest that design A, with the largest F_{SiN} , has more noise than design E, which has the least F_{SiN} . Importantly, annealing improves the noise of both resonators by about an order of magnitude at high powers. Comparing the noise of designs A-E with annealing

Si_3N_4 , there is no clear trend of power-dependent total circuit occupation on F_{SiN} , which implies that the remaining noise is not due to the bulk Si_3N_4 .

Preliminary measurements of a transducer (named Aspen 1) with annealed Si_3N_4 suggest that annealing reduced the power-dependent internal circuit noise in the full opto-electromechanical system. In Fig. 6.6, we compare the internal cavity noise of the transducer used in [Brubaker et al. 2022] to that of a transducer with annealed Si_3N_4 and partial screening of the Si_3N_4 . These are the same devices whose power-dependent loss is reported in Fig. 5.10. While the phase noise corrections described in the last section were needed to extract the internal cavity noise of the transducer with as-deposited Si_3N_4 , this was not required for the transducer with annealed Si_3N_4 as the filter cavity described in Sec. 6.1.1 reduced the generator technical noise by more than 30 dB. In addition to an overall reduction in the internal cavity noise in the transducer circuit, we find that the residual noise after annealing is comparable with the annealed planar devices, which supports that the conclusion that the residual noise is not related to bulk Si_3N_4 participation. The data presented in this section are strong evidence that annealing and screening of the Si_3N_4 have significantly reduced the pump-induced microwave internal cavity noise. In concert with the results of the last chapter and filtering of the generator phase noise, these results suggest substantial improvements to the electromechanical-pump-induced microwave circuit loss and noise which were the main limits on the performance of our doubly-parametric, mechanically-mediated transducer architecture.

6.3 Modeling internal cavity noise

The TLS model predicts that, in addition to dissipation, a microwave resonator coupled to a bath of TLS will also experience fluctuations, in accordance with the fluctuation-dissipation theorem [Behunin et al. 2016; Constantin et al. 2009; Shnirman et al. 2005]. As we saw with the TLS-induced microwave dissipation, the TLS model alone does not predict the power-dependent behavior which we observe. Up to the writing of thesis, we do not have extensive dataset so it is difficult to parameterize the TLS model of noise in the same manner we did for the circuit loss. We instead develop a simple coupled-mode model of the internal cavity noise based on the self-heating extracted from the circuit

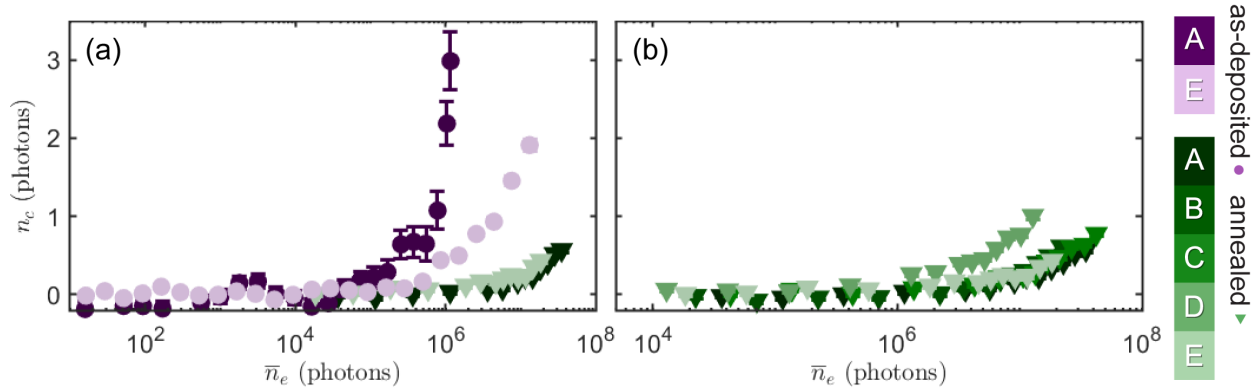


Figure 6.5: **Internal cavity noise in planar resonators.** Power-dependent internal cavity noise of planar, capacitive Si_3N_4 designs. (a) Designs A and E with either as-deposited (blue circles) or annealed (green triangles) Si_3N_4 . Design A (with larger F_{SiN}) has more noise than design E for the as-deposited resonators. Annealing reduces the power-dependent total circuit occupation from internally generated noise. (b) Designs A-E with annealed Si_3N_4 . No clear relation is seen between F_{SiN} and noise occupation, suggesting the remaining noise is not due to the Si_3N_4 . Error bars represent $\pm 1\sigma$ uncertainty.




Figure 6.6: **Internal cavity noise in transducers.** Power-dependent internal cavity noise of transducer devices. The transducers have either as-deposited (purple circles) [Brubaker et al. 2022] or annealed (green triangles) Si_3N_4 (from device Aspen 1), and are the same devices whose power-dependent loss is reported in Fig. 5.10. We see that annealing also reduces internal cavity noise in the transducer devices.

loss data, which predicts the observed occupation of the planar designs with reasonable accuracy but underestimates the observed occupation of the transducer circuit.

First, we review the predictions that the TLS model would have for the temperature-dependent internal cavity noise. When the TLS dipole moment is driven by random thermal noise, the TLS-resonator coupling can generate electromagnetic noise. The power spectrum of the dipole fluctuations is given by Eqn. 19 of [Behunin et al. 2016] for a TLS of energy E with relaxation (dephasing) time T_1 (T_2):

$$S_{ij}(\omega) = \delta_{ij} p_0^2 \times \left[\frac{\Delta_0^2}{E^2} \left(p_e(E) \frac{2T_2}{1 + T_2^2(E/\hbar + \omega)^2} + p_g(E) \frac{2T_2}{1 + T_2^2(E/\hbar - \omega)^2} \right) + \frac{\Delta^2}{E^2} \text{sech}^2(E/2k_B T) \frac{2T_1}{1 + T_1^2 \omega^2} \right] \quad (6.6)$$

where $p_g(E)(p_e(E))$ is the probability of the defect to be in the ground (excited) state, given by Fermi-Dirac statistics. The first two terms (proportional to Δ_0) capture the contribution from resonant damping, while the third term (proportional to Δ) models the contribution from relaxation damping. Noise from resonant damping is only relevant when the TLS and circuit are in resonance $\hbar\omega = E$, and otherwise fluctuations from relaxation damping dominate the fluctuations induced by a single TLS.

Now, we must integrate over all TLS energies, for which we again assume a uniform energy distribution $f(\Delta, \Delta_0) = P(E)/\Delta_0$, although more detailed models of the energy distribution can lead to different temperature dependencies of the fluctuations from relaxation damping [Shnirman et al. 2005]. The resonant term is negligible at low-frequencies and in the high frequency limit $\omega T_2 \gg 1$ is proportional to the number of defects:

$$S_{\text{res}}(\omega) \propto V_D P_D(\hbar\omega) p_g(\hbar\omega). \quad (6.7)$$

In this limit, the probability of being in the ground state approaches unity for sufficiently low temper-

atures, so the noise from resonant damping acquires a temperature- and dimension-independent value. The high-frequency (low-temperature) limit of the noise from relaxation damping $\omega T_{1,\min}(k_B T) \gg 1$ is:

$$S_{\text{rel}}(\omega) \propto V_D P_D(k_B T) \frac{T^{D+1}}{\omega^2} \quad (6.8)$$

for an idealized D -dimensional bulk. In this limit, relaxation processes result in $1/f^2$ noise, which eventually runs into the resonant damping at sufficiently high frequencies. Given that the relaxation damping suggested a two-dimensional phonon bath, we would expect the noise to scale as T^3 in this regime. In the low-frequency (high-temperature) limit, $\omega T_{1,\min}(k_B T) \ll 1$, the noise has a $1/f$ character and scales linearly with temperature:

$$S_{\text{rel}}(\omega) \propto V_D P_D(k_B T) \frac{T}{\omega}. \quad (6.9)$$

The full expression for noise from relaxation processes can be solved analytically as we did for the relaxation damping. As a final note, it can be shown that the cross-over between the high- and low-frequency regimes occurs at $\hbar\omega \sim k_B T$ [Shnirman et al. 2005].

The only explicitly power-dependent behavior of the noise is the saturation of the resonant damping, which takes $p_g(\hbar\omega)$ from $1 \rightarrow 0.5$, only a factor of two change. Ideally, we could use the self-heating ansatz to create power-dependent noise from the relaxation term as we did for the pump-induced circuit loss (Sec. 5.2.5). Unlike the dissipation, however, we have not taken careful data of the temperature-dependent noise in circuits with either as-deposited or annealed Si_3N_4 at the time of writing this thesis. Without the ability to constrain the parameters of the noise model from relaxation damping, we cannot construct a self-heating prediction of the power-dependent noise from the temperature-dependent noise using the predictions of the standard TLS model. We should still expect annealing to reduce the noise from TLS, as we found that the defect density was drastically reduced in Sec. 5.11.

A more general model, developed by [Faoro and Ioffe 2015b], which accounts for a weakly energy dependent TLS distribution $P(E) \propto E^\mu$ and TLS-TLS interactions predicts that at high-field

intensities, the fractional frequency noise from TLS coupled to a two-dimensional phonon bath should scale as:

$$\frac{S_{\delta f}(\omega)}{f^2} \propto \frac{1}{\omega \sqrt{n_e} T_1} T^{-(1+\mu)/3} \quad (6.10)$$

which also only predicts reduction of the fractional frequency noise with larger drive powers. This model could more accurately reflect our system, given that we expect strong TLS-TLS interactions and a non-uniform TLS energy distribution, but needs to be compared to temperature-dependent noise data. On an empirical front, [Gao 2008] observed a

$$\frac{S_{\delta f}(\omega)}{f^2} \propto \frac{T^{-3/2}}{\sqrt{n_e}} \quad (6.11)$$

scaling for the fractional frequency noise of superconducting aluminum coplanar waveguide resonators between $0.3 \text{ K} < T < 0.9 \text{ K}$. The noise increased by an order of magnitude when plasma-enhanced chemical vapor deposition (PECVD) SiN_x was deposited onto the resonators.

In the absence of the appropriate temperature-dependent noise data, we developed a phenomenological model of the internal cavity noise based on the self-heating ansatz employed to model the power-dependent dissipation. This model is inspired by the coupled mode picture of the mechanical mode occupation and gives approximately the right scale of the power-dependent total circuit occupation. The occupancy of the microwave circuit can be written as a weighted average of the occupancies of the bath to which it is coupled. The coupling rates are given by the circuit's loss rates to these various baths. For this model, we will consider three baths: the transmission line with coupling Q_e^{-1} , the background loss with coupling Q_{bg}^{-1} , and the TLS bath in the Si_3N_4 with coupling $Q_{\text{res}}^{-1} + Q_{\text{rel}}^{-1} = Q_{\text{TLS}}^{-1}$. The occupancy of each bath is given by the Planck factor. We assume that the transmission line n_{TL} and background n_{bg} occupancy remain cold, while the Si_3N_4 heats up and the occupancy n_{SiN} is given by T_{eff} . The occupancy of the microwave circuit is then:

$$n_c = \frac{n_{\text{TL}} Q_e^{-1} + n_{\text{bg}} Q_{\text{bg}}^{-1} + n_{\text{SiN}} Q_{\text{TLS}}^{-1}}{Q^{-1}} \quad (6.12)$$

which is reminiscent of the Eqn. 2.23.

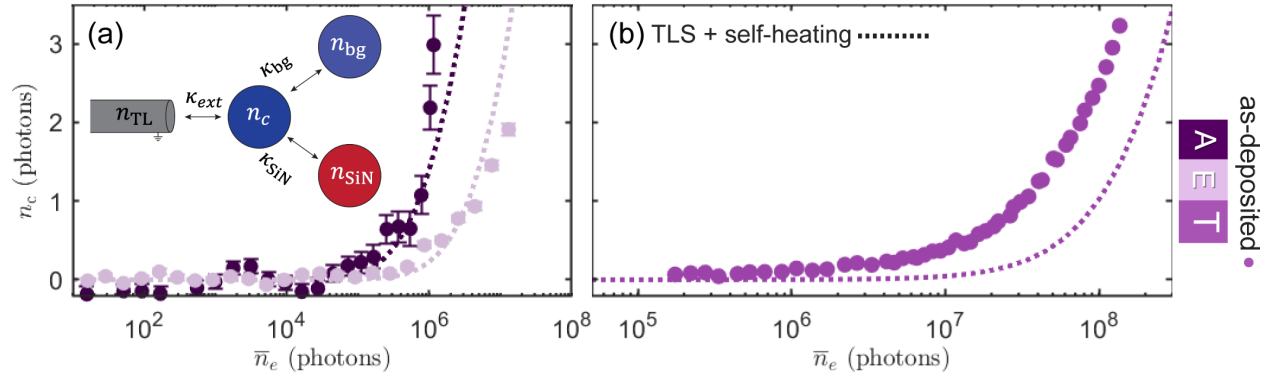


Figure 6.7: **Self-heating based noise model.** (inset) Coupled mode schematic of internal circuit noise. The transmission line and background loss source are assumed to stay cold as a function of drive power. (a) Power-dependent noise of planar designs A and E as-deposited Si₃N₄, with the self-heating model prediction (dotted line). (b) Power-dependent noise of transducer microwave circuit with as-deposited Si₃N₄. Our model underestimates the internal circuit noise. in the transducer.

A schematic description of the model, along with its predictions for designs with as-deposited Si₃N₄ is detailed in Fig. 6.7. While we find reasonable agreement of our model with the power-dependent internal circuit noise. of the planar designs A and E, our model underestimates that internal circuit noise. of the transducer. One possible explanation for this discrepancy is that the temperature-dependent loss of the transducer was not measured at consistent photon numbers. As a result, the parameters of the relaxation damping we extract and the self-heating map could be skewed. In fact, if the self-heating non linearity was 1.5 times larger, there would be much better agreement between our phenomenological model and the measured noise. It also possible that the transducer geometry is susceptible to an additional noise source that the planar designs are insensitive too. The metal-air interface is certainly host to some oxide, which could host its own TLS defects. The flip chip transducer strongly samples these TLS, as the electric fields mainly live between the capacitor pads (see Table 5.1).

In the absence of a complete model of the internal circuit noise. which agrees with all of the observed data, we can compare the scale of noise in the various devices by fitting a linear relationship between n_c and \bar{n}_e , which provides reasonable fits for each design. The slope of these fits is compared

in Fig. 6.8. This comparison emphasizes the conclusions we drew from the Fig. 6.5 and Fig. 6.6. Noise of the planar designs A and E is reduced by at least an order of magnitude, and the annealed designs do not demonstrate any relation between F_{SiN} and the power-dependent internal cavity noise. Most importantly, a combination of annealing and partial screening reduces the internal cavity noise in transducer devices by a factor of three. In combination with filtering of the generator technical noise which accounted for about a quarter of $n_{\text{eff,e}}$ in [Brubaker et al. 2022], the total circuit occupation seen by the mechanical mode is reduced by a factor of four. We expect that a transducer with similar mechanical and optical properties to [Brubaker et al. 2022] but without power-dependent microwave loss and four times less power-dependent microwave noise will be capable of quantum transduction, paving the way for the demonstration of quantum communication protocols between the microwave and optical domain.




Figure 6.8: **Annealing reduces pump-induced internal cavity noise.** Slope of linear fits to n_c vs. \bar{n}_e . Annealing improves the internal cavity noise in the planar designs A and E by at least an order of magnitude. Internal noise of the transducer circuit is improved by about a factor of three by annealing and partial screening, based on linear fits to the data in Fig. 6.6.

6.4 Upgrading microwave architecture

In this section, I will briefly introduce two topics which are worthy of future exploration. The first is a redesign of the transducer microwave circuit to isolate the electromechanical pump from the signal using symmetric and anti symmetric microwave modes. The second is the implementation of a near quantum-limited kinetic inductance traveling wave parametric amplifier (KIT) into the microwave readout chain to reduce the measurement noise/increase the measurement efficiency of microwave readout.

6.4.1 Pump and signal mode isolation

In our current circuit design for the transducer, the electromechanical pump and any potential signal live in the same electromagnetic mode. As a result, the pump-induced heating and generator technical noise at the frequency of the signal will add noise to the signal. We propose a two-mode circuit such that the pump and signal are isolated by the symmetry of the circuit. This also requires a capacitor pad design on the membrane chip that has equal but opposite displacement at two




Figure 6.9: **Two-mode circuit design** (a) Current circuit design. The pump and signal live in the same electromagnetic mode. (b) Two-mode circuit, where the signal lives in the differential (anti-symmetric) mode and the pump interacts with the common (symmetric) mode. We choose a capacitor pad design such C_1 and C_2 experience equal but opposite changes in gap spacing. (c) Adding an inductor and capacitor to the common mode allows independent control of the common mode frequency and external coupling. Purple lines indicate metal on the membrane chip. Insets show outline of proposed capacitor pad design on membrane.

different locations, which can be accomplished by using two out-of-phase anti-nodes of the $(2, 2)$ mode.

This design is described in Fig. 6.9, with insets showing proposed capacitor pad designs on a square membrane. To see how the design in (b) provides isolation between the pump and signal modes, we can write the Hamiltonian of this circuit:

$$\mathcal{H} = \frac{Q_{in}^2}{2C_{in}} + \frac{Q_1^2}{2C_1} + \frac{Q_2^2}{2C_2} + \frac{\phi^2}{2L}. \quad (6.13)$$

We can define the charge of the common and differential modes, $Q_c = Q_1 + Q_2 = Q_{in}$ and $Q_d = Q_1 - Q_2$. Additionally, the two mechanically-mediated capacitors will see equal but opposite changes in the the gap spacing x : $C_1 = C_0/(1 + x/d)$ and $C_2 = C_0/(1 - x/d)$, where C_0 is the unperturbed capacitance and d is the unperturbed gap spacing. To all orders of the displacement, the Hamiltonian becomes:

$$\mathcal{H} = \frac{\phi^2}{2L} + \frac{1}{2} \left(\frac{2C_0 + C_{in}}{2C_0 C_{in}} \right) Q_c^2 + \frac{Q_d^2}{4C_0} + \frac{x}{2C_0 d} Q_c Q_d. \quad (6.14)$$

The last term gives pure three-wave mixing between the common, differential, and mechanical modes.

We choose the common mode as the pump mode and the differential mode as the signal mode. In this configuration, we can apply a strong electromechanical pump on the pump mode, and none of the pump-induced internal cavity noise or generator technical would directly leak into the signal. These noise processes no longer contribute to the mechanical occupancy. This design also improves the promptly reflected noise term for input-referred added noise in downconversion (Eqn. 2.36), which could be important for communication protocols that require downconversion of optical signals.

Now, we analyze the complications of this design. The lack of direct interaction between the pump and signal modes is dependent on the circuit symmetry. Thus, any asymmetry in C_1 and C_2 from misalignment of the capacitor pads during transducer construction could reduce the isolation. If we define $C_1 = C_0/(1 + x/d)$ and $C_2 = C'_0/(1 - x/d)$, the circuit Hamiltonian becomes:

$$\mathcal{H} = \frac{\phi^2}{2L} + \frac{Q_c^2}{2C_{in}} + \frac{1}{8} \frac{C_0 + C'_0}{C_0 C'_0} \left[(Q_c^2 + Q_d^2) + 2 \frac{x}{d} Q_c Q_d \right] + \frac{1}{8} \frac{C_0 - C'_0}{C_0 C'_0} \left[\frac{x}{d} (Q_c^2 + Q_d^2) + Q_c Q_d \right]. \quad (6.15)$$

The third term contains the desired three-wave mixing interaction, while the last term describes a second-order interaction between the mechanical mode and each circuit mode, as well as direct coupling between the pump and signal modes $Q_c Q_d$. The ratio of the last two terms, $\frac{C_0 + C'_0}{C_0 - C'_0}$, describes the ratio of the desired and undesired couplings. If we want a 30 dB ratio of these terms, the capacitances C_0 and C'_0 must be matched to within a part in a thousand. This corresponds to less than $1 \mu\text{m}$ misalignment of the capacitor pads during transducer construction. While we can reliably achieve less than $10 \mu\text{m}$ misalignment with our current assembly process, computer-aided assembly tools and making the membrane chip capacitor pad larger than the circuit chip capacitor pad could increase the alignment accuracy.

In the two-mode circuit design, the coupling capacitor C_{in} sets both the external coupling rate and frequency of the pump mode. Ideally, the pump mode is detuned from the signal mode by the mechanical frequency. It must also have a linewidth $\kappa \ll 4\omega_m$ to avoid backaction effects. Both of these conditions cannot be achieved simultaneously with a single tuning knob C_{in} . Fig. 6.9c describes a circuit design with an additional inductor L_s and capacitor C_s that only affects

the common mode frequency. Further design work to understand the feasibility of constructing such a circuit with the appropriate lumped element components is required, but this is a promising technique for isolating the pump and signal modes, beyond that realized by the coherent pump cancellation.

6.4.2 Incorporating near-quantum limited amplifier

The noise floor of our microwave measurement chain is set by the input-referred of the HEMT amplifier, which is usually 10-20 photons. Thus, resolving the thermomechanical spectrum of the mechanical mode when it is deeply ground state cooled can become time-consuming. This is also true for any weak signal down converted from the optical domain. There are many options for quantum-limited or near-quantum limited amplifiers [Castellanos-Beltran and Lehnert 2007; Roy and Devoret 2016; Frattini et al. 2018] whose input referred added noise can approach the quantum-limit of 0.5 photons. The unpredictable microwave frequency of our transducers as well as the strong electromechanical pump required poses significant constraints on the amplifier. It must either be broadband or tunable, and have a 1 dB compression point $P_{1\text{dB}}$ sufficiently large to handle the electromechanical pump.

In recent years, travelling-wave parametric amplifiers (TWPA) have demonstrated wideband operation over a few GHz, with near-quantum limited noise performance. The two most common implementations use either the non-linearity of the Josephson junction (JTWPA) [Macklin et al. 2015] or the kinetic inductance non-linearity (KITWPA or KIT for short) [Malnou et al. 2022] to generate gain. JTWPAs saturate at powers around -100 dBm as they rely on tunnel junctions, but KITs have generally demonstrated much better power-handling, saturating at -75 to -65 dBm input power. Our transducers have -50 dBm incident pump power, which is then coherently canceled by at least 40 dB, placing us comfortably below the saturation of KITs.³

With the help of Logan Howe at NIST Boulder, we acquired a KIT and verified the gain.

³In [Youssefi et al. 2023], a JTWPA was employed and in [Teufel et al. 2011], a Josephson parametric amplifier (JPA) was used to increase the system measurement efficiency. There is also the possibility of using the two-mode squeezing interaction to use the mechanics as its own amplifier, as in [R. D. Delaney, Reed, et al. 2019],

Unlike the JTWPA, the KIT requires a DC bias to create the kinetic inductance nonlinearity. Then, a three-wave mixing nonlinearity with one pump photon at roughly two times the signal frequency generates gain. The KIT requires a strong microwave pump to provide this gain, so it will be important to manage the reflections and dissipation of this pump tone. The wiring diagram for the KIT is shown in [Fig. 6.10](#), along with a gain measurement. Diplexers are used to direct the KIT pump away from the HEMT, as the KIT pump could also saturate the HEMT gain. We dissipate the KIT pump power at the 4 K plate of the fridge, which has much larger cooling power than the base plate. Circulators before and after the KIT reduce any reflections due to mismatched impedances within the gain bandwidth.

With the KIT pump at 12.599 GHz and 0.3 mA of DC bias current, we measure 10 dB of gain over nearly 2 GHz. Assuming the KIT input-referred added noise is about 3 photons [[Malnou et al. 2022](#)], this would increase the measurement efficiency from $\eta_{\text{meas}} \approx 5\%$ to $\eta_{\text{meas}} \approx 30\%$. This signal-to-noise improvement would facilitate more precise measurements of electromechanical cooling, pump-induced internal cavity noise, and a potential superconducting qubit used for communication protocols. In order to realize this improvement, we need to find ways to work around the ~ 8 MHz gain ripple and reduce the likelihood that the KIT is driven normal by fluctuations of the bias conditions.

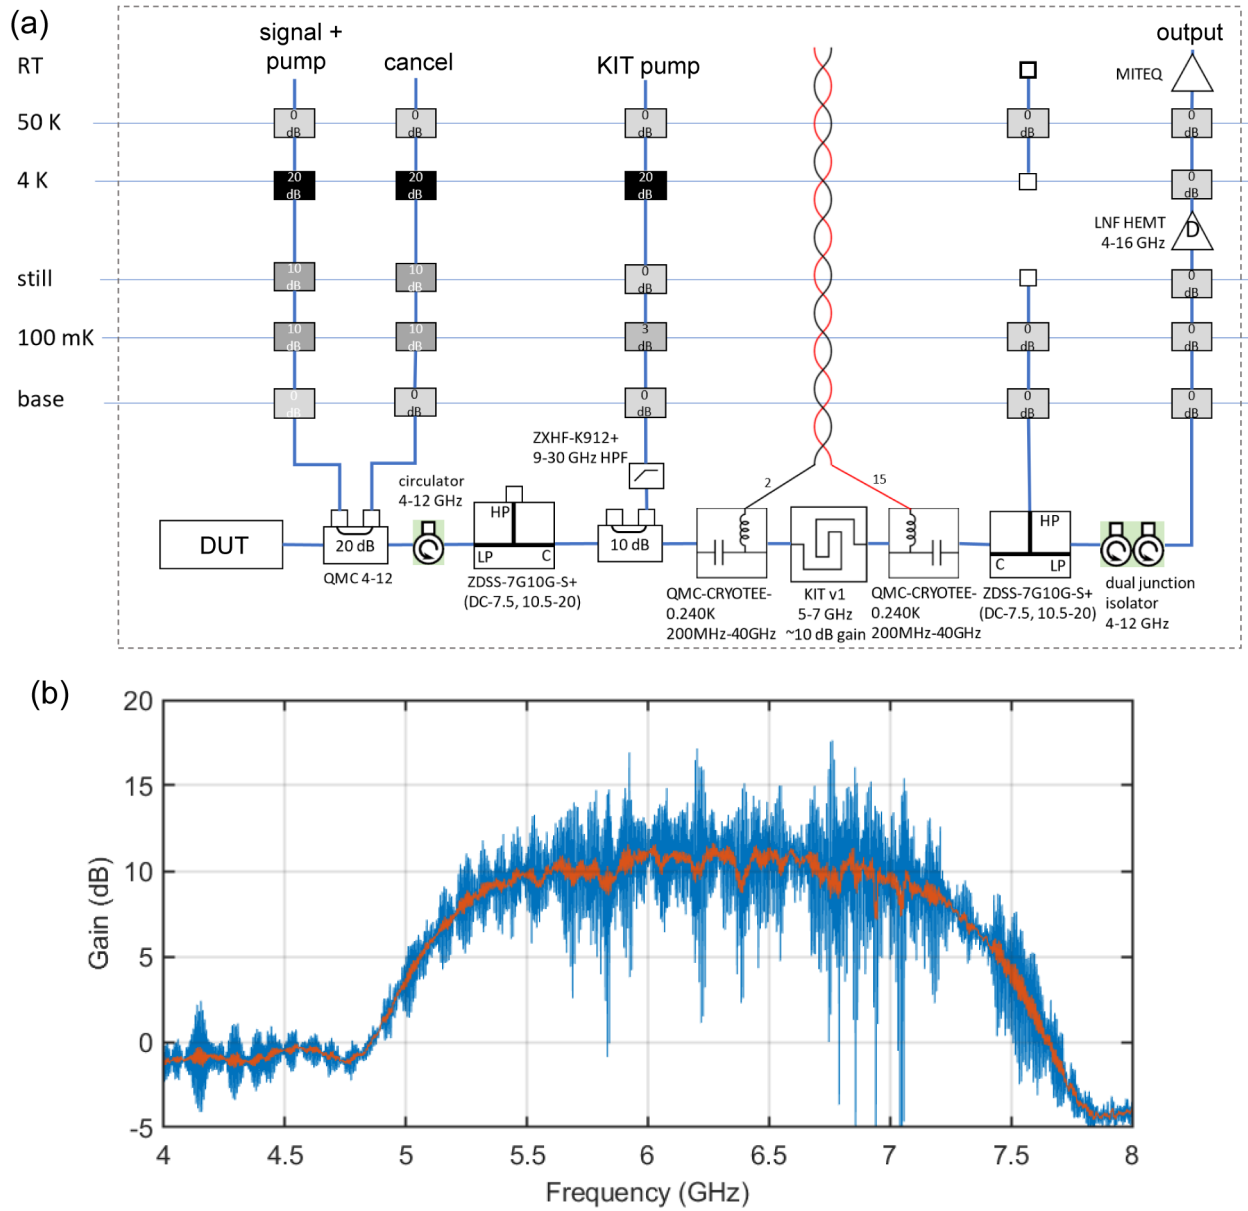


Figure 6.10: **KIT setup and gain** (a) Wiring diagram of KIT setup, including coherent electromechanical pump cancellation and KIT pump lines. The ZDSS-7G10G-S+ are diplexers, components which separate signals at different frequencies. DC current is applied via twisted pair and two bias tees. (b) Gain of KIT, demonstrating 10 dB of gain over nearly 2 GHz. The KIT pump is at 12.599 GHz. Raw data in blue, smoothed in red.

References

- [1] A. G. Bell, “On the production and reproduction of sound by light”, [American Journal of Science](#) **s3-20**, 305 (1880) (Cited on p. 1).
- [2] P. Polishuk, *Fiber Optics Weekly Update* (Information Gatekeepers Inc, 1981) (Cited on p. 1).
- [3] A. Somoroff, Q. Ficheux, R. A. Mencia, H. Xiong, R. Kuzmin, and V. E. Manucharyan, “Millisecond Coherence in a Superconducting Qubit”, [Physical Review Letters](#) **130**, 267001 (2023) (Cited on p. 1).
- [4] S. Ganjam, Y. Wang, Y. Lu, A. Banerjee, C. U. Lei, L. Krayzman, K. Kisslinger, C. Zhou, R. Li, Y. Jia, M. Liu, L. Frunzio, and R. J. Schoelkopf, “Surpassing millisecond coherence in on chip superconducting quantum memories by optimizing materials and circuit design”, [Nature Communications](#) **15**, 3687 (2024) (Cited on pp. 1, 191).
- [5] L. Ding, M. Hays, Y. Sung, B. Kannan, J. An, A. Di Paolo, A. H. Karamlou, T. M. Hazard, K. Azar, D. K. Kim, B. M. Niedzielski, A. Melville, M. E. Schwartz, J. L. Yoder, T. P. Orlando, S. Gustavsson, J. A. Grover, K. Serniak, and W. D. Oliver, “High-Fidelity, Frequency-Flexible Two-Qubit Fluxonium Gates with a Transmon Coupler”, [Physical Review X](#) **13**, 031035 (2023) (Cited on p. 1).
- [6] I. N. Moskalenko, I. A. Simakov, N. N. Abramov, A. A. Grigorev, D. O. Moskalev, A. A. Pishchimova, N. S. Smirnov, E. V. Zikiy, I. A. Rodionov, and I. S. Besedin, “High fidelity two-qubit gates on fluxoniums using a tunable coupler”, [npj Quantum Information](#) **8**, 1 (2022) (Cited on p. 1).
- [7] A. Grimm, N. E. Frattini, S. Puri, S. O. Mundhada, S. Touzard, M. Mirrahimi, S. M. Girvin, S. Shankar, and M. H. Devoret, “Stabilization and operation of a Kerr-cat qubit”, [Nature](#) **584**, 205 (2020) (Cited on p. 1).
- [8] V. V. Sivak, A. Eickbusch, B. Royer, S. Singh, I. Tsioutsios, S. Ganjam, A. Miano, B. L. Brock, A. Z. Ding, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret, “Real-time quantum error correction beyond break-even”, [Nature](#) **616**, 50 (2023) (Cited on p. 1).
- [9] S. Storz, J. Schär, A. Kulikov, P. Magnard, P. Kurpiers, J. Lütolf, T. Walter, A. Copetudo, K. Reuer, A. Akin, J.-C. Besse, M. Gabureac, G. J. Norris, A. Rosario, F. Martin, J. Martinez, W. Amaya, M. W. Mitchell, C. Abellan, J.-D. Bancal, N. Sangouard, B. Royer, A. Blais, and A. Wallraff, “Loophole-free Bell inequality violation with superconducting circuits”, [Nature](#) **617**, 265 (2023) (Cited on p. 1).
- [10] J. F. Dynes, A. Wonfor, W. W.-S. Tam, A. W. Sharpe, R. Takahashi, M. Lucamarini, A. Plews, Z. L. Yuan, A. R. Dixon, J. Cho, Y. Tanizawa, J.-P. Elbers, H. Greißer, I. H. White, R. V. Pentty, and A. J. Shields, “Cambridge quantum network”, [npj Quantum Information](#) **5**, 1 (2019) (Cited on p. 2).

- [11] T.-Y. Chen, X. Jiang, S.-B. Tang, L. Zhou, X. Yuan, H. Zhou, J. Wang, Y. Liu, L.-K. Chen, W.-Y. Liu, H.-F. Zhang, K. Cui, H. Liang, X.-G. Li, Y. Mao, L.-J. Wang, S.-B. Feng, Q. Chen, Q. Zhang, L. Li, N.-L. Liu, C.-Z. Peng, X. Ma, Y. Zhao, and J.-W. Pan, “Implementation of a 46-node quantum metropolitan area network”, [npj Quantum Information](#) **7**, 1 (2021) (Cited on p. 2).
- [12] C. H. Bennett and G. Brassard, “Quantum cryptography: Public key distribution and coin tossing”, [Theoretical Computer Science, Theoretical Aspects of Quantum Cryptography – Celebrating 30 Years of BB84](#) **560**, 7 (1984) (Cited on p. 2).
- [13] D. Gottesman, T. Jennewein, and S. Croke, “Longer-Baseline Telescopes Using Quantum Repeaters”, [Physical Review Letters](#) **109**, 070503 (2012) (Cited on p. 2).
- [14] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, “Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels”, [Physical Review Letters](#) **70**, 1895 (1993) (Cited on p. 2).
- [15] L. Jiang, J. M. Taylor, A. S. Sørensen, and M. D. Lukin, “Distributed quantum computation based on small quantum registers”, [Physical Review A](#) **76**, 062323 (2007) (Cited on p. 2).
- [16] A. Kumar, A. Suleymanzade, M. Stone, L. Taneja, A. Anferov, D. I. Schuster, and J. Simon, “Quantum-enabled millimetre wave to optical transduction using neutral atoms”, [Nature](#) **615**, 614 (2023) (Cited on p. 2).
- [17] J. Verdú, H. Zoubi, C. Koller, J. Majer, H. Ritsch, and J. Schmiedmayer, “Strong Magnetic Coupling of an Ultracold Gas to a Superconducting Waveguide Cavity”, [Physical Review Letters](#) **103**, 043603 (2009) (Cited on p. 2).
- [18] T. Vogt, C. Gross, J. Han, S. B. Pal, M. Lam, M. Kiffner, and W. Li, “Efficient microwave-to-optical conversion using Rydberg atoms”, [Physical Review A](#) **99**, 023832 (2019) (Cited on p. 2).
- [19] B. D. Smith, B. Babaei, A. Narayanan, and L. J. LeBlanc, “Microwave-to-optical conversion in a room-temperature 87Rb vapor for frequency-division multiplexing control”, [Communications Physics](#) **6**, 1 (2023) (Cited on p. 2).
- [20] S. Borówka, U. Pylypenko, M. Mazelanik, and M. Parniak, “Continuous wideband microwave-to-optical converter based on room-temperature Rydberg atoms”, [Nature Photonics](#) **18**, 32 (2024) (Cited on p. 2).
- [21] A. Imamoglu, “Cavity QED Based on Collective Magnetic Dipole Coupling: Spin Ensembles as Hybrid Two-Level Systems”, [Physical Review Letters](#) **102**, 083602 (2009) (Cited on p. 2).
- [22] D. Marcos, M. Wubs, J. M. Taylor, R. Aguado, M. D. Lukin, and A. S. Sørensen, “Coupling Nitrogen-Vacancy Centers in Diamond to Superconducting Flux Qubits”, [Physical Review Letters](#) **105**, 210501 (2010) (Cited on p. 3).
- [23] L. A. Williamson, Y.-H. Chen, and J. J. Longdell, “Magneto-Optic Modulator with Unit Quantum Efficiency”, [Physical Review Letters](#) **113**, 203601 (2014) (Cited on p. 3).
- [24] J. G. Bartholomew, J. Rochman, T. Xie, J. M. Kindem, A. Ruskuc, I. Craiciu, M. Lei, and A. Faraon, “On-chip coherent microwave-to-optical transduction mediated by ytterbium in YVO₄”, [Nature Communications](#) **11**, 3266 (2020) (Cited on p. 3).
- [25] J. Rochman, T. Xie, J. G. Bartholomew, K. C. Schwab, and A. Faraon, “Microwave-to-optical transduction with erbium ions coupled to planar photonic and superconducting resonators”, [Nature Communications](#) **14**, 1153 (2023) (Cited on p. 3).

- [26] R. Hisatomi, A. Osada, Y. Tabuchi, T. Ishikawa, A. Noguchi, R. Yamazaki, K. Usami, and Y. Nakamura, “Bidirectional conversion between microwave and light via ferromagnetic magnons”, [Physical Review B](#) **93**, 174427 (2016) (Cited on p. 3).
- [27] N. Zhu, X. Zhang, X. Han, C.-L. Zou, C. Zhong, C.-H. Wang, L. Jiang, and H. X. Tang, “Waveguide cavity optomagnonics for microwave-to-optics conversion”, [Optica](#) **7**, 1291 (2020) (Cited on p. 3).
- [28] C. Wang, I. Gonin, A. Grassellino, S. Kazakov, A. Romanenko, V. P. Yakovlev, and S. Zorzetti, “High-efficiency microwave-optical quantum transduction based on a cavity electro-optic superconducting system with long coherence time”, [npj Quantum Information](#) **8**, 1 (2022) (Cited on p. 3).
- [29] R. Sahu, W. Hease, A. Rueda, G. Arnold, L. Qiu, and J. M. Fink, “Quantum-enabled operation of a microwave-optical interface”, [Nature Communications](#) **13**, 1276 (2022) (Cited on p. 3).
- [30] L. Qiu, R. Sahu, W. Hease, G. Arnold, and J. M. Fink, “Coherent optical control of a superconducting microwave cavity via electro-optical dynamical back-action”, [Nature Communications](#) **14**, 3784 (2023) (Cited on p. 3).
- [31] J. Holzgrafe, N. Sinclair, D. Zhu, A. Shams-Ansari, M. Colangelo, Y. Hu, M. Zhang, K. K. Berggren, and M. Lončar, “Cavity electro-optics in thin-film lithium niobate for efficient microwave-to-optical transduction”, [Optica](#) **7**, 1714 (2020) (Cited on p. 3).
- [32] M. Xu, C. Li, Y. Xu, and H. X. Tang, “Light-induced microwave noise in superconducting microwave-optical transducers”, [Physical Review Applied](#) **21**, 014022 (2024) (Cited on p. 3).
- [33] H. M. Doeleman, T. Schatteburg, R. Benevides, S. Vollenweider, D. Macri, and Y. Chu, “Brillouin optomechanics in the quantum ground state”, [Physical Review Research](#) **5**, 043140 (2023) (Cited on p. 3).
- [34] W. Jiang, F. M. Mayor, S. Malik, R. Van Laer, T. P. McKenna, R. N. Patel, J. D. Witmer, and A. H. Safavi-Naeini, “Optically heralded microwave photon addition”, [Nature Physics](#) **19**, 1423 (2023) (Cited on p. 3).
- [35] M. J. Weaver, P. Duivesteyn, A. C. Bernasconi, S. Scharmer, M. Lemang, T. C. van Thiel, F. Hijazi, B. Hensen, S. Gröblacher, and R. Stockill, “An integrated microwave-to-optics interface for scalable quantum computing”, [Nature Nanotechnology](#) **19**, 166 (2024) (Cited on p. 3).
- [36] M. Mirhosseini, A. Sipahigil, M. Kalaei, and O. Painter, “Superconducting qubit to optical photon transduction”, [Nature](#) **588**, 599 (2020) (Cited on pp. 3, 70).
- [37] H. Zhao, A. Bozkurt, and M. Mirhosseini, “Electro-optic transduction in silicon via gigahertz-frequency nanomechanics”, [Optica](#) **10**, 790 (2023) (Cited on p. 3).
- [38] S. Meesala, S. Wood, D. Lake, P. Chiappina, C. Zhong, A. D. Beyer, M. D. Shaw, L. Jiang, and O. Painter, “Non-classical microwave–optical photon pair generation with a chip-scale transducer”, [Nature Physics](#) **20**, 871 (2024) (Cited on p. 3).
- [39] C. A. Regal and K. W. Lehnert, “From cavity electromechanics to cavity optomechanics”, [Journal of Physics: Conference Series](#) **264**, 012025 (2011) (Cited on p. 4).
- [40] R. W. Andrews, R. W. Peterson, T. P. Purdy, K. Cicak, R. W. Simmonds, C. A. Regal, and K. W. Lehnert, “Bidirectional and efficient conversion between microwave and optical light”, [Nature Physics](#) **10**, 321 (2014) (Cited on pp. 4, 9, 15, 17, 18, 43, 44).

- [41] E. Planz, Y. Seis, T. Capelle, X. Xi, E. Langman, and A. Schliesser, “Towards quantum electro-optic transduction with an embedded 100 ms coherence time quantum memory”, in [Quantum 2.0 Conference and Exhibition \(2022\)](#), paper QM4B.7 (2022), QM4B.7 (Cited on p. 4).
- [42] J. Suh, “Coupled Dynamics of a Nanomechanical Resonator and Superconducting Quantum Circuits”, PhD thesis (California Institute of Technology, 2011) (Cited on pp. 4, 32, 130–132).
- [43] J. D. Teufel, T. Donner, D. Li, J. W. Harlow, M. S. Allman, K. Cicak, A. J. Sirois, J. D. Whittaker, K. W. Lehnert, and R. W. Simmonds, “Sideband cooling of micromechanical motion to the quantum ground state”, [Nature](#) **475**, 359 (2011) (Cited on pp. 4, 21, 32, 131, 150).
- [44] F. Massel, S. U. Cho, J.-M. Pirkkalainen, P. J. Hakonen, T. T. Heikkilä, and M. A. Sillanpää, “Multimode circuit optomechanics near the quantum limit”, [Nature Communications](#) **3**, 987 (2012) (Cited on pp. 4, 32, 131).
- [45] E. E. Wollman, “Quantum Squeezing of Motion in a Mechanical Resonator”, PhD thesis (California Institute of Technology, 2015) (Cited on pp. 4, 32, 131).
- [46] C. U. Lei, “Circuit Cavity Electromechanics in the Quantum Regime”, PhD thesis (California Institute of Technology, 2017) (Cited on pp. 4, 32).
- [47] A. P. Higginbotham, P. S. Burns, M. D. Urmey, R. W. Peterson, N. S. Kampel, B. M. Brubaker, G. Smith, K. W. Lehnert, and C. A. Regal, “Harnessing electro-optic correlations in an efficient mechanical converter”, [Nature Physics](#) **14**, 1038 (2018) (Cited on pp. 5, 23, 33, 43, 44, 66, 73, 76).
- [48] R. D. Delaney, M. D. Urmey, S. Mittal, B. M. Brubaker, J. M. Kindem, P. S. Burns, C. A. Regal, and K. W. Lehnert, “Superconducting-qubit readout via low-backaction electro-optic transduction”, [Nature](#) **606**, 489 (2022) (Cited on pp. 5, 66).
- [49] B. M. Brubaker, J. M. Kindem, M. D. Urmey, S. Mittal, R. D. Delaney, P. S. Burns, M. R. Vissers, K. W. Lehnert, and C. A. Regal, “Optomechanical Ground-State Cooling in a Continuous and Efficient Electro-Optic Transducer”, [Physical Review X](#) **12**, 021062 (2022) (Cited on pp. 5, 16, 21, 22, 24, 34, 37, 43, 49, 51, 56, 58, 62, 66, 120, 131–133, 138, 140, 141, 146, 170, 192, 194).
- [50] S. Mittal, K. Adachi, N. Frattini, M. Urmey, S.-X. Lin, A. Emser, C. Metzger, L. Talamo, S. Dickson, D. Carlson, S. Papp, C. Regal, and K. Lehnert, “Annealing reduces Si₃N₄ microwave-frequency dielectric loss in superconducting resonators”, [Physical Review Applied](#) **21**, 054044 (2024) (Cited on pp. 5, 94).
- [51] E. Zeuthen, A. Schliesser, A. S. Sørensen, and J. M. Taylor, “Figures of merit for quantum transducers”, [Quantum Science and Technology](#) **5**, 034009 (2020) (Cited on p. 7).
- [52] M. D. Urmey, “Quantum Optomechanics With a Stable Fabry-Pérot Cavity in a Microwave-to-Optical Transducer”, PhD thesis (University of Colorado Boulder, 2024) (Cited on pp. 10, 26, 32, 54, 56, 58, 66, 75).
- [53] C.-H. Wang, F. Li, and L. Jiang, “Quantum capacities of transducers”, [Nature Communications](#) **13**, 6698 (2022) (Cited on p. 10).
- [54] M. Aspelmeyer, T. J. Kippenberg, and F. Marquardt, “Cavity optomechanics”, [Reviews of Modern Physics](#) **86**, 1391 (2014) (Cited on pp. 10, 19).

- [55] C. M. Caves, “Quantum limits on noise in linear amplifiers”, [Physical Review D](#) **26**, 1817 (1982) (Cited on p. 16).
- [56] A. Metelmann and A. A. Clerk, “Quantum-Limited Amplification via Reservoir Engineering”, [Physical Review Letters](#) **112**, 133904 (2014) (Cited on p. 16).
- [57] R. Andrews, “Quantum signal processing with mechanical oscillators”, PhD thesis (University of Colorado Boulder, Boulder, CO, 2015), p. 134 (Cited on p. 16).
- [58] H. P. Yuen and V. W. S. Chan, “Noise in homodyne and heterodyne detection”, [Optics Letters](#) **8**, 177 (1983) (Cited on p. 21).
- [59] M. A. Castellanos-Beltran and K. W. Lehnert, “Widely tunable parametric amplifier based on a superconducting quantum interference device array resonator”, [Applied Physics Letters](#) **91**, 083509 (2007) (Cited on pp. 21, 150).
- [60] A. Roy and M. Devoret, “Introduction to parametric amplification of quantum signals with Josephson circuits”, [Comptes Rendus. Physique](#) **17**, 740 (2016) (Cited on pp. 21, 150).
- [61] A. M. Jayich, J. C. Sankey, K. Børkje, D. Lee, C. Yang, M. Underwood, L. Childress, A. Petrenko, S. M. Girvin, and J. G. E. Harris, “Cryogenic optomechanics with a Si₃N₄ membrane and classical laser noise”, [New Journal of Physics](#) **14**, 115018 (2012) (Cited on pp. 22, 27, 42).
- [62] A. H. Safavi-Naeini, J. Chan, J. T. Hill, S. Gröblacher, H. Miao, Y. Chen, M. Aspelmeyer, and O. Painter, “Laser noise in cavity-optomechanical cooling and thermometry”, [New Journal of Physics](#) **15**, 035007 (2013) (Cited on pp. 22, 30).
- [63] A. J. Weinstein, C. U. Lei, E. E. Wollman, J. Suh, A. Metelmann, A. A. Clerk, and K. C. Schwab, “Observation and Interpretation of Motional Sideband Asymmetry in a Quantum Electromechanical Device”, [Physical Review X](#) **4**, 041003 (2014) (Cited on p. 22).
- [64] T. P. Purdy, K. E. Grutter, K. Srinivasan, and J. M. Taylor, “Quantum correlations from a room-temperature optomechanical cavity”, [Science](#) **356**, 1265 (2017) (Cited on p. 25).
- [65] S. Krinner, S. Storz, P. Kurpiers, P. Magnard, J. Heinsoo, R. Keller, J. Lütolf, C. Eichler, and A. Wallraff, “Engineering cryogenic setups for 100-qubit scale superconducting circuit systems”, [EPJ Quantum Technology](#) **6**, 1 (2019) (Cited on p. 27).
- [66] Y. J. Joshi, N. Sauerwein, A. Youssefi, P. Urich, and T. J. Kippenberg, “Automated wide-ranged finely tunable microwave cavity for narrowband phase noise filtering”, [Review of Scientific Instruments](#) **92**, 034710 (2021) (Cited on pp. 32, 133).
- [67] P. Burns, “Reducing added noise in a microwave-mechanical-optical converter”, PhD thesis (University of Colorado Boulder, Boulder, 2019), p. 139 (Cited on pp. 32, 33, 49, 52, 75, 87, 90, 138).
- [68] J. Gao, “The Physics of Superconducting Microwave Resonators”, PhD thesis (California Institute of Technology, 2008) (Cited on pp. 33, 80, 93, 144).
- [69] H. Paik and K. D. Osborn, “Reducing quantum-regime dielectric loss of silicon nitride for superconducting quantum circuits”, [Applied Physics Letters](#) **96**, 072505 (2010) (Cited on pp. 33, 97, 113, 118, 123, 174).
- [70] J. Burnett, L. Faoro, I. Wisby, V. L. Gurtovoi, A. V. Chernykh, G. M. Mikhailov, V. A. Tulin, R. Shaikhaidarov, V. Antonov, P. J. Meeson, A. Y. Tzalenchuk, and T. Lindström, “Evidence for interacting two-level systems from the 1/f noise of a superconducting resonator”, [Nature Communications](#) **5**, 4119 (2014) (Cited on pp. 33, 105).

- [71] R. O. Behunin, F. Intravaia, and P. T. Rakich, “Dimensional transformation of defect-induced noise, dissipation, and nonlinearity”, [Physical Review B](#) **93**, 224110 (2016) (Cited on pp. 33, 93, 98, 99, 103, 106, 116, 127, 128, 140, 142).
- [72] D. J. Wilson, C. A. Regal, S. B. Papp, and H. J. Kimble, “Cavity Optomechanics with Stoichiometric SiN Films”, [Physical Review Letters](#) **103**, 207204 (2009) (Cited on p. 42).
- [73] I. Wilson-Rae, R. A. Barton, S. S. Verbridge, D. R. Southworth, B. Ilic, H. G. Craighead, and J. M. Parpia, “High- Q Nanomechanics via Destructive Interference of Elastic Waves”, [Physical Review Letters](#) **106**, 047205 (2011) (Cited on pp. 42, 44).
- [74] P.-L. Yu, K. Cicak, N. S. Kampel, Y. Tsaturyan, T. P. Purdy, R. W. Simmonds, and C. A. Regal, “A phononic bandgap shield for high- Q membrane microresonators”, [Applied Physics Letters](#) **104**, 023510 (2014) (Cited on pp. 42, 43).
- [75] M. S. Kushwaha, P. Halevi, L. Dobrzynski, and B. Djafari-Rouhani, “Acoustic band structure of periodic elastic composites”, [Physical Review Letters](#) **71**, 2022 (1993) (Cited on p. 43).
- [76] M. Eichenfield, J. Chan, R. M. Camacho, K. J. Vahala, and O. Painter, “Optomechanical crystals”, [Nature](#) **462**, 78 (2009) (Cited on p. 43).
- [77] D. Hatanaka, I. Mahboob, K. Onomitsu, and H. Yamaguchi, “Phonon waveguides for electromechanical circuits”, [Nature Nanotechnology](#) **9**, 520 (2014) (Cited on p. 43).
- [78] C. Reetz, R. Fischer, G. Assumpção, D. McNally, P. Burns, J. Sankey, and C. Regal, “Analysis of Membrane Phononic Crystals with Wide Band Gaps and Low-Mass Defects”, [Physical Review Applied](#) **12**, 044027 (2019) (Cited on pp. 44, 45, 59).
- [79] I. Wilson-Rae, “Intrinsic dissipation in nanomechanical resonators due to phonon tunneling”, [Physical Review B](#) **77**, 245418 (2008) (Cited on p. 44).
- [80] L. G. Villanueva and S. Schmid, “Evidence of Surface Loss as Ubiquitous Limiting Damping Mechanism in SiN Micro- and Nanomechanical Resonators”, [Physical Review Letters](#) **113**, 227201 (2014) (Cited on pp. 44, 51).
- [81] P.-L. Yu, “Quantum optomechanics with engineered membrane resonators”, PhD thesis (University of Colorado Boulder, Boulder, CO, 2016), p. 119 (Cited on p. 45).
- [82] D. Høj, U. B. Hoff, and U. L. Andersen, *Ultra-coherent nanomechanical resonators based on density phononic crystal engineering*, 2022 (Cited on p. 47).
- [83] G. Huang, A. Beccari, N. J. Engelsen, and T. J. Kippenberg, “Room-temperature quantum optomechanics using an ultralow noise cavity”, [Nature](#) **626**, 512 (2024) (Cited on p. 47).
- [84] B. D. Hauer, J. Maciejko, and J. P. Davis, “Nonlinear power spectral densities for the harmonic oscillator”, [Annals of Physics](#) **361**, 148 (2015) (Cited on p. 48).
- [85] R. Delaney, “Superconducting Qubit Readout via Low-Backaction Electro-Optomechanical Transduction”, PhD thesis (University of Colorado at Boulder, United States – Colorado, 2022) (Cited on pp. 49, 56, 58, 62, 66, 120, 170).
- [86] M. Yuan, M. A. Cohen, and G. A. Steele, “Silicon nitride membrane resonators at millikelvin temperatures with quality factors exceeding 108”, [Applied Physics Letters](#) **107**, 263501 (2015) (Cited on p. 51).
- [87] A. Plöchl and G. Kräuter, “Wafer direct bonding: tailoring adhesion between brittle materials”, [Materials Science and Engineering: R: Reports](#) **25**, 1 (1999) (Cited on p. 52).

- [88] *FiveNine Optics – Boulder, Colorado* / $R > 99.999\%$, <https://www.fivenineoptics.com/> (Cited on pp. 54, 173).
- [89] N. Jin, C. A. McLemore, D. Mason, J. P. Hendrie, Y. Luo, M. L. Kelleher, P. Kharel, F. Quinlan, S. A. Diddams, and P. T. Rakich, “Micro-fabricated mirrors with finesse exceeding one million”, *Optica* **9**, 965 (2022) (Cited on p. 56).
- [90] Hamel, “Georg Duffing, Ingenieur: Erzwungene Schwingungen bei veränderlicher Eigenfrequenz und ihre technische Bedeutung. Sammlung Vieweg. Heft 41/42, Braunschweig 1918. VI+134 S”, *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik* **1**, 72 (1921) (Cited on p. 63).
- [91] “The Duffing Oscillator”, in *Chaos: A Program Collection for the PC*, edited by H. J. Korsch, H.-J. Jodl, and T. Hartmann (Springer, Berlin, Heidelberg, 2008), pp. 157–184 (Cited on p. 63).
- [92] K. Y. Fong, W. H. P. Pernice, and H. X. Tang, “Frequency and phase noise of ultrahigh Q silicon nitride nanomechanical resonators”, *Physical Review B* **85**, 161410 (2012) (Cited on p. 63).
- [93] L. Catalini, M. Rossi, E. C. Langman, and A. Schliesser, “Modeling and Observation of Nonlinear Damping in Dissipation-Diluted Nanomechanical Resonators”, *Physical Review Letters* **126**, 174101 (2021) (Cited on p. 63).
- [94] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf, “Introduction to quantum noise, measurement, and amplification”, *Reviews of Modern Physics* **82**, 1155 (2010) (Cited on p. 70).
- [95] M. Hatridge, S. Shankar, M. Mirrahimi, F. Schackert, K. Geerlings, T. Brecht, K. M. Sliwa, B. Abdo, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret, “Quantum Back-Action of an Individual Variable-Strength Measurement”, *Science* **339**, 178 (2013) (Cited on p. 73).
- [96] 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”, *Physical Review Letters* **95**, 210503 (2005) (Cited on p. 80).
- [97] C. Müller, J. H. Cole, and J. Lisenfeld, “Towards understanding two-level-systems in amorphous solids: insights from quantum circuits”, *Reports on Progress in Physics* **82**, 124501 (2019) (Cited on pp. 80, 97, 126).
- [98] C. R. H. McRae, H. Wang, J. Gao, M. R. Vissers, T. Brecht, A. Dunsworth, D. P. Pappas, and J. Mutus, “Materials loss measurements using superconducting microwave resonators”, *Review of Scientific Instruments* **91**, 091101 (2020) (Cited on pp. 80, 86, 93).
- [99] G. Stan, S. B. Field, and J. M. Martinis, “Critical Field for Complete Vortex Expulsion from Narrow Superconducting Strips”, *Physical Review Letters* **92**, 097003 (2004) (Cited on pp. 80, 82).
- [100] C. Song, T. W. Heitmann, M. P. DeFeo, K. Yu, R. McDermott, M. Neeley, J. M. Martinis, and B. L. T. Plourde, “Microwave response of vortices in superconducting thin films of Re and Al”, *Physical Review B* **79**, 174512 (2009) (Cited on p. 80).
- [101] J. M. Kreikebaum, A. Dove, W. Livingston, E. Kim, and I. Siddiqi, “Optimization of infrared and magnetic shielding of superconducting TiN and Al coplanar microwave resonators”, *Superconductor Science and Technology* **29**, 104002 (2016) (Cited on p. 82).

- [102] I. Nsanzineza, “Vortices and Quasiparticles in Superconducting Microwave Resonators”, Dissertations - ALL (2016) (Cited on p. 82).
- [103] S. H. Lohaus, M. Heine, P. Guzman, C. M. Bernal-Choban, C. N. Saunders, G. Shen, O. Hellman, D. Broido, and B. Fultz, “A thermodynamic explanation of the Invar effect”, [Nature Physics](#) **19**, 1642 (2023) (Cited on p. 82).
- [104] C. Song, M. P. DeFeo, K. Yu, and B. L. T. Plourde, “Reducing microwave loss in superconducting resonators due to trapped vortices”, [Applied Physics Letters](#) **95**, 232501 (2009) (Cited on p. 83).
- [105] D. Bothner, T. Gaber, M. Kemmler, D. Koelle, and R. Kleiner, “Improving the performance of superconducting microwave resonators in magnetic fields”, [Applied Physics Letters](#) **98**, 102504 (2011) (Cited on p. 83).
- [106] B. Chiaro, A. Megrant, A. Dunsworth, Z. Chen, R. Barends, B. Campbell, Y. Chen, A. Fowler, I. C. Hoi, E. Jeffrey, J. Kelly, J. Mutus, C. Neill, P. J. J. O’Malley, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, and J. M. Martinis, “Dielectric surface loss in superconducting resonators with flux-trapping holes”, [Superconductor Science and Technology](#) **29**, 104006 (2016) (Cited on p. 83).
- [107] G. H. C. J. Wijts, “Magnetic Resonance Force Microscopy at milliKelvin Temperatures”, PhD thesis (Leiden University, 2013) (Cited on p. 83).
- [108] C. Lee, H. S. Jo, C. S. Kang, G. B. Kim, I. Kim, Y. H. Kim, H. J. Lee, and J. H. So, “Vibration Mitigation for a Cryogen-Free Dilution Refrigerator for the AMoRE-Pilot Experiment”, [Journal of Low Temperature Physics](#) **193**, 786 (2018) (Cited on p. 83).
- [109] D. J. Goldie and S. Withington, “Non-equilibrium superconductivity in quantum-sensing superconducting resonators”, [Superconductor Science and Technology](#) **26**, 015004 (2012) (Cited on p. 84).
- [110] P. J. de Visser, D. J. Goldie, P. Diener, S. Withington, J. J. A. Baselmans, and T. M. Klapwijk, “Evidence of a Nonequilibrium Distribution of Quasiparticles in the Microwave Response of a Superconducting Aluminum Resonator”, [Physical Review Letters](#) **112**, 047004 (2014) (Cited on p. 84).
- [111] A. Giachero, M. R. Vissers, J. D. Wheeler, M. Malnou, J. E. Austermann, J. Hubmayr, A. Nucciotti, J. N. Ullom, and J. Gao, “Characterization of NbTiN Films With Thicknesses Below 20 nm for Low Power Kinetic Inductance Amplifiers”, [IEEE Transactions on Applied Superconductivity](#) **33**, 1 (2023) (Cited on p. 84).
- [112] D. C. Mattis and J. Bardeen, “Theory of the Anomalous Skin Effect in Normal and Superconducting Metals”, [Physical Review](#) **111**, 412 (1958) (Cited on p. 86).
- [113] R. Barends, N. Vercruyssen, A. Endo, P. J. de Visser, T. Zijlstra, T. M. Klapwijk, and J. J. A. Baselmans, “Reduced frequency noise in superconducting resonators”, [Applied Physics Letters](#) **97**, 033507 (2010) (Cited on p. 87).
- [114] L. Zhang, L. You, L. Ying, W. Peng, and Z. Wang, “Characterization of surface oxidation layers on ultrathin NbTiN films”, [Physica C: Superconductivity and its Applications](#) **545**, 1 (2018) (Cited on p. 87).

- [115] S. A. Nazib, T. A. Hutchins-Delgado, M. V. Reymatias, H. Lee, E. Jamil, G. A. Smolyakov, T.-M. Lu, S. A. Ivanov, J. Nogan, I. Komissarov, R. Sobolewski, and M. Osiński, “Native Oxide Formation in NbTiN Thin Films for Traveling-Wave Superconducting Nanostripe Single-Photon Detectors”, in [Optica Quantum 2.0 Conference and Exhibition \(2023\)](#), paper [QW2A.34](#) (2023), QW2A.34 (Cited on p. 87).
- [116] N. P. Lockyer, S. Aoyagi, J. S. Fletcher, I. S. Gilmore, P. A. W. van der Heide, K. L. Moore, B. J. Tyler, and L.-T. Weng, “Secondary ion mass spectrometry”, [Nature Reviews Methods Primers](#) **4**, 1 (2024) (Cited on p. 87).
- [117] *QSpec Technology, Inc.* http://qspec.us/About_Us.html (Cited on p. 87).
- [118] M. V. Burdastyh, S. V. Postolova, T. Proslier, S. S. Ustavshikov, A. V. Antonov, V. M. Vinokur, and A. Y. Mironov, “Superconducting phase transitions in disordered NbTiN films”, [Scientific Reports](#) **10**, 1471 (2020) (Cited on p. 87).
- [119] *X-Ray Photoelectron Spectroscopy (Allentown Lab)*, <https://www.intertek.com/analytical-laboratories/allentown/x-ray-photoelectron-spectroscopy/> (Cited on p. 88).
- [120] C. T. Earnest, J. H. Béjanin, T. G. McConkey, E. A. Peters, A. Korinek, H. Yuan, and M. Mariantoni, “Substrate surface engineering for high-quality silicon/aluminum superconducting resonators”, [Superconductor Science and Technology](#) **31**, 125013 (2018) (Cited on p. 89).
- [121] C. U. Lei, S. Ganjam, L. Krayzman, A. Banerjee, K. Kisslinger, S. Hwang, L. Frunzio, and R. J. Schoelkopf, “Characterization of Microwave Loss Using Multimode Superconducting Resonators”, [Physical Review Applied](#) **20**, 024045 (2023) (Cited on p. 89).
- [122] D. S. Wisbey, J. Gao, M. R. Vissers, F. C. S. da Silva, J. S. Kline, L. Vale, and D. P. Pappas, “Effect of metal/substrate interfaces on radio-frequency loss in superconducting coplanar waveguides”, [Journal of Applied Physics](#) **108**, 093918 (2010) (Cited on p. 89).
- [123] G. Koley and M. G. Spencer, “On the origin of the two-dimensional electron gas at the AlGaIn/GaN heterostructure interface”, [Applied Physics Letters](#) **86**, 042107 (2005) (Cited on p. 90).
- [124] J. W. Park, D. F. Bogorin, C. Cen, D. A. Felker, Y. Zhang, C. T. Nelson, C. W. Bark, C. M. Folkman, X. Q. Pan, M. S. Rzchowski, J. Levy, and C. B. Eom, “Creation of a two-dimensional electron gas at an oxide interface on silicon”, [Nature Communications](#) **1**, 94 (2010) (Cited on p. 90).
- [125] N. C. Bristowe, P. Ghosez, P. B. Littlewood, and E. Artacho, “The origin of two-dimensional electron gases at oxide interfaces: insights from theory”, [Journal of Physics: Condensed Matter](#) **26**, 143201 (2014) (Cited on p. 90).
- [126] H. Zhou, E. Li, K. Godeneli, Z.-H. Zhang, S. Jahanbani, K. Yu, M. Odeh, S. Aloni, S. Griffin, and A. Sipahigil, *Observation of Interface Piezoelectricity in Superconducting Devices on Silicon*, 2024 (Cited on p. 90).
- [127] W. A. Phillips, “Two-level states in glasses”, [Reports on Progress in Physics](#) **50**, 1657 (1987) (Cited on pp. 93, 97, 102, 109).
- [128] S. Hunklinger and W. Arnold, “3 - Ultrasonic Properties of Glasses at Low Temperatures”, in [Physical Acoustics](#), Vol. 12, edited by W. P. Mason and R. N. Thurston (Academic Press, 1976), pp. 155–215 (Cited on p. 93).

- [129] P. w. Anderson, B. I. Halperin, and c. M. Varma, “Anomalous low-temperature thermal properties of glasses and spin glasses”, [The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics](#) **25**, 1 (1972) (Cited on p. 93).
- [130] J. Wenner, R. Barends, R. C. Bialczak, Y. Chen, J. Kelly, E. Lucero, M. Mariantoni, A. Megrant, P. J. J. O’Malley, D. Sank, A. Vainsencher, H. Wang, T. C. White, Y. Yin, J. Zhao, A. N. Cleland, and J. M. Martinis, “Surface loss simulations of superconducting coplanar waveguide resonators”, [Applied Physics Letters](#) **99**, 113513 (2011) (Cited on p. 94).
- [131] 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”, [Physical Review Applied](#) **19**, 034064 (2023) (Cited on pp. 94, 95).
- [132] *Ansys HFSS / 3D High Frequency Simulation Software* (Cited on p. 94).
- [133] A. Padovani, P. La Torraca, J. Strand, L. Larcher, and A. L. Shluger, “Dielectric breakdown of oxide films in electronic devices”, [Nature Reviews Materials](#) **9**, 607 (2024) (Cited on p. 97).
- [134] S. M. Sze, “Current Transport and Maximum Dielectric Strength of Silicon Nitride Films”, [Journal of Applied Physics](#) **38**, 2951 (1967) (Cited on p. 97).
- [135] R. Apodaca, *Electrical Characterization of Silicon-Rich Nitride and Silicon Oxynitride Films Deposited by Low-Pressure Chemical Vapor Deposition*. Tech. rep. SAND2005-7280P (Sandia National Lab. (SNL-NM), Albuquerque, NM (United States), 2005) (Cited on p. 97).
- [136] E. A. Wollack, A. Y. Cleland, P. Arrangoiz-Arriola, T. P. McKenna, R. G. Gruenke, R. N. Patel, W. Jiang, C. J. Sarabalis, and A. H. Safavi-Naeini, “Loss channels affecting lithium niobate phononic crystal resonators at cryogenic temperature”, [Applied Physics Letters](#) **118**, 123501 (2021) (Cited on pp. 98, 116).
- [137] 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**, 840 (2020) (Cited on p. 98).
- [138] A. L. Emser, C. Metzger, B. C. Rose, and K. W. Lehnert, *Thin-film quartz for high-coherence piezoelectric phononic crystal resonators*, 2024 (Cited on pp. 98, 100).
- [139] M. Hegedüs, R. Banerjee, A. Hutcheson, T. Barker, S. Mahashabde, A. V. Danilov, S. E. Kubatkin, V. Antonov, and S. E. de Graaf, *In-situ scanning gate imaging of individual two-level material defects in live superconducting quantum circuits*, 2024 (Cited on p. 98).
- [140] M. Trif, O. Dmytruk, H. Bouchiat, R. Aguado, and P. Simon, “Dynamic current susceptibility as a probe of Majorana bound states in nanowire-based Josephson junctions”, [Physical Review B](#) **97**, 041415 (2018) (Cited on p. 100).
- [141] L. Faoro and L. B. Ioffe, “Interacting tunneling model for two-level systems in amorphous materials and its predictions for their dephasing and noise in superconducting microresonators”, [Physical Review B](#) **91**, 014201 (2015) (Cited on p. 105).
- [142] M. S. Khalil, M. J. A. Stoutimore, F. C. Wellstood, and K. D. Osborn, “An analysis method for asymmetric resonator transmission applied to superconducting devices”, [Journal of Applied Physics](#) **111**, 054510 (2012) (Cited on pp. 111, 189, 190).
- [143] S. Probst, F. B. Song, P. A. Bushev, A. V. Ustinov, and M. Weides, “Efficient and robust analysis of complex scattering data under noise in microwave resonators”, [Review of Scientific Instruments](#) **86**, 024706 (2015) (Cited on pp. 111, 189).

- [144] E. T. Swartz and R. O. Pohl, “Thermal boundary resistance”, [Reviews of Modern Physics](#) **61**, 605 (1989) (Cited on pp. 116, 117).
- [145] S.-M. Lee and D. G. Cahill, “Heat transport in thin dielectric films”, [Journal of Applied Physics](#) **81**, 2590 (1997) (Cited on p. 116).
- [146] S. Bai, Z. Tang, Z. Huang, and J. Yu, “Thermal Characterization of Si₃N₄ Thin Films Using Transient Thermoreflectance Technique”, [IEEE Transactions on Industrial Electronics](#) **56**, 3238 (2009) (Cited on p. 117).
- [147] J. L. Black and B. I. Halperin, “Spectral diffusion, phonon echoes, and saturation recovery in glasses at low temperatures”, [Physical Review B](#) **16**, 2879 (1977) (Cited on p. 117).
- [148] A. Würger and D. Bodea, “Thermal conductivity by two-level systems in glasses”, [Chemical Physics](#) **296**, 301 (2004) (Cited on p. 117).
- [149] S. Shu, N. Klimovich, B. H. Eom, A. D. Beyer, R. B. Thakur, H. G. Leduc, and P. K. Day, “Nonlinearity and wide-band parametric amplification in a (Nb,Ti)N microstrip transmission line”, [Physical Review Research](#) **3**, 023184 (2021) (Cited on p. 118).
- [150] M. Checchin, D. Frolov, A. Lunin, A. Grassellino, and A. Romanenko, “Measurement of the Low-Temperature Loss Tangent of High-Resistivity Silicon Using a High- Q Superconducting Resonator”, [Physical Review Applied](#) **18**, 034013 (2022) (Cited on p. 118).
- [151] 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”, [Physical Review Applied](#) **19**, 034064 (2023) (Cited on p. 118).
- [152] K. Luke, A. Dutt, C. B. Poitras, and M. Lipson, “Overcoming Si₃N₄ film stress limitations for high quality factor ring resonators”, [Optics Express](#) **21**, 22829 (2013) (Cited on p. 118).
- [153] X. Ji, S. Roberts, M. Corato-Zanarella, and M. Lipson, “Methods to achieve ultra-high quality factor silicon nitride resonators”, [APL Photonics](#) **6**, 071101 (2021) (Cited on p. 118).
- [154] H. Kurita, M. Nakamura, H. Miyagawa, and Y. Kamigaki, “Effect of N₂-Anneal Temperature on Silicon Nitride Film: (I) Time-Dependent Dielectric Breakdown and ESR Evaluations”, [ECS Transactions](#) **108**, 69 (2022) (Cited on p. 121).
- [155] I. Parkhomenko, L. Vlasukova, F. Komarov, N. Kovalchuk, S. Demidovich, A. Zhussupbekova, K. Zhussupbekov, I. V. Shvets, O. Milchanin, D. Zhigulin, and I. Romanov, “Effect of Rapid Thermal Annealing on Si-Based Dielectric Films Grown by ICP-CVD”, [ACS Omega](#) **8**, 30768 (2023) (Cited on p. 121).
- [156] S. Tönnberg, “Optimisation and characterisation of LPCVD silicon nitride thin film growth”, (2006) (Cited on p. 121).
- [157] *kSA MOS UltraScan/ ThermalScan*, <https://k-space.com/product/mos-scan/> (Cited on p. 121).
- [158] W. Jiang, D. Xu, S. Yao, B. Xiong, and Y. Wang, “Effect of hyperthermal annealing on LPCVD silicon nitride”, [Materials Science in Semiconductor Processing](#) **43**, 222 (2016) (Cited on p. 121).
- [159] E. Herth, B. Legrand, L. Buchaillot, N. Rolland, and T. Lasri, “Optimization of SiNX:H films deposited by PECVD for reliability of electronic, microsystems and optical applications”, [Microelectronics Reliability](#) **50**, 1103 (2010) (Cited on p. 123).

- [160] R. Boyle, S. Sukumaran, and R. Chen, *Quantitative determination of hydrogen concentration in silicon nitride dielectric films on silicon wafers using FTIR spectroscopy*, Application Note (Thermo Fisher Scientific, 2022) (Cited on p. 123).
- [161] V. Mazet, C. Carteret, D. Brie, J. Idier, and B. Humbert, “Background removal from spectra by designing and minimising a non-quadratic cost function”, *Chemometrics and Intelligent Laboratory Systems* **76**, 121 (2005) (Cited on p. 124).
- [162] W. A. Lanford and M. J. Rand, “The hydrogen content of plasma-deposited silicon nitride”, *Journal of Applied Physics* **49**, 2473 (1978) (Cited on p. 124).
- [163] A. M. Holder, K. D. Osborn, C. J. Lobb, and C. B. Musgrave, “Bulk and Surface Tunneling Hydrogen Defects in Alumina”, *Physical Review Letters* **111**, 065901 (2013) (Cited on p. 126).
- [164] L. Gordon, H. Abu-Farsakh, A. Janotti, and C. G. Van de Walle, “Hydrogen bonds in Al₂O₃ as dissipative two-level systems in superconducting qubits”, *Scientific Reports* **4**, 7590 (2014) (Cited on p. 126).
- [165] W. Morr, A. Müller, G. Weiss, H. Wipf, and B. Golding, “Isotope dependence of hydrogen tunneling in niobium”, *Physical Review Letters* **63**, 2084 (1989) (Cited on p. 126).
- [166] D. Steinbinder, H. Wipf, A.-J. Dianoux, A. Magerl, K. Neumaier, D. Richter, and R. Hempelmann, “Quantum Diffusion of Trapped-Hydrogen Interstitials in Nb: The Role of the Tunnel Splitting”, *Europhysics Letters* **16**, 211 (1991) (Cited on p. 126).
- [167] G. Cannelli, R. Cantelli, F. Cordero, and F. Trequattrini, “Tunneling of H and D in Metals and Semiconductors”, in *Tunneling Systems in Amorphous and Crystalline Solids*, edited by P. Esquinazi (Springer, Berlin, Heidelberg, 1998), pp. 389–458 (Cited on p. 126).
- [168] L. Faoro and L. B. Ioffe, “Interacting tunneling model for two-level systems in amorphous materials and its predictions for their dephasing and noise in superconducting microresonators”, *Physical Review B* **91**, 014201 (2015) (Cited on pp. 127, 143).
- [169] S. Hunklinger and A. K. Raychaudhuri, “Chapter 3: Thermal and Elastic Anomalies in Glasses at Low Temperatures”, in *Progress in Low Temperature Physics*, Vol. 9, edited by D. F. Brewer (Elsevier, 1986), pp. 265–344 (Cited on p. 127).
- [170] S. T. Skacel, C. Kaiser, S. Wuensch, H. Rotzinger, A. Lukashenko, M. Jerger, G. Weiss, M. Siegel, and A. V. Ustinov, “Probing the density of states of two-level tunneling systems in silicon oxide films using superconducting lumped element resonators”, *Applied Physics Letters* **106**, 022603 (2015) (Cited on pp. 127, 128).
- [171] M. Göppl, A. Fragner, M. Baur, R. Bianchetti, S. Filipp, J. M. Fink, P. J. Leek, G. Puebla, L. Steffen, and A. Wallraff, “Coplanar waveguide resonators for circuit quantum electrodynamics”, *Journal of Applied Physics* **104**, 113904 (2008) (Cited on p. 127).
- [172] S. Gevorgian, L. Linner, and E. Kollberg, “CAD models for shielded multilayered CPW”, *IEEE Transactions on Microwave Theory and Techniques* **43**, 772 (1995) (Cited on p. 127).
- [173] C. E. Murray, J. M. Gambetta, D. T. McClure, and M. Steffen, “Analytical Determination of Participation in Superconducting Coplanar Architectures”, *IEEE Transactions on Microwave Theory and Techniques* **66**, 3724 (2018) (Cited on p. 127).
- [174] R. Garg, I. Bahl, M. Bozzi, and V. Kane, *Microstrip Lines and Slotlines* (1979) (Cited on p. 127).

- [175] T. Capelle, E. Flurin, E. Ivanov, J. Palomo, M. Rosticher, S. Chua, T. Briant, P.-F. Cohadon, A. Heidmann, T. Jacqmin, and S. Deléglise, “Probing a Two-Level System Bath via the Frequency Shift of an Off-Resonantly Driven Cavity”, [Physical Review Applied](#) **13**, 034022 (2020) (Cited on p. 128).
- [176] B. D. Hauer, P. H. Kim, C. Doolin, F. Souris, and J. P. Davis, “Two-level system damping in a quasi-one-dimensional optomechanical resonator”, [Physical Review B](#) **98**, 214303 (2018) (Cited on p. 128).
- [177] X. Zhou, D. Cattiaux, R. R. Gazizulin, A. Luck, O. Maillet, T. Crozes, J.-F. Motte, O. Bourgeois, A. Fefferman, and E. Collin, “On-chip Thermometry for Microwave Optomechanics Implemented in a Nuclear Demagnetization Cryostat”, [Physical Review Applied](#) **12**, 044066 (2019) (Cited on p. 128).
- [178] Y. Seis, T. Capelle, E. Langman, S. Saarinen, E. Planz, and A. Schliesser, “Ground state cooling of an ultracoherent electromechanical system”, [Nature Communications](#) **13**, 1507 (2022) (Cited on pp. 128, 131).
- [179] R. Kubo, “The fluctuation-dissipation theorem”, [Reports on Progress in Physics](#) **29**, 255 (1966) (Cited on p. 130).
- [180] A. Youssefi, S. Kono, M. Chegnizadeh, and T. J. Kippenberg, “A squeezed mechanical oscillator with millisecond quantum decoherence”, [Nature Physics](#) **19**, 1697 (2023) (Cited on pp. 131, 150).
- [181] J. M. Fink, M. Kalaei, A. Pitanti, R. Norte, L. Heinzle, M. Davanço, K. Srinivasan, and O. Painter, “Quantum electromechanics on silicon nitride nanomembranes”, [Nature Communications](#) **7**, 12396 (2016) (Cited on p. 131).
- [182] Y. Liu, J. Zhou, L. M. de Lépinay, and M. A. Sillanpää, “Quantum backaction evading measurements of a silicon nitride membrane resonator”, [New Journal of Physics](#) **24**, 083043 (2022) (Cited on p. 131).
- [183] M. Constantin, C. C. Yu, and J. M. Martinis, “Saturation of two-level systems and charge noise in Josephson junction qubits”, [Physical Review B](#) **79**, 094520 (2009) (Cited on p. 140).
- [184] A. Shnirman, G. Schön, I. Martin, and Y. Makhlin, “Low- and High-Frequency Noise from Coherent Two-Level Systems”, [Physical Review Letters](#) **94**, 127002 (2005) (Cited on pp. 140, 142, 143).
- [185] N. E. Frattini, V. V. Sivak, A. Lingenfelter, S. Shankar, and M. H. Devoret, “Optimizing the Nonlinearity and Dissipation of a SNAIL Parametric Amplifier for Dynamic Range”, [Physical Review Applied](#) **10**, 054020 (2018) (Cited on p. 150).
- [186] C. Macklin, K. O’Brien, D. Hover, M. E. Schwartz, V. Bolkhovskiy, X. Zhang, W. D. Oliver, and I. Siddiqi, “A near-quantum-limited Josephson traveling-wave parametric amplifier”, [Science](#) **350**, 307 (2015) (Cited on p. 150).
- [187] M. Malnou, J. Aumentado, M. Vissers, J. Wheeler, J. Hubmayr, J. Ullom, and J. Gao, “Performance of a Kinetic Inductance Traveling-Wave Parametric Amplifier at 4 Kelvin: Toward an Alternative to Semiconductor Amplifiers”, [Physical Review Applied](#) **17**, 044009 (2022) (Cited on pp. 150, 151).
- [188] R. D. Delaney, A. P. Reed, R. W. Andrews, and K. W. Lehnert, “Measurement of Motion beyond the Quantum Limit by Transient Amplification”, [Physical Review Letters](#) **123**, 183603 (2019) (Cited on p. 150).

- [189] K. Williams and R. Muller, “Etch rates for micromachining processing”, *Journal of Microelectromechanical Systems* **5**, 256 (1996) (Cited on p. 167).
- [190] K. Williams, K. Gupta, and M. Wasilik, “Etch rates for micromachining processing-Part II”, *Journal of Microelectromechanical Systems* **12**, 761 (2003) (Cited on p. 167).
- [191] M. R. Vissers, J. Gao, D. S. Wisbey, D. A. Hite, C. C. Tsuei, A. D. Corcoles, M. Steffen, and D. P. Pappas, “Low loss superconducting titanium nitride coplanar waveguide resonators”, *Applied Physics Letters* **97**, 232509 (2010) (Cited on p. 167).
- [192] K. Sato, M. Shikida, Y. Matsushima, T. Yamashiro, K. Asaumi, Y. Iriye, and M. Yamamoto, “Characterization of orientation-dependent etching properties of single-crystal silicon: effects of KOH concentration”, *Sensors and Actuators A: Physical, Tenth IEEE International Workshop on Micro Electro Mechanical Systems* **64**, 87 (1998) (Cited on p. 177).
- [193] *Spectral Thin Film Reflectance Calculator for Thin-Film Stacks*, <https://www.filmetrics.com/reflectance-calculator> (Cited on p. 187).
- [194] P. G. Baity, C. Maclean, V. Seferai, J. Bronstein, Y. Shu, T. Hemakumara, and M. Weides, “Circle fit optimization for resonator quality factor measurements: Point redistribution for maximal accuracy”, *Physical Review Research* **6**, 013329 (2024) (Cited on pp. 189, 191).
- [195] D. Rieger, S. Günzler, M. Spiecker, A. Nambisan, W. Wernsdorfer, and I. Pop, “Fano Interference in Microwave Resonator Measurements”, *Physical Review Applied* **20**, 014059 (2023) (Cited on p. 189).
- [196] *SteeleLab / Technical Note Fano · GitLab*, <https://gitlab.tudelft.nl/steelelab/technical-note-fano>, 2022 (Cited on p. 190).

Appendix A

Transducer Fabrication

In this appendix, I will step through the fabrication process for our transducers, with details of the etch chemistry and comments on likely pitfalls. I will provide gas flow rates and plasma powers where applicable, but these are specific to the tools in the NIST BMF cleanroom. These can also change overtime, so consider the values provided here as historical record, as opposed to eternal truth. A good resource of etch chemistries and material compatibilities is [Williams and Muller 1996; Williams, Gupta, et al. 2003] as well as the NIST BMF wiki. For all fabricated devices, we start with high-resistivity, "floatzone", three-inch Si wafers aligned to the (1-0-0) crystal axis and double-side polished (DSP). The wafers have a resistivity $> 10,000 \Omega \cdot \text{cm}$ and are $381 \pm 10 \mu\text{m}$ thick. Before fabrication, we optionally perform a 10 minute nanostrip clean at 80 C to remove any contaminants.

A.1 Circuit and mirror wafer

The process flow for the wafer of circuit chips, which eventually is bonded to a mirror wafer and diced into chips, is described in Fig. A.2. The NbTiN deposition is performed by Michael Vissers, using process similar to that developed for TiN [Vissers et al. 2010]. Right before entering the deposition chamber of the tool "PVD Big", an anhydrous vapor HF etch of the Si wafer is performed to remove the native oxide and hydrogen terminate the surface. The deposition is performed at 500 C, and it is thought that a thin $\approx 2 \text{ nm}$ layer of SiN forms in the predeposition steps, which aids in lattice matching of the superconductor to the substrate and generally higher T_c . After deposition of the 200 nm film, we use a four-point probe to measure the sheet-resistivity, which is $6\text{-}8 \Omega/\square$

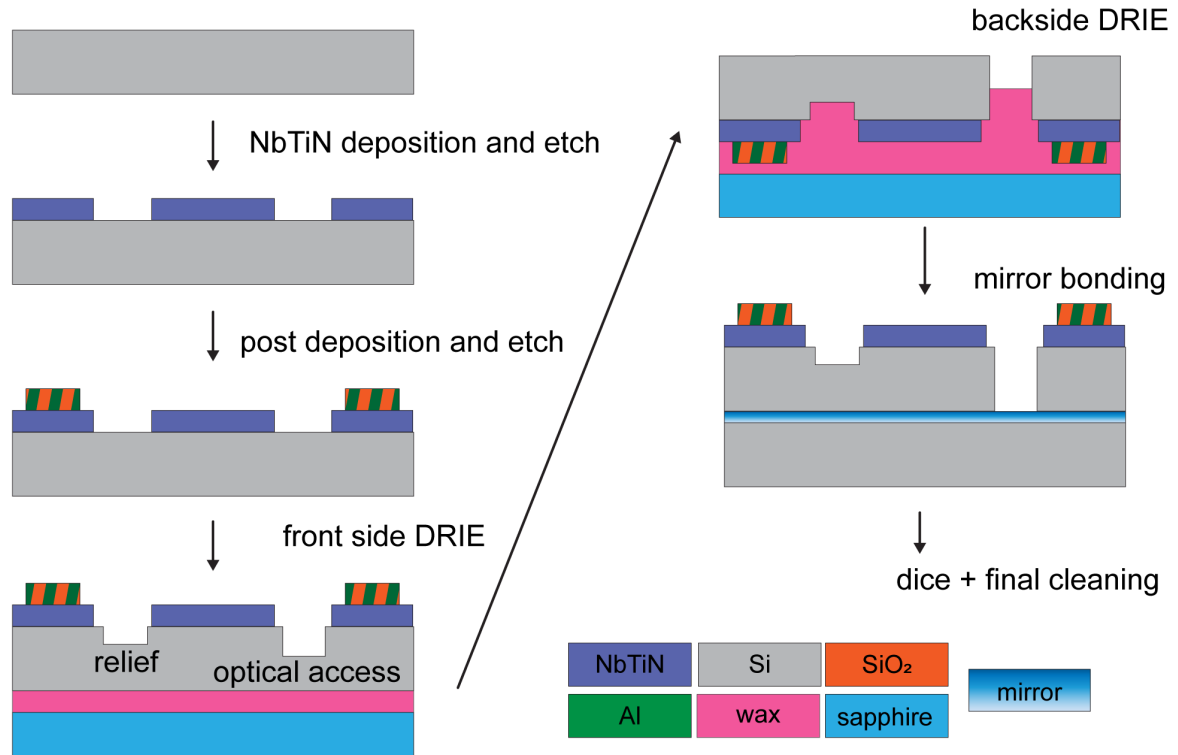


Figure A.1: **Circuit and mirror wafer fabrication flow.** Step-by-step description of process developed to fabricate wafer of circuit chips bonded to mirror chips. Photolithography steps are left out for clarity, but are described in the text. The striped structures are made of either Al or SiO₂, depending on how whether the circuit chip and membrane chip are designed for bonding. Geometries are not drawn to scale.

for a properly grown film. We will reserve a section of the wafer to measure the sheet resistivity throughout the process.

Then, we pattern a layer of photoresist in order to etch the circuit geometry. We spin and pattern a $1\text{ }\mu\text{m}$ layer of SPR660L using the following procedure:

- (optional) Oxygen plasma in Asher with 50 sccm O_2 at 60 W for 2-3 min to promote resist adhesion
- Spin P-20 primer and then SPR660L resist at 2600 rpm for 40 s
- Pre-exposure bake at 95 C for 60 s
- Expose design using Heidelberg MLA (Mask-Less Aligner) with 405 nm in "Quality Mode" and standard dosage and focus for SPR660L.
- Post-exposure bake at 110 C for 60 s
- Develop with MF-26A in two 30 s intervals, spin-dry
- (optional) Additional spin-rinse dry to remove resist residue on backside of wafer from dirty chuck in the previous step.
- Inspect, first with red light, then white light

In this first photolithography step, there are no alignment marks as nothing has been patterned yet. The design for the first layer should include alignment marks so that subsequent designs are placed properly. The standard pattern used in this thesis was "+" signs at $(x, y) = (0, 33), (0, -33), (33, 0)$, and $(-33, 0)$ mm where $(0, 0)$ represents the center of the wafer. The minimum feature size that can be patterned with the MLA is $2.5\text{ }\mu\text{m}$, which is smaller than any of the features in the resonator designs used in this thesis. This process will be our standard recipe for photolithography with $1\text{ }\mu\text{m}$ resist.

Now, we etch the NbTiN in the Oxford Fluorine etcher using the recipe, "CMB TiN undr less O_2 ". This process uses 2.5 sccm of O_2 , 25 sccm of CHF_3 and 90 sccm of Ar. The inductively coupled plasma (ICP) power is 1000 W and the RF power is 100 W. For plasma etches or cleans that could affect the NbTiN surface, we generally use plasmas with lower oxygen gas flows in order to prevent oxide growth on the metal surface. After the etch is complete, which we track using the endpoint detector, the photoresist is stripped using our standard procedure:

- Oxygen plasma in Asher with 50 sccm O_2 at 60 W for 2-5 min to etch some of the resist
- 2 min sonication at 70% power in acetone
- 2 min sonication at 70% power in fresh acetone
- Dip into IPA
- Spin-rinse dry

The height of the NbTiN is measured using a profilometer. The next step is to deposit and pattern the structural posts which define the spacing of the mechanically-compliant capacitor pad. In the device used for [Brubaker et al. 2022] and [R. Delaney 2022], these posts were made by sputtering Al in the SIS tool with the recipe "QubitBaseCleanAlDep". The sputtering performed at 5 mTorr with 300 W power on the Al gun. As these posts are purely structural, use of the highly-controlled and clean SIS is likely overkill, and this process could be completed more quickly in one of the evaporators. The Al is patterned using our standard $1\ \mu\text{m}$ resist procedure and then etched in Transene Type A at 50 C for 30 s, after which the resist is stripped.

To fabricate devices compatible with bonding (section 3.2.2), the posts are instead made of SiO_2 deposited in the Oxford ECR tool. The standard oxide recipe is tweaked to remove oxygen from the pre-deposition RF clean. The new recipe, "SiO2_rf_clean_cu_PSB", uses 139 sccm of Ar, 118 sccm of O_2 , and 18.5 sccm of SiH_4 (silane) during the deposition, with 720 W of microwave power and 220 W of RF power. The SiO_2 pattern is defined with $1\ \mu\text{m}$ resist and then etched in buffered oxide etchant (BOE), which is a mix of HF and some stabilizing chemicals, after which the resist is stripped. In either the Transene or BOE etches, we do not agitate the wafer, as the small posts have been observed to wash away if the wafer is agitated too strongly in the etchant.

Now, we wax the backside of the circuit wafer to a sapphire carrier in order to protect the back surface from the upcoming deep reactive ion etch (DRIE) on the front side.¹ A 2:1 ratio by weight of CrystalBond wax to EBR solvent is mixed a few days in advance to give time for the wax to dissolve into the solvent. To wax the wafer to the sapphire carrier:

¹We observed that the DRIE roughened the edge of the backside of the wafer, which will make bonding to the mirror wafer difficult around the edge of the circuit wafer

- Ash the "wax-side" of the double-side polished sapphire carrier wafer and Si sample wafer with 50 sccm O₂ at 60 W for 2-3 min
- Using plastic disposable droppers with the tips cut off (helps prevent air bubbles), cover the 2/3 - 3/4 "wax-side" of each wafer with wax, and then spin at 2000 rpm for 45 s.
- Place each wafer into its own waxer disc and bake in vacuum oven at 148 C for 20-30 min.
- Remove wafers from oven and let cool for 2 min while propped out of the waxer disc.
- Using the flats of each wafer to align them, carefully sandwich together the wafers.
- Place the sandwiched wafers into a waxer disc with the Si wafer pointing downwards. Top with a weight disc and bake in the vacuum oven at 148 C for 40-60 min.
- After the wafer cools, spin clean with acetone and IPA to remove any wax that leaked out or found its way to the Si surface.

It is important to avoid air bubbles in the waxing process, as air can violently escape from them when the hardened wax is placed into the vacuum chambers of the cleanroom tools. The wax layer should also be a uniform thickness, so that alignment in the MLA is accurate.

At this point, we want to relieve the Si such that the posts and microwave circuit are up on a "mesa" of sorts. This prevents dust particles or other contaminants from setting the spacing of the capacitor pads, as the surrounding Si is etched away using the STS DRIE. The pattern for the DRIE is defined using 3 or 7 μm resist, depending on the thickness of Si to be etched. Considering the wafer is waxed to a carrier, we must also be careful to not remelt the wax during the baking steps. To pattern this thicker resist layer:

- (optional) Oxygen plasma in Asher with 50 sccm O₂ at 60 W for 2-3 min to promote resist adhesion
- Spin P-20 primer and then SPR-222-3.0 or SPR-220-7.0 resist at 1500 rpm for 40 s
- Pre-exposure bake at 115 C for 90 s using a "no touch" recipe that avoids direct contact between the hotplate and wafer (recipe: SPR220-7 3IN N0 TOUCH KAT)
- Expose design using Heidelberg MLA (Mask-Less Aligner) with 405 nm in "Quality Mode" and standard dosage and focus for SPR-222-3.0 or SPR-220-7.0.
- No post-exposure bake
- Develop with MF-26A in several 60 s intervals, inspecting after every few intervals to see if resist has cleared, spin-dry

- (optional) Additional spin-rinse dry to remove resist residue on backside of wafer from dirty chuck in the previous step.
- Inspect, first with red light, then white light

We etch the top 12-15 μm of the Si away in the DRIE using the "specbnod" recipe. This recipe has two steps per etch cycle: a 12 s etch step using 130 sccm of SF_6 and a 8 s passivation step using 85 sccm of C_4F_4 which coats the etch sidewalls in a teflon like material to prevent horizontal etching. This is also known as the Bosch process and is used to etch high-aspect ratio structures in Si. During the resist stripping procedure after the etch is complete, we sonicate for 5 min to remove the thicker resist.

Ideally, the next step is to DRIE the optical access hole which will go all the way through the wafer. However, it is difficult to spin resist over the now non-uniform top surface. There will be thick layers of resist in the flat areas of the wafer, but it is difficult to get sufficiently thick layers over the posts and metal, which are now elevated above the surrounding Si by several microns. Instead, we remove the wafer from the sapphire carrier and flip it over. The bottom side of the wafer is flat, and we can align the pattern using backside alignment in the MLA as the sapphire and wax are both transparent.² To remove the Si wafer from the wax,

- Place the wafer sandwich on a hotplate at 148 C. Use the vacuum on the hotplate to hold the wafers in place.
- Once the wax is sufficiently melted, use a pair of metal wafer tweezers to push the Si wafer off of the sapphire wafer
- 5 min sonication at 100% power in acetone
- 5 min sonication at 100% power in fresh acetone
- Dip into IPA
- Spin-rinse dry

Now, we repeat the same steps to wax the wafer to the sapphire carrier, except now with the frontside buried in the wax. Using 7 μm resist, we etch the optical access hole in the DRIE.

²Recently, we have attempted to etch as much of the optical access hole from the front side as possible while carefully tracking the resist thickness. We believe this will prevent overetching of the DRIE that occurs in the procedure as described.

Approximately every 100 cycles of the Bosch process, we rotate the wafer orientation in the chamber for etch uniformity. After removing the wafer from the wax again, we are ready to wafer bond the mirror wafer which was described in [section 3.2.1](#). The mirror wafer is a Si wafer onto which a dielectric Bragg mirror has been deposited by [*FiveNine Optics – Boulder, Colorado* / $R > 99.999\%$ 2024]. The top layer of the mirror is SiO₂, and 3 μm layer of SiO₂ is also deposited on the backside as a stress-compensation layer. To bond the mirrors together:

- 2 min sonication at 70% power in acetone, 2 min sonication at 70% power in fresh acetone, dip IPA, and SRD both wafers
- In the Ontos Atmospheric Plasma tool, oxygen ash the surfaces to be bonded together³
- In the Ontos Atmospheric Plasma tool, hydrogen active the surfaces to be bonded together
- Spin-rinse dry to create -OH bonds at the surface
- Using the wafer bonding apparatus, align the two wafers using the flats. Using the back of a set of wafer tweezers, strongly press the two wafers together in the middle and retract the fingers suspending the circuit wafer above the mirror wafer.
- Without releasing pressure at the center of the wafer, use a separate set of tweezers to press the wafers together, being careful to avoid pressing on the microwave circuit or structural posts.
- Construct a makeshift oven with two wafer cooling plates and place on a hotplate. Anneal the bonded wafers in the makeshift oven at 150 C for at least one hour.
- Inspect under an infrared camera to look for voids

At this point, the wafers should be rigidly attached to each other. The stress compensation layer is removed using an anhydrous HF vapor for 90-100 min with then help of a chuck that makes the etch one-sided, protecting the Al or SiO₂ posts on the front side of the wafer. The last step is to dice the wafer into 9 mm chips using the DISCO DAD dicing saw. We etch the dicing lines into the front side of the wafer during the front side DRIE that creates the relieved structures. Before dicing, we spin 1 μm onto the wafer and run the post-exposure bake to hard bake the resist and protect the circuit from the dicing process. The dicing resist is easily removed by our standard photoresist stripping procedure. Finally, the chips are cleaned in nanostrip at 80 C for 2 min and spin-dried with

³The purpose of the ash is to clean the surfaces, so it can be removed or replaced by a nanostrip clean before the solvent clean.

DI water using a custom vacuum chuck capable of holding chips. No organics should be used to clean the chip after the last nanostrip clean. Note, this nanostrip clean is not compatible with the Al posts, and is skipped for those chips.

A.2 Membrane wafer

Fabrication of the membrane wafer is more involved than the circuit chip as the membrane must be released and significant micromaching with the DRIE is required to define the PNC (Fig. A.2). Similar to the circuit wafer, we start with a clean FZ Si wafer, which is optionally cleaned in nanostrip at 80 C for 10 min. The Si_3N_4 deposition is performed by Jim Beall in the Tystar TYTAN low-pressure chemical vapor deposition (LPCVD) tool. Using the precursor gasses ammonia (NH_3) and dichlorosilane (SiCl_2H_2) in a 3:1 ratio, we deposit a 100 nm film of Si_3N_4 at 770°C. Post-deposition, we measure a film stress of around 1 GPa and a refractive index at 623 nm of 2.1. These measurements are consistent with a 3:4 stoichiometric ratio of Si and N [Paik and Osborn 2010]. For the annealed planar resonators, we perform a four hour anneal of the Si_3N_4 at 1100°C in a pure O_2 atmosphere. The atmosphere of O_2 is recommended by the furnace manufacturer to drive out the hydrogen impurities better than N_2 or vacuum. For the annealed transducer devices discussed towards the end of this thesis, the Si_3N_4 is deposited with a 7:1 ratio of $\text{NH}_3:\text{SiCl}_2\text{H}_2$ in order to achieve high tensile stress (see section 5.3.1) and then annealed.

We next sputter a 30 nm layer of NbTiN, again performed by Mike Vissers. As there is already a layer of Si_3N_4 , we skip the anhydrous vapor HF clean before the sputter. The difference in morphology of the high-stress Si_3N_4 and the Si_3N_4 grown in the chamber for bare Si wafers could explain the difference in T_c we observe for NbTiN on Si vs NbTiN on Si_3N_4 on Si (see Fig. 4.9). For this film, we find the sheet resistance to be 42-48 Ω/\square , and again, a section of the wafer is reserved to track this sheet resistance throughout the fabrication process.

To etch the NbTiN, we cannot use the Oxford Fl etcher, as that recipe also etches Si_3N_4 very quickly (>10 nm/s). Instead, we first deposit and pattern a 100 nm layer of SiO_2 in the Oxford ECR which will serve as a hard mask. We use the standard 1 μm photoresist process, and etch the SiO_2

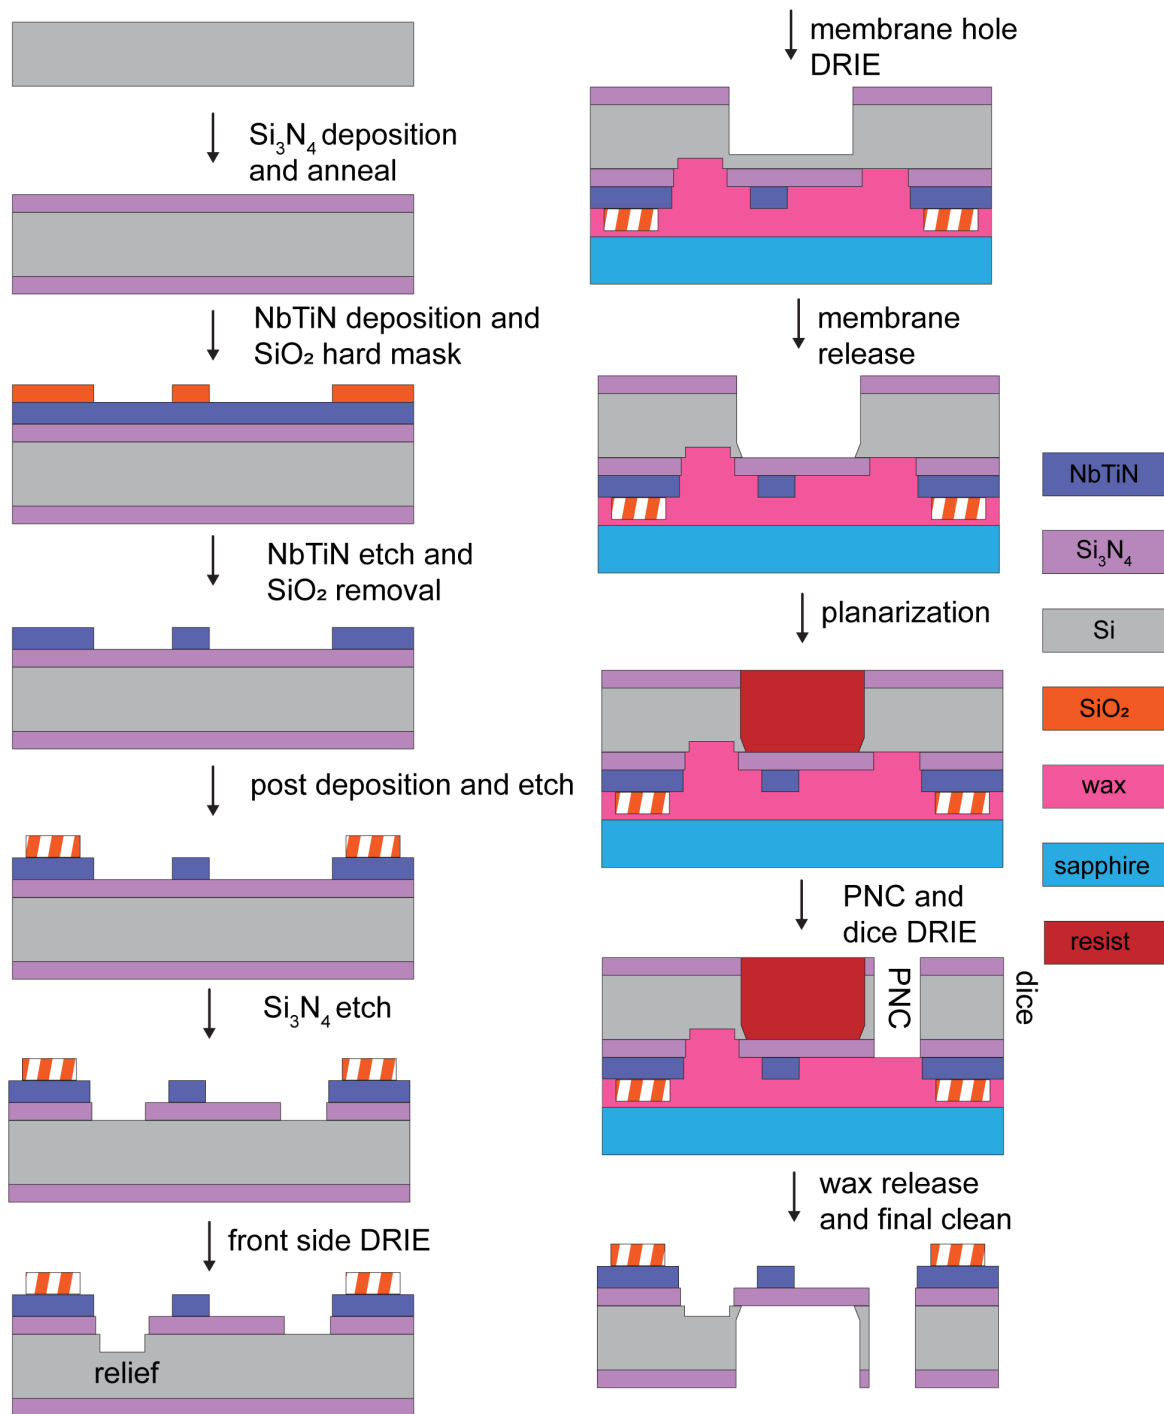


Figure A.2: **Membrane wafer fabrication flow.** Step-by-step description of process developed to fabricate wafer of membrane chips. Photolithography steps are left out for clarity, but are described in the text. The striped boxes are made of either SiO₂ or absent, depending on how whether the circuit chip and membrane chip are designed for bonding. When Al posts are used, the full post height is on the circuit wafer. Geometries are not drawn to scale.





Figure A.3: **Trion Fl etch non-uniformity** The Trion Fl "SiO₂" recipe struggles to clear the SiO₂ at the edges of the wafer. This photo is taken after the NbTiN etch, and we can see that NbTiN on the edge of the wafer due to the remaining SiO₂. Significant overetching when using the Trion Fl is required to clear the edge of the wafer.

in BOE without agitating the wafer in order to ensure small features do not get washed away. In some fabrication runs, the SiO₂ was etched in the Trion Fl etcher using the "SiO₂" recipe although this etch is known to be highly non-uniform (see Fig. A.3). In general, any etch process can be used for this step as long as it does not etch the NbTiN too quickly. Some overetching is tolerable, as the NbTiN that gets etched away was slated to be removed anyway.

After stripping the resist, the NbTiN is etched in a mixture of 5:2:1 H₂O:H₂O₂(30%):NH₄OH(30%) at 62 C for 20 s. This chemistry does not etch the Si₃N₄ or the SiO₂ hard-mask. The hard mask is then removed in BOE. Note, BOE also etches Si₃N₄, albeit very slowly at 1 nm/min compared >100 nm/min for SiO₂. Thus, BOE etches with exposed Si₃N₄ should be kept below 1 min etch time, and often the SiO₂ is cleared in the first 30 s.

For devices with Al posts, the full post height is on the circuit chip, and no posts need to be deposited on the membrane chip, although the designs should include locations where the Al posts can contact the circuit chip. If the devices are to be compatible with bonding, we deposit and pattern a thin 20 nm layer of SiO₂, again using the ECR, the 1 μm photoresist process, and BOE without agitation. Note, this choice limits the cleaning steps compatible with the membrane chip, as we can no longer BOE clean the chips after finishing the fabrication process.

Now, we pattern the Si₃N₄ using our 1 μm photoresist process and the Axic RIE with recipe

Figure A.4: **KOH etch angle**

"JM_nitride". This recipe uses 42 sccm of CF_4 and 2 sccm of O_2 at a chamber pressure of 60 mTorr with 150 W of RIE power and -360 V of DC bias. The design of the Si_3N_4 should be larger than the desired membrane size and include the locations where the posts will contact between the chips.

With either 3 or $7 \mu\text{m}$ photoresist, we use the DRIE to relieve the Si $\sim 10 \mu\text{m}$ away from the membrane and the post locations, akin to the front side relief on the circuit wafer. The front side of the wafer (side with all of the features up to this point) is waxed to a sapphire carrier wafer in the same manner described in the previous section. We now turn to releasing the Si_3N_4 membrane, which we perform in a two-step process. First, we etch most of the way through the Si using the DRIE, and then perform the final release with heated KOH, which hard stops of Si_3N_4 . This two-step process is faster than fully etching the hole and releasing the membrane with KOH. The pattern for the membrane hole DRIE will depend on how far we plan to etch into the Si before moving to KOH. To calculate the correct hole size, we note that KOH etches the 1-0-0 Si at 54.74° from the plane of the wafer as shown in Fig. A.4 [Sato et al. 1998]. Thus, if we want a membrane with side length $W \mu\text{m}$ and we plan to etch $(381 - y) \mu\text{m}$ into the Si with the DRIE, the side length D of the membrane hole in the pattern should be:

$$D = W + 2y \cot(55.74^\circ) \quad (\text{A.1})$$

With $7\mu\text{m}$ photoresist, we define the pattern for the membrane hole using this formula (remembering to use the "no touch" recipe for the pre-exposure bake). Importantly, the KOH etches Si along a specific crystal axis, so the designs need to be rotated 45° from the wafer flat in order to ensure the correct etch profile. As we etch this hole, we track the depth of the hole using an optical microscope. By comparing the microscope stage height that focuses on either the top surface of the wafer or the bottom of the etched hole, we can measure etch depth. In most of our fabrication runs, we aim to leave $20\mu\text{m}$ of Si (etch $360\mu\text{m}$ down). Some etch rate non-uniformity across the wafer will lead to non-uniformity in the membrane size, which in the future could be fixed by scaling the size of the patterned hole.

We spin-clean the wafer with acetone and IPA to remove the resist, as submersion in the solvents would lead to the wax dissolving. We heat up a beaker of water and a beaker of 30% KOH to 80°C . The wafer is heated in the water to avoid thermal shock when transferring to the heated KOH. The etch rate of 1-0-0 Si KOH at 80°C is $\sim 1\mu\text{m}/\text{min}$. With the membrane holes facing downwards and the wafer propped at an angle such that gas can escape the membrane holes, we etch the remaining Si, which takes 30-45 min to ensure complete membrane "release" (the membranes are still held fixed by the wax). To again avoid thermal shock, the wafer is slowly cooled by transferring back to the heated water and diluting with room temperature water after the etch.

The final step of the membrane wafer process is to etch the PNC and dice, which are performed at the same time using the DRIE. However, the wafer is now highly non-uniform. If we attempted to spin resist onto the wafer, the high-aspect ratio holes would not be filled, and thus the membranes would be etched by the DRIE before the PNC cutouts were defined. We experimented with several techniques to planarize the wafer, including filling the holes with wax or using roll-on resist tape. Ultimately, we developed a manual process which involves using a small paintbrush to fill the membrane holes with resist and pumping/slowly baking the resist to remove airbubbles. The process is as follows:

- Using a fine-tipped paintbrush (like those used to paint miniatures), fill holes with SPR-660
- Pump on resist using dessicator/bell jar for at least 5 min. This should help air bubbles

escape from the resist

- No contact ("proximity") bake at 95 C for 60 s to harden resist
- Inspect, and record with which holes are filled. Refer to [Fig. A.5](#) for visual indicators of whether the hole is appropriately filled.
- Repeat the previous four steps until all holes are filled.
- Using an acetone-soaked cleanroom "qtip", wipe away excess resist that is around the holes.
- Expose/develop a design using the MLA that is the inverse of the membrane holes to remove last bits of excess resist. As the resist is unknown, use a large dosage (at least 1.5-2 times the standard dosage). At this point, the membrane holes should be filled and the rest of the wafer clear of resist, a good proxy for "planarization"
- Spin a very thick layer 7 μm resist without P-20 (which dissolves the resist). Dispense a large puddle of resist, allow it to sit on the wafer for 10 s, spin slowly for a few secs, allow to rest again, and then spin at 1500 rpm for 40 s with a slow ramp-up. This should help achieve uniform coverage. The recipe "3I_SPR220_7_NOP20-1500-SLOWSPINUP_SM_SHAKE" at the resist bench implements this technique.
- Pump on resist using dessicator/bell jar for 10 min to remove air bubbles.
- Long-step down bake to slowly heat up the resist and again remove air bubbles. (recipe: SPR220-7 3IN N0 TOUCH SM)
- Expose/develop the PNC cutouts and dice pattern. Again, use a very high dosage as the resist layer is much thicker than 7 μm in most places.
- Inspect. If some holes are still not full, repeat the first four steps.

By removing the excess resist after the first few rounds of painting, we allow the PNC pattern to be defined by the MLA, rather than the excess resist left over from painting. The last round of painting is a final effort to save any membranes from the DRIE, but these membranes are unlikely to have a perfect PNC pattern. Using this process, we have successfully gotten 62 surviving membranes out of the 85 dies on the membrane wafer, much better than the 25% survival rate with the previous technique.

This final DRIE step etches away much more Si than the previous steps, so consequently it is the slowest. It is important to track the resist thickness as well, which can be done using the white light spectrophotometer from Filmetrics. In addition, the design should have resist protect any areas that do not need to be etched by the DRIE. This accelerates the Si etch rate and prevents the resist from getting etched before the Si is finished etching. The fabrication limits and uncertainties of the




Figure A.5: **Filling membrane hole with resist.** Examples of holes that are not-filled, partially-filled, and sufficiently filled with resist after painting the resist into the membrane hole using a fine-tipped paintbrush. The images are taken looking down on the membrane hole from the backside of the wafer. The resist around the holes will be removed before defining the PNC and dice pattern.




Figure A.6: **DRIE damage** (a) Example of a membrane etched through by the DRIE. This image is looking through the sapphire carrier and wax. (b) On the exposed side of the wafer, the resist pattern is faithfully transferred into the Si, even with long, narrow tethers (c) On the side of the wafer buried in the wax, the design is "blown-out" due to lateral etching of the Si. This limits the minimum feature size of the PNC.

PNC are also set by this final DRIE step. For example, we observe some lateral etching of the Si once the DRIE has finished etching vertically. This leads to a "blown-out" design on the side of the wafer that is buried in the wax, as shown in Fig. A.6. This leads to slightly different realizations of the PNC, which shifts the frequency of the substrate modes in an unpredictable manner.

At this point, the membrane wafer has been separated into individual membrane chips, but all of the chips are still waxed to the sapphire carrier. To carefully release the chips from the wax, we let the wafer soak in an acetone bath with the sapphire carrier side above the membrane chips. The bath is setup so once the chips release from the wax, they come to rest on a cleanroom wipe with the membranes facing upwards. We let the chips soak for at least one day, preferably over a weekend, and then pick the chips out of the acetone into a chip holder submerged in IPA. The chips undergo a solvent clean of 5 min in acetone and then another 5 min in fresh acetone and then into an IPA bath. These cleans are performed without sonication, as the membranes are fragile to vibrations.⁴ The chips are then put into two subsequent DI water baths to remove most of the solvents. Lastly, the chips are cleaned in heated nanostrip at 80 C for 2 min before being spun dry with DI water. If there is no SiO₂ on the membrane chips, they can also be cleaned in BOE for 1 min before the final nanostrip clean.

⁴Recent tests suggest the membranes can be sonicated at 40% power without breaking.

A.3 Transducer construction

There are two methods for transducer construction used in this thesis: glued or bonded. In this section, I will describe how the bonded circuit and mirror chips are assembled with the membrane chip to make the electromechanical etalon (every part of the transducer except for the curved mirror).

A.3.1 Glued construction

To assemble the chips together, we use a WestBond Manual Eutectic diebonder, which has dual heads shown in [Fig. A.7](#). One of the heads is a mini-vacuum chuck and the other is a pair of mechanical tweezers. It also has a work platform that can be moved up and down by means of a single thumb screw. Onto a two-axis stage from Thorlabs, we use double-sided kapton tape to hold the circuit chip in place. A glob of Stycast 2850 CAT 9 is placed on another piece of kapton tape nearby, such that it can be reached by the diebonder tool head. The circuit chip is held on one corner by the vacuum chuck tool. After centering the membrane chip in the microscope field of view, the remaining alignment is performed with the two-axis stage, not the diebonder vacuum head. Rotational alignment is achieved by aligning the bottom edge of both chips by rotating the two-axis stage. Then the micrometers on the two-axis stage are used to align the capacitor pads together by looking through the microscope objective. The work platform is raised slowly until the two chips are in contact. The vacuum will release the membrane chip if sufficient force is applied by raising the work platform.

Before beginning assembly, we grip a small piece of wirebonding wire with the mechanical tweezers, which will serve as an applicator for the stycast. After the membrane chip is placed on the circuit chip, we switch the diebonder head to the tweezers. The manual manipulator is then used to scoop a glob of epoxy onto the wire. Being careful not to contact the membrane chip, the epoxy is painted near the membrane chip and we allow capillary forces to pull the epoxy onto the edge of the membrane chip. Repeating in every desired location, the epoxy is then left to cure overnight. Generally, we apply the epoxy at the middle of the east and west parts of the membrane chip, where




Figure A.7: **Transducer construction** (a) Westbond diebonder tool used for transducer assembly. (b) Membrane chip is held by vacuum chuck on one corner. (c) Alignment is done by eye by trying to align the capacitor pads on the circuit and membrane chip, which can be seen in the center of this image.




Figure A.8: **Example of epoxied transducer.** This electromechanical chip is also wirebonded to a flexible PCB that mates with an SMA connector. The chip is held in place using more epoxy.

north denotes the side with the capacitor pad (see Fig. A.8).

The location of the epoxy and the location of the posts that define the chip spacing determine whether the thermal contraction of the epoxy will make the membrane bow up away from the circuit chip, or bow down towards the circuit chip. We can gain an intuitive picture by considering a simple one-dimensional model, where the post and epoxy blob are projected onto a line going from the center of the chip to one of the corners (blue dashed line in Fig. A.9). The posts are generally placed along this line. We model the post as infinitely rigid and expect the epoxy to shrink upon cooling down, which applies a downward force to the membrane chip at the location of the epoxy. In this projection, if the epoxy is farther (closer) to the chip center than the post, the center of the membrane chip will bow away (towards) the circuit chip. Thus, if the epoxy blobs are placed on the sides of the chip, if the posts are inside (outside) the black-dashed line, the center of the chip will bow up (down). Controlling the spacing is a key challenge for transducer construction, and moving to a bonded construction could increase reproducibility by removing the need for epoxy, which can change dramatically between cooldowns.

A.3.2 Bonded construction

To prepare the chips for bonding, we first ensure that both circuit and membrane chips have been recently cleaned in heated nanostrip to remove any organic contaminants. Then, we proceed

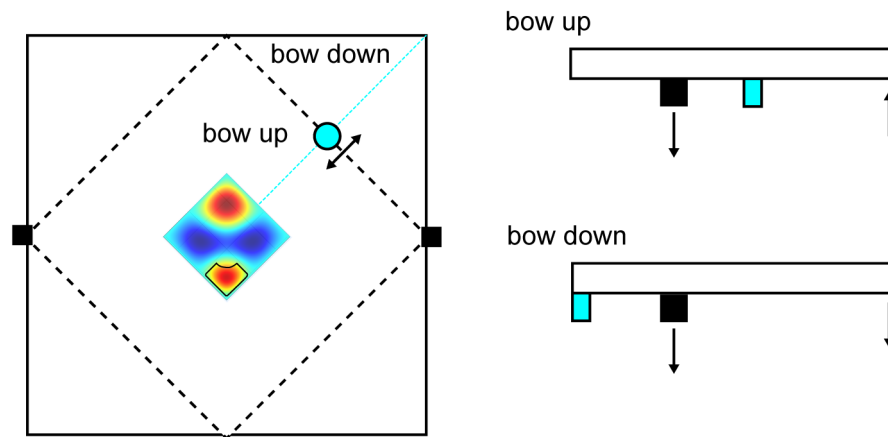


Figure A.9: **One-dimensional model of chip bowing.** The post (blue circle) and epoxy (black square) can be projected onto a line running from the center of the chip to the corner (blue dashed line). If the epoxy is farther away from the chip center than the post, the center of the chip will bow up. If the epoxy is closer to the chip center than the post, the center of the chip will bow down.




Figure A.10: **Applying force for circuit-membrane chip bonding.** Here, a specially machined piece of Al with four legs is used to push the chips together, although a clean Si chip works as well.

directly to hydrogen activation in the Ontos Atmospheric plasma, after which both chips are rinsed in DI water and spun dry. The setup and alignment procedure is identical to the glued construction. After placing the membrane chip on the circuit chip, we grip a third chip or a specially designed weight with the vacuum chuck and use it to apply a uniform force to the membrane chip, pushing the SiO_2 posts together (see Fig. A.10). The bonded chips are then annealed on a hotplate at 150 C for at least an hour. A good test of the bond quality is turning the bonded etalon upside down and seeing if the membrane chip falls off.

A.3.3 Determining gap separation

After assembling the chips, we attempt to measure the spacing of the capacitor pads to avoid cooling down collapsed electromechanical devices. Remarkably, this can be done with the help of a microscope and our color vision. The metal acts like a bad mirror, forming a low finesse optical cavity between the capacitor pads. The shortest wavelength of light that can fit into a cavity of length d is $\lambda = 2d$. Thus, if we are aiming for a capacitor gap spacing of 200 nm, the gap between the capacitor pads will act as a filter which picks out 400 nm light, or the color purple. If the spacing is closer to a $1\ \mu\text{m}$, the filter will not pick out any visible wavelengths, and metal on the membrane chip should look opaque. If the capacitor pads are in contact ("collapsed"), no such filter is formed between the capacitor pads. In practice, we find that the area around the capacitor pad (which

has a spacing equal to the combined thickness of the two capacitor pads 230 nm) appears blue. Examples are shown in [Fig. A.11](#). We make this technique more quantitative by measuring and fitting the reflectance spectrum between 400-1400 nm to a model of the material stack using the Filmetrics white light spectrophotometer [*Spectral Thin Film Reflectance Calculator for Thin-Film Stacks* 2024]. In a more general sense, this tool is a key technique for non-destructive measurements of film thicknesses, in particular measuring the thickness of photoresist during deep-reactive ion etching.




Figure A.11: **Capacitor gap spacing from color.** Using the color of the capacitor as seen with a microscope, we can determine if the capacitor pads are (a) collapsed, (b) correctly spaced, or (c) greatly bowed apart. We can make this technique more quantitative by fitting the reflectance spectrum between 400-1400 nm to a model of the material stack.

Appendix B

Numerical Fitting Techniques and Results

B.1 Resonator characterization

A key procedure employed throughout this thesis is fitting of the two-tone spectroscopy measurement to the S-parameters of a resonant circuit, measured either in reflection or "hanger" style. The transducer devices, and all flip chip designs in fact, are measured in reflection, whereas most of the planar devices are measured in "hanger" style, with several resonators inductively coupled to a shared transmission line. To mimic electromechanical cooling, a strong pump tone is applied at a fixed detuning and a weak probe tone with at most 1% the power of the pump tone is swept through frequency. We record the magnitude and phase of the probe tone at each frequency, which is then fit to the theoretical S-parameter to extract the resonator frequency, external coupling, and internal loss. For reflection and "hanger", those S-parameter expressions are:

$$S_{11}(\omega) = \frac{(\kappa_{\text{int,e}} - \kappa_{\text{ext,e}}) - 2i(\omega - \omega_e)}{(\kappa_{\text{int,e}} + \kappa_{\text{ext,e}}) - 2i(\omega - \omega_e)} \quad (\text{B.1})$$

$$S_{21}(\omega) = \frac{\kappa_{\text{int,e}} - 2i(\omega - \omega_e)}{(\kappa_{\text{int,e}} + \kappa_{\text{ext,e}}) - 2i(\omega - \omega_e)}. \quad (\text{B.2})$$

In its simplest form, this procedure is fragile to Fano interference [Khalil et al. 2012; Probst et al. 2015; Baity et al. 2024; Rieger et al. 2023], which arises from unintended impedance mismatches that cause partial reflections of the probe tone or finite directivities of directional components such as directional couplers or circulators. The resulting interference can lead to asymmetric or otherwise distorted lineshapes. A good technical note with an interactive python notebook was developed




Figure B.1: **Example of numerical fit.** The data and fit for as-deposited device A at 10 mK are shown for a power of $\bar{n} \sim 1$, in good agreement with Eq. B.3. The complex response is shown on the left, and the equivalent magnitude and phase response is on the right.

by the Steele lab to gain intuition into this effect [*SteeleLab / Technical Note Fano · GitLab 2022*]. The "diameter-correction method", introduced by [Khalil et al. 2012], circumvents this effect by analyzing the S-parameter measurements in real and imaginary space as opposed to magnitude and phase space. In this picture, the interference amounts to a rotation of the resonator response about some arbitrary point in phase space, which can be corrected. We point out that this correction is only accurate for a hanger measurement geometry, and not pure reflection, so the material properties extracted from circuits measured in hanger are more trustworthy.

We first pre-process the frequency-dependent data by normalizing the off-resonance values to unity transmission. We then fit the response to

$$S_{21}(\omega) = 1 - \frac{Q|Q_e^{-1}|e^{i\phi}}{1 + 2iQ\frac{\omega - \omega_e}{\omega_e}} \quad (\text{B.3})$$

to extract the quality factor of the resonator Q , the external quality factor Q_e , and center frequency ω_e . We use a complex external quality factor $Q_e^{-1} = |Q_e^{-1}|e^{i\phi}$ to account for unintended reflections that lead to asymmetry in the frequency response and define the internal quality factor as $Q_i^{-1} = Q^{-1} - \Re\{Q_e^{-1}\}$. An example fit is shown in Fig. B.1.

Another technique that increases the speed and accuracy of resonator characterization and data

acquisition is a homophasal point distribution $\{f_n\}_{n \in [-(N-1)/2, (N-1)/2]}$ of the sample frequencies, introduced by [Ganjam et al. 2024; Baity et al. 2024]. Standard vector network analyzer (VNA) measurements sample points linearly in frequency space which when projected onto the real and imaginary plane, leads to a high point density away from the circuits resonance frequency and a low point density on resonance. One can design a nonlinear frequency sweep, where the frequency spacing of the sample points increases quadratically with detuning from the expected resonance frequency:

$$f_n = f_r + \frac{\Delta F}{2} \frac{\tan\left(\frac{n}{N-1}\Delta\theta\right)}{\tan(\Delta\theta/2)} \quad (\text{B.4})$$

with $\Delta\theta = \arctan \frac{2W}{1-W^2}$. In this expression $W = \Delta F/\kappa_e$ where ΔF is the measurement span centered on a resonator with frequency f_r and linewidth κ_e . This point distribution is evenly spaced in phase, and results in a more efficient frequency sampling scheme which requires fewer overall points to achieve a similar fit uncertainty. Only approximate knowledge of the resonator frequency and linewidth are required for this technique, and in the limit $W \rightarrow 0$, we recover the standard linear sweep. This nonlinear point distribution can be implemented on a VNA using the "segmented sweep" option, which allows the user to define an arbitrary set of sample frequencies. There is often a minimum frequency spacing between consecutive points set by hardware limitations, so for very narrow resonances (such as the lightly damping mechanical oscillator), we must revert to a linear frequency sweep.

B.2 Tables of fit parameters

For the fits to the TLS model, we fit to the expressions derived in [chapter 5](#), reproduced here with some small changes to reflect the fitting routine. Each resonator design A-E has an independent background term, but $\tan \delta_{0,\text{res}}$, n_c , a background TLS term $F \tan \delta_{\text{other}}$, the pre-factor for the full expression for relaxation damping B , the pre-factor A_2 for the expression of $T_{1,\text{min}}$, and the self-heating pre-factor A are fit globally for each set of resonators.

	Device	$\omega_e/2\pi$ (GHz)	$Q_i^{-1} \times 10^5$	$ Q_e^{-1} \times 10^5$
as-deposited	A	5.968	18.4	9.3
	B	6.133	16.2	10.3
	C	6.289	12.5	22.4
	D	6.384	7.5	9.8
	E	6.480	1.9	15.3
annealed	A	5.959	2.4	9.8
	B	6.103	2.2	12.4
	C	6.271	1.9	23.1
	D	6.362	1.3	8.7
	E	6.443	0.8	16.6
as-deposited transducer [Brubaker et al. 2022]	—	7.938	5.56 at $T_{\text{bp}} = 40$ mK	17.9

Table B.1: **Fitted resonator parameters.** The center frequency, internal quality factor, and external quality factor for the studied devices measured at $T_{\text{bp}} = 10$ mK and $\bar{n}_e \sim 1$, except where otherwise noted. For all resonators, the external quality factor is roughly within an order of magnitude of the internal loss, which allows for accurate determination of the relative coupling rates from the numerical fits.

$$Q_{\text{res}}^{-1}(T) = \frac{F_{\text{SiN}} \tan \delta_{0,\text{res}} + F \tan \delta_{\text{other}}}{\sqrt{1 + \bar{n}_e/n_c}} \tanh \left(\frac{\hbar \omega_e}{2k_{\text{B}}T} \right) \quad (\text{B.5})$$

$$Q_{\text{rel}}^{-1} = \frac{F_{\text{SiN}} B \omega_e^2}{2T} \int_0^{1000T} dE \int_{T_{1,\min}(E)}^{\infty} dT_1 \text{sech}^2 \left(\frac{E}{2T} \right) \sqrt{1 - \frac{T_{1,\min}(E)}{T_1}} \frac{1}{1 + \omega_e^2 T_1^2} \quad (\text{B.6})$$

$$T_{1,\min}^{-1}(E) = A_2 E^2 \coth(E/2T) \quad (\text{B.7})$$

$$T_{\text{eff}}(\bar{n}_e) = A\sqrt{\bar{n}_e} + T_{\text{bp}} \quad (\text{B.8})$$

Note the extra factor of ω_e in the expression for relaxation damping. This will just rescale the values extracted for the pre-factor B . The Boltzmann constant is also wrapped into A_2 as we redefined $E k_{\text{B}} \rightarrow E$. We cutoff the integral for the relaxation damping term at $E = 1000T$ in order to increase the speed of our fit routine. This should still sample all relevant TLS energies. As a final technical note, we fit logarithm (base 10) of many of these parameters such that Matlab converges more easily.

device description	Q_{bg}^{-1}	$\tan \delta_{0,\text{res}}$	$F \tan \delta_{\text{other}}$	n_{c}
as-deposited planar	$Q_{\text{bg},\text{A}}^{-1} = 1.33 \times 10^{-4}$ $Q_{\text{bg},\text{B}}^{-1} = 1.16 \times 10^{-4}$ $Q_{\text{bg},\text{C}}^{-1} = 8.28 \times 10^{-5}$ $Q_{\text{bg},\text{D}}^{-1} = 4.86 \times 10^{-5}$ $Q_{\text{bg},\text{E}}^{-1} = 1.38 \times 10^{-5}$	1.37×10^{-3}	6.11×10^{-6}	12.85
annealed planar	$Q_{\text{bg},\text{A}}^{-1} = 2.08 \times 10^{-6}$ $Q_{\text{bg},\text{B}}^{-1} = 2.49 \times 10^{-6}$ $Q_{\text{bg},\text{C}}^{-1} = 3.38 \times 10^{-6}$ $Q_{\text{bg},\text{D}}^{-1} = 2.67 \times 10^{-6}$ $Q_{\text{bg},\text{E}}^{-1} = 2.62 \times 10^{-6}$	4.82×10^{-4}	4.41×10^{-6}	23.1
as-deposited transducer [Brubaker et al. 2022]	2.88×10^{-5}	$F \tan \delta_{\text{tot}} = 3.19 \times 10^{-5}$	—	1.48×10^3

device description	B	A_2	A
as-deposited planar	1.11×10^{-12}	2.19×10^{-10}	7.97×10^{-4}
annealed planar	5.05×10^{-16}	3.12×10^{-11}	—
as-deposited transducer [Brubaker et al. 2022]	3.78×10^{-10}	2.83×10^{-5}	1.45×10^{-4}

Table B.2: **TLS model parameters.** Fitted model parameters used throughout this thesis. The annealed planar devices did not have any self-heating, so no coefficient is reported. The F_{SiN} of the transducer is unknown due to misalignment of the capacitor pads, so we just report the scale of the resonant damping from all TLS sources.