

**Quantum optomechanics with a stable Fabry-Pérot cavity
in a microwave-to-optical transducer**

by

Maxwell D. Urmey

B.A., Amherst College, 2009

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:

Cindy Regal, Chair

Konrad Lehnert

John Teufel

Tasshi Dennis

Joshua Combes

Urney, Maxwell D. (Ph.D., Physics)

Quantum optomechanics with a stable Fabry-Pérot cavity in a microwave-to-optical transducer

Thesis directed by Prof. Cindy Regal

Continued development of complementary capabilities in two quantum technologies—information processing with microwave-frequency superconducting circuits and networks based on optical-frequency photonic links—have motivated an effort to bridge the frequency gap between them with a transducer capable of preserving quantum signals. At the same time, mechanical resonators have been engineered to interact with electromagnetic waves across the spectrum in a variety of tailorable physical implementations. This thesis describes advances in a transducer using the mechanical mode of a silicon nitride membrane as an intermediary permitting coupling between a superconducting resonant LC circuit and a Fabry-Pérot optical cavity. With this choice of architecture, we have developed transducers of unparalleled efficiency and the unique capability to run continuously without appreciably degrading the performance of the superconductor from scattered optical photons. We have been able to laser-cool the mechanical mode to an occupation of less than one photon, and achieved transducer performance capable of converting a single microwave photon to optical frequencies with a signal-to-noise ratio of $1/3$. Furthermore, operating the transducer does not cause appreciable heating of a superconducting transmon qubit linked to its microwave input port. With progress towards a quantum-enabled interface, we have started to develop capabilities towards implementing protocols that optically verify the quantum performance of the transducer.

Dedication

To my parents Tammy and Bill, and Steffi, for all of your support.

And to Oslo, for the kind of support that comes from gently laying fish skins on my feet.

Acknowledgements

I have many people to thank for their help on contributing to the work described here, and, more importantly, to my training as a scientist. I'm extremely fortunate to have worked on such a collaborative project with such wonderful teammates. The success of our experiments has demanded it, but I know that wasn't the primary motivation behind the generosity and diligence of everyone I've worked with.

First of all I would like to thank Cindy. I'm so grateful that you gave me the opportunity to not only work in your lab, but to work with you and learn from you. I have learned so much from your clarity of thought, experimental intuition, and seemingly innate knowledge of which problems are the ones most worthy of spending time on. I can't think of someone better to model how to be a member of the scientific community, nor can I think of an advisor with as much dedication and commitment: thanks for zooming with me on a Christmas Eve for final considerations before submitting a paper, for sending me revisions on Thanksgiving, and for hand-delivering my poster to me in Trieste.

Konrad, I am very fortunate to have signed on to a project with you as a second advisor. I also owe you thanks for that opportunity. I appreciate your patience in exploring every idea with thoroughness, your excitement in teaching a new concept, and your sense of humor. You ingrained many of the questions I now ask myself whenever I look at the results of any measurement, or to address when presenting results to others.

Thank you Andrew and Pete for being a welcoming team to join when I began working on the project. Nir, thank you for giving me a crash course on working with optics, and you Bob for

teaching me the secrets and lore of making optomechanical cavities.

I've been really fortunate to work with and learn from extremely talented post docs. Andrew, you've been a role-model to aspire to in many ways: on how to move fast and how to lead with excitement and generosity. Ben, thank you for being such a patient and unyielding thinker, for your determination that we could do this, and for great musical taste and keeping a (the?) chestnut on-hand. Jon, thank you for being so practical and so quietly capable, for introducing me to Sludgefest, and for getting angry on my behalf when needed. And Nick thanks for being a natural leader, an excellent and excited scientist, and a great friend. Bonus points to each of you for being extremely fun to work with.

The microwave to optics team has been a constantly-changing, usually-growing, consistently-great set of people. Rob, we came to JILA at the same time and same age, but from our first year you already seemed like a fully-formed scientist. Though we started out spending more time in the command center in Duane, I'm really glad we got to work more closely together in the lab in the later portion of your time here. I really appreciate your teaching me the fundamentals about superconducting qubit measurement. Sarang, it's been a pleasure to work closely with you throughout our time here, thanks for teaching me Comsol, and for the inspiration roll, which I still need to cash in on.

I'm already half a page over the limit for this section, so I'm going to have to speed this up: Kazemi, thanks for your thoughtfulness, on matters scientific and otherwise, for the doggy time, the musical selection, and the deep dives on graphic design. Luca, it was fun iterating on coupling optics and then banging out the Frankenboard very quickly once we realized we needed a maxi-design. I'm looking forward to what we do next. Sheng-Xiang thanks for teaching me about the intricacies of qubit fabrication, and for playing with Oslo. Sarah, it's really been a pleasure getting to work with such a natural team-player over the last year, and seeing you quickly take more on. Jake thanks for giving a much clearer window into your side of the collaboration. It's been really great to work with you and get to know you.

To everyone in the instrument shop, I always have too much fun up there: Kim, Hans, Kyle,

Adam, Calvin, and James I can distinctly remember asking each of you to do something borderline-unphysical for me. And, sometimes after laughing in my face, you patiently talk about what I need, bring in some new considerations, try to sell me on some more reasonable alternative, and then make exactly what I need. Terry, thank you for teaching me control theory early on and later about pemmican, and I'm still amazed at how quickly you turned out the impedance matching circuit board for the SAW filter described in Sec. 6.2.3, when we needed it immediately.

Thank you to the rest of the Regal group, especially Chris, Ravid, Ting Wei, Gabriel, Dylan, Sanjay, Remi, and Brian. And thanks to the Lehnert group, Eric, Dan, Ben, Lucas, Ceasar, Max, Jeremie, Kelly, Christian, Brad, Joyce, Brendon, Pablo, Alec, Liz, Ziyi, Kyle

Steffi, I can't list all that I want to thank you for. Your support, encouragement, and commitment have been so important to me. You're a great teammate.

Thanks to Mohammad Alhejji, Wes Brand, Jose Valencia, Drew Morrill and Kyle Matsuda, and Korn and Naomi, Mike and Laura, Kaela and Seth, Adam and Annalise for all the great times over the last several years.

I'd like to especially thank Grover, Peter, Monty and Ben for always being there, and for some really important conversations.

Contents

Chapter

1	Introduction	1
1.1	Current state of microwave-optical transduction	2
1.2	Thesis overview	4
2	Transducer requirements in a quantum network	8
2.1	Efficiency and added noise	8
2.2	Networking rate: bandwidth and duty cycle	9
2.3	Comparison between several transducer platforms	11
2.3.1	Transduction via low-frequency mechanical mode	11
2.3.2	Microwave-frequency resonant piezo-optomechanical transducers	14
2.3.3	Direct electro-optic transduction with nonlinear materials	15
2.3.4	Millimeter-wave-to-optical transduction with trapped atoms	16
3	Primitive networking demonstrations	18
3.1	Entanglement protocols	19
3.1.1	Direct state transfer	19
3.1.2	Singly-heralded protocols with weak excitation	20
3.1.3	Doubly-heralded protocol	24
3.2	Single-node demonstrations	25
3.2.1	Intensity correlation measurement with a deterministic quantum source . . .	26

3.2.2	Dual-rail interference using a frequency-bin encoding	32
3.2.3	Testing single-optical-photon counting without a qubit source.	36
3.3	Nonclassical intensity and higher-order correlations with pure optomechanical system	42
4	Fabry-Pérot cavities in membrane-based transducers	48
4.1	Operating theory of a mechanical transducer	48
4.1.1	The optomechanical and electromechanical interactions	48
4.1.2	Efficient transduction through a mechanical mode	50
4.1.3	Optomechanical sideband cooling	52
4.1.4	Optical output spectrum	54
4.1.5	Effects of technical noise	56
4.2	Transducer design	59
4.2.1	Transducer architecture	59
4.2.2	Membrane misalignment	61
4.2.3	Selecting values for mirror transmissions	69
4.2.4	Fully fixed optical cavity	72
4.2.5	Thermal environment of the mechanical mode	74
4.2.6	Membrane collapse and intentional pinning	78
5	Optomechanical ground-state cooling in a transducer	81
5.1	Transducer details	82
5.1.1	Mechanical isolation with phononic shield	82
5.2	Optomechanical ground-state cooling	85
5.3	Transducer characterization	86
5.4	Effects of Laser Illumination	90
5.5	Experimental setup	92
5.6	Details of transducer characterization	98
5.6.1	Microwave circuit characterization	98

5.6.2	Optical cavity characterization	98
5.6.3	Optomechanical and electromechanical characterization	101
5.6.4	Characterization of other parameters	103
5.7	Analysis Details	105
5.7.1	Transducer efficiency	105
5.7.2	Calibration of spectra and temperature sweeps	106
5.7.3	Ground-state cooling measurements	110
5.7.4	Added noise measurements	115
5.8	Microwave pump effects on superconducting circuit	116
5.9	Dispersive readout of a transmon qubit using optical heterodyne detection via the transducer	118
6	Subsequent improvements towards a quantum-enabled transducer	122
6.1	Optomechanics in a cryogen-free dilution refrigerator	122
6.1.1	Coupling optics	123
6.1.2	Vibrations from the pulse-tube	125
6.1.3	Challenges working with the helium battery	126
6.1.4	Spring-mounted curved cavity mirror	127
6.1.5	Clamping the cold lenses	129
6.1.6	Optomechanical ground-state cooling in a cryogen-free dilution refrigerator	129
6.1.7	Mechanical ringdown with stroboscopic optical readout	131
6.1.8	Inferring the mechanical environmental temperature	131
6.2	Noise considerations for future experiments	133
6.2.1	Laser technical noise	133
6.2.2	Increasing filter cavity lock bandwidth	134
6.2.3	Additional filter cavity	135
6.2.4	Leakage of the coherent pump	138

6.2.5	Thermal motion of cavity mirrors	138
6.2.6	Phononically engineering mirror substrates	143
References		146
 Appendix		
A	Details of optical cavity assembly	161
A.1	Mounting the curved mirror	161
A.2	Cavity assembly setup	163
A.3	Cavity assembly procedure	165
A.4	Notes on measuring cavity linewidth in the assembly setup	165
B	Details of operating the optical detection setup	167
B.1	Detail on the optical measurement setup	167
B.2	Isolating the polarizations between the damp and lock	171
B.3	Procedure for locating the cavity mode with the pump and lock beams	173
B.4	Procedure for optimizing modematching	174
B.5	Impact of modematching on reflection measurements made with optical heterodyne detection	176
B.6	Transducer cavity lock loop	177

Tables

Table

2.1	Comparison of transducer platforms	12
3.1	Parameters for calculation of $g^{(2)}$ and T_{in} when sweeping coherent microwave input.	42
5.1	Electro-optic transducer parameters	99
5.2	Contributions to η_q in qubit readout experiment.	121
6.1	Parameters used for ground-state cooling model in Fig. 6.4	132
6.2	Parameters for estimation of noise required for $\text{SNR} = 1$	134

Figures

Figure

2.1	Transduction in upconversion and downconversion.	9
2.2	Efficiency and added noise.	10
2.3	Bandwidth and duty cycle.	11
3.1	Direct state transfer.	18
3.2	Heralded entanglement generation with the DLCZ protocol.	20
3.3	Doubly heralded entanglement protocol	23
3.4	Intensity correlation measurement setup	28
3.5	Measurement of $g^{(2)}$ with a single-photon source and coherent and thermal noise contributions	31
3.6	Generic dual-rail interference experiment	33
3.7	Time- and frequency-bin interference experiments	35
3.8	Methods for calibrating p_N	38
3.9	Single-photon counting with a coherent microwave source.	40
3.10	Optomechanical intensity correlation experiments	43
4.1	Principles of transducer operation	49
4.2	Technical pump noise	58
4.3	Transducer architecture	60
4.4	Stable cavity geometries	61

4.5	Optomechanical cavity geometry	62
4.6	Input mirror choice	65
4.7	Optical cavity architecture improvements	67
4.8	Locking beam backaction	71
4.9	Fully fixed optical cavity	73
4.10	Thermal modeling of membrane heating from laser absorption	76
4.11	Evidence of a collapsed membrane	80
5.1	Mechanical design and characterization	83
5.2	Ground-state cooling of electro-optomechanical transducer.	84
5.3	Minimizing input-referred added noise.	89
5.4	Effects of laser illumination on microwave circuit	91
5.5	Experimental setup	93
5.6	Cavity birefringence and beam detunings.	95
5.7	Photographs of transducer mounted in cryostat	97
5.8	Rendition of transducer assembly	97
5.9	Losses and modematchings in efficiency calibration	106
5.10	Temperature sweeps	107
5.11	Analysis of optomechanical spectra	111
5.12	Microwave pump effects on superconducting circuit	116
5.13	Setup to perform dispersive readout via electro-optic transduction and optical heterodyne detection.	119
5.14	Dispersive readout measurement results	120
5.15	Using qubit readout to calibrate the total quantum efficiency.	120
6.1	Dilution refrigerator measurement setup photographs	124
6.2	Spring-mounted mirror design	128
6.3	Clamping the cold lenses.	129

6.4	Ground state cooling in a cryogen-free dilution refrigerator	130
6.5	Decreasing laser technical noise	136
6.6	Optical spectra showing thermomechanical mirror motion	140
6.7	Phononically engineered mirror substrates.	144
A.1	Epoxying the curved mirror	162
A.2	Cavity assembly setup	164
B.1	Detail on optical measurement setup	168
B.2	Jones matrices describing the desired polarization configuration of the experiment. .	172
B.3	Simplified transducer cavity lock closed loop response	178

Chapter 1

Introduction

A transducer capable of linking microwave and optical frequencies with sufficiently low noise and high efficiency would directly port our exceptional control over quantum states using superconducting circuits to the more robust energy scales of optical photons. The primary motivation behind this effort is the networking of remote superconducting quantum processors, a platform with exciting recent advances in engineering favorable error rate hierarchies towards fault tolerant quantum computing [Teoh et al. 2023] and an exciting demonstration of favorable scaling of logical qubit performance with surface code size [Acharya et al. 2023]. Because of the microwave operating frequencies of superconducting quantum circuits, their signals are isolated to the mK environment at the base plate of a dilution refrigerator. Quantum circuits were recently used to perform a loophole-free measurement violating Bell’s inequality, with an impressive 30 m length of cryogenic transmission line [Storz et al. 2023]. However, longer networks of superconducting quantum circuits would require promotion of the information to optical photons. Upconverting deterministically created microwave photons to optical frequencies could furthermore be an additional tool in quantum key distribution protocols [Bennett and Brassard 1984] (which are instead currently implemented with weak coherent optical sources over long distances [Hwang 2003; Liao et al. 2018]), and could enhance the baseline of photon-starved astronomical telescopes [Gottesman, Jennewein, and Croke 2012].

The difference in energy scales that makes microwave-to-optical transduction useful also renders it quite technically challenging to pull off. The promotion of quantum information from

microwave frequencies to optical frequencies requires additional energy, which is supplied by an optical pump. However, even a single optical photon has enough energy to disrupt superconductivity, and any stray pump photons absorbed by the transducer can add noise to the transduced signal or dephase the superconducting qubit to which it is linked. An ideal transducer would simultaneously permit strong enough coupling between the microwave and optical excitations, while keeping the optical mode far enough from the superconductor.

1.1 Current state of microwave-optical transduction

There has been exciting work on a variety of platforms capable of linking microwave and optical frequencies, including nonlinear materials [Ilchenko et al. 2003; Strekalov et al. 2009; Xiong, Pernice, and Tang 2012; Fan et al. 2018], microwave-frequency optomechanical crystals [Bochmann et al. 2013; Balram et al. 2016], membrane-optomechanical systems [Regal and Lehnert 2011; Andrews et al. 2014; Planz et al. 2022], neutral atoms [Verdú et al. 2009; Hafezi et al. 2012; Covey, Sipahigil, and Saffman 2019; Vogt et al. 2019; Kumar et al. 2023], emitters embedded in solids [Imamoğlu 2009; Marcos et al. 2010; Kubo et al. 2010; Williamson, Chen, and Longdell 2014; Bartholomew et al. 2020], and magnons [Hisatomi et al. 2016].

I'll further detail three of the above platforms that have emerged as particularly successful recently. The first couples an optical whispering gallery mode of a lithium niobate disk to the mode of a three dimensional microwave cavity [Rueda et al. 2016; Hease et al. 2020]. The second platform uses the microwave-frequency mechanical resonance in an optomechanical crystal (OMC) [Eichenfield et al. 2009; Chan et al. 2011]. Transducers based around this technology channel microwave excitations along a transmission line, and piezoelectrically transduce them to phonons resonant with the OMC. A laser pump then supplies the energy to convert that microwave-frequency mechanical excitation to an optical photon via the optomechanical interaction [Bochmann et al. 2013; Balram et al. 2016]. These first two platform permit MHz-scale transduction bandwidth, but the guided optical modes are susceptible to absorption, and typically require pulsing the operation of the transducer, in practice a reduction of the duty cycle by many orders of magnitude. By pulsing

fast enough, both of these platforms have reduced the noise below one input-referred photon, a necessary threshold for quantum networking protocols [Mirhosseini et al. 2020; Sahu et al. 2022; Meesala et al. 2023].

The third platform also uses a mechanical intermediate mode, though with a much lower resonant frequency [Regal and Lehnert 2011; Andrews et al. 2014; Planz et al. 2022]. This platform has been the focus of my experimental efforts, and is the one detailed in this thesis. The lower mode frequency allows for a much larger mechanical oscillator, and greater separation between the optical mode and the superconducting circuit, and therefore potentially reduced sensitivity to the optical pump. Because of the lower mode frequency, an additional microwave-frequency pump is needed. At the time of writing, the effects of the strong microwave pump are the primary limitation preventing membrane-based transducers from achieving input-referred noise of less than one photon, and preclude any quantum demonstration.

The transducer developed by the Regal and Lehnert groups is based around a MHz-scale vibrational mode of a highly-tensioned silicon nitride membrane [Andrews et al. 2014]. The choice to use a membrane is compatible with the optical mode being defined by a Fabry-Pérot cavity, in which the optical field propagates primarily in vacuum, rather than through a dielectric as a guided mode. Propagation through vacuum along with the excellent mode control afforded by such a highly developed optical technology allows transducers using Fabry-Pérot cavities to greatly reduce the optical scattering toward the superconductor, and indirect heating of the circuit due to optical absorption, relative to the more integrated alternative platforms. Furthermore, membrane-based transducers benefit from the improvements that the membrane optomechanical community has developed since the platform was first pioneered [Thompson et al. 2008]. These include improvements to the mechanical loss rates due to material improvements [Wilson et al. 2009], from engineering an acoustic band-gap in the support structure [Yu et al. 2014; Tsaturyan et al. 2014] in order to suppress radiation loss [Wilson-Rae et al. 2007], and from reducing intrinsic loss rates due to material bending by patterning the membrane itself [Tsaturyan et al. 2017; Ghadimi et al. 2018; Reetz et al. 2019].

1.2 Thesis overview

During the course of my tenure working on this collaborative project to develop a quantum microwave-to-optical transducer, we have developed an increasingly nuanced understanding of the performance of the transducer as we refine its capability. We have progressed from focusing on the narrowband noise performance at the maximum transducer efficiency with matched optomechanical and electromechanical rates, to an understanding of several, often competing, noise sources to optimize relative to one another. As we have reduced the noise of our transducer towards quantum performance, transduced pulses encoding the state of a superconducting qubit, and brought demonstrations of verifying a quantum signal within reach, we have considered more fully the time and frequency dependence of the signal to be transduced. As we look toward demonstrations that use optical single-photon-detection, we have also considered additional, and increasingly feeble, sources of noise that will contribute substantially to these more sensitive experiments.

More concretely, in my time working on this project, we have developed the ability to make increasingly stable and reliable cryogenic Fabry-Pérot cavities. My responsibilities have centered on the performance of the optical cavity and the optical measurement setup. When I first joined the project, I learned from Bob Peterson to construct the high-finesse optomechanical cavities we employ in the transducer, by imaging the membrane within the optical cavity, aligning it by sweeping the cavity length and minimizing its perturbation of the cavity modes, and epoxying it in place. I also learned to measure the first transducer devices that we operated at dilution refrigerator temperatures from Andrew Higginbotham. In characterizing the cryogenic performance of the optical cavity I had assembled, I built up a picture of the device being operated as well as what aspects of the cavity needed improvement. The culmination of these efforts in assembly and measurement was a transducer of unprecedentedly low added noise, and sufficient efficiency to allowed us to observe strong correlations in the thermal noise emanating from the microwave and optical ports of the transducer, which we realized could be leveraged to remove noise by monitoring the reflection from the input port and feeding it forward [Higginbotham et al. 2018].

Despite improved performance of the transducer, superconducting circuit noise due to both the microwave pump and the optical pump prohibited operation without tens of noise photons that would overwhelm any single-quantum signal.

We then worked to improve the transducer performance towards quantum operation. For the design of the next-generation device, I evaluated the mirror transmission asymmetry of the optical cavity make the transducer more tolerant of membrane misalignment. Iterating between design and measurement of new devices allowed the operation of the optical cavity to move from a mode of exhaustive searching for a desirable operating configuration, to one of fuller understanding that allowed me to converge more quickly and predictably on an optical resonance with cavity properties and optomechanical coupling most suitable for the desired measurement. Simultaneously, efforts in the clean room within our team enabled chemical bonding of the membrane to one cavity mirror, for cavity construction that dramatically reduced the misalignment of the membrane upon cooling the devices to cryogenic temperatures (see Oliver Wipfli’s thesis for early work on this concept [Wipfli 2015] and Pete Burns’s thesis for a description of the chemical bonding process [Burns 2019]). With the membrane affixed to one mirror, we could no longer tune the spacing between the membrane and that mirror using piezoelectric actuation to control the optomechanical coupling strength, and instead tuned the wavelength of an external cavity diode laser, which we incorporated into the setup while ensuring it met our strict demands on laser stability and noise. With the mirror coating now deposited on a silicon chip, we had to understand the consequences of the silicon’s absorption of our 1 μm laser source, both the effect on the optical cavity assembly and alignment process at room temperature, as well as the absorption’s impact on device thermalization during operation at cryogenic temperatures.

With improved devices we needed to continually keep additional sources of noise were revealed at bay. Our optical cavity performance became limited by the noise introduced through the leads of the piezo actuation of the second cavity mirror, and we began operating transducer cavities without any tunability. I modified our laser servo to account for this change, and designed cavity geometries that permitted a distribution of optomechanical couplings within the tuning range of our laser. A

significant improvement in the design of the mechanical device was Sarang Mittal’s integration of a single-unit-cell phononic filter in the silicon chip suspending the membrane to better isolate it from its thermal environment, reducing the power requirement of the optical pump in cooling the mechanical mode.

With these improvements, we were able to optomechanically cool the mechanical mode to close to its quantum ground state [Brubaker et al. 2022], dramatically reduce the noise on the superconducting circuit due to the laser by approximately two orders of magnitude, and reduce the total transduction noise by a factor of ten, bringing our transducer to the cusp of quantum performance. While characterizing the device in this experiment, we began to leverage the full capability of a device with both microwave and optical connectivity, for example, using the superior measurement efficiency on the optical side for sensitive noise measurements, while using the inherent stability of the microwave signals and circuit to precisely calibrate the loss rate of our membrane’s mechanical mode. Transmission measurements from the microwave port to the optical port helped us to calibrate the optical cavity response and beam detunings with unprecedented ease. With the same device, we also demonstrated compatibility with a superconducting qubit, reading out its quantum state classically via the transducer [Delaney et al. 2022] (see also Rob Delaney’s thesis [Delaney 2022]). Remarkably, we could operate the transducer continuously without appreciably increasing the decoherence rate of the qubit, and the optical pump didn’t measurably affect the qubit performance.

The primary impediment to quantum-enabled transduction was additional loss and noise from the strongly-driven microwave circuit. Our team has recently investigated the source of this microwave pump-induced loss [Mittal et al. 2023], as will be described in Sarang Mittal’s thesis.

At the same time, we are preparing for a next set of measurements and building up capability to use single-photon detection with our transducer to demonstrate quantum performance of the transducer, and allow measurement-based quantum control over the optomechanical system. The details of that detection setup will be described in Luca Talamo’s thesis. With the new capability afforded by single-photon detection apparatus comes more stringent requirements on the noise

of our laser, which Luca, Sarah Dickson, and I have worked to characterize and further reduce. Whereas all of our previous transducer measurements were done in the low-vibration environment of a dilution refrigerator precooled by a liquid helium reservoir, we have also built up the ability to perform our sensitive optomechanical measurements in a new cryogen-free dilution refrigerator that uses a “helium battery,” which allows approximately two hours of low-vibration measurement time [BlueFors 2023; Uhlig 2023]. This system allows for more rapid cycling and measurement of devices, and for more turn-key operation and the practical capability to keep a device at cryogenic temperatures for a period of months, which is valuable as the complexity of our setup increases and includes more experimental modules. Orienting towards future transducer demonstrations, we have extended our ability to optomechanically ground-state cool our transduction mode to this new cryogenic system. We have also measured the thermal motion of optical cavity mirrors at dilution-refrigerator temperatures, which will contribute additional noise counts in an experiment involving single-photon detection.

In Ch. 2, I introduce the figures of merit used to evaluate quantum transducers, and compare the most successful recent demonstrations achieved by different platforms. In Ch. 3, I investigate potential experiments to be done with transducers near quantum performance and what the experiments require of the setup and the transducer. I start with entanglement distribution protocols that require two microwave nodes and two transducers, and then explore demonstrations that evaluate subsections of such a rudimentary network. In Ch. 4, I discuss the design constraints on the optomechanical portion of the transducer, and what I’ve chosen to optimize for achieving ground state cooling and quantum transduction goals, as well as broader indications for future, more stringent networking goals. In Ch. 5, I discuss the transducer device whose mechanical mode we first cooled to a phononic occupation of less than one photon, which we used to demonstrate compatibility of the platform with a superconducting qubit system by reading out the state of the qubit via the transducer. And in Ch. 6, I discuss subsequent devices and improvements to the optical measurement setup as we look towards more sensitive experiments using single-photon detection, and detail transducer operation in a cryogen-free dilution refrigerator.

Chapter 2

Transducer requirements in a quantum network

In this chapter, I'll introduce the figures of merit that characterize a transducer, as working to improve those metrics towards a quantum-enabled device has been the primary goal that has oriented my work at JILA. I'll also briefly indicate how these metrics affect performance as they pertain to enabling quantum applications, for example quantum state transfer or remote entanglement generation via the communication channel (Ch. 3 describes specific protocols and experiments in more detail). I'll then review several prominent recent transducer demonstrations that use different platforms, and assess the tradeoffs between those platforms as they relate to the transducer figures of merit.

Transducers are generally characterized by three metrics: their efficiency η_t , their input-referred added noise N_{add} , and their bandwidth Γ_t . An additional consideration that is relevant for communication rates is the duty cycle D , which may be limited by e.g. absorption of the optical pump disrupting a superconductor in platforms with an optical guided mode.

2.1 Efficiency and added noise

By transducer efficiency, I refer to the photon number efficiency of the transduced signal, i.e. $\eta_{\text{up}} = \left| \hat{a}_{\text{out}} / \hat{b}_{\text{in}} \right|^2$ in upconversion, where \hat{b}_{in} is the incident microwave signal field, and \hat{a}_{out} is the transduced optical signal field, as illustrated in Fig. 2.1. The efficiency in downconversion η_{down} is defined analogously. Because it can be calibrated more precisely, I'll generally refer to the geometric mean of these two quantities, the transducer's bidirectional efficiency $\eta_t = \sqrt{\eta_{\text{up}} \eta_{\text{down}}}$,




Figure 2.1: **Transduction in upconversion and downconversion.** Microwave-to-optical transduction in upconversion with efficiency $\eta_{\text{up}} = |\hat{a}_{\text{out}}/\hat{b}_{\text{in}}|^2$ is represented on the left, and downconversion with efficiency $\eta_{\text{down}} = |\hat{b}_{\text{out}}/\hat{a}_{\text{in}}|^2$ is represented on the right.

which describes the probability that a given incident input photon is converted to an output photon, as illustrated in Fig. 2.2 [Andrews et al. 2014].

A transducer will in general add noise to the output signal, quantified by the average photon number flux per Hz of bandwidth N_{out} , in units of photons/s/Hz. In order to compare the noise directly with an incident signal it would be competing with, the more useful number for evaluating the performance of a transducer is the input-referred added noise $N_{\text{add}} = N_{\text{out}}/\eta_t$, illustrated in Fig. 2.2. $N_{\text{add}} = 1$ photon/s/Hz is considered a threshold for quantum performance, the noise value required to resolve a single transduced signal photon with a signal-to-noise ratio (SNR) of one, assuming detection optimized for the bandwidth and temporal envelope of the signal photon.

The efficiency and input-referred added noise are the crucial figures of merit in evaluating a transducer’s performance in a quantum network [Zeuthen et al. 2020]. Whereas both metrics are important for direct state transfer, N_{add} can be regarded to be of greater importance, as employing heralding in a protocol to distribute entanglement ensures that η_t does not affect the fidelity of the entanglement created (see Ch. 3). Instead any inefficiency reduces the success probability of the protocol, so it impacts the rate of entanglement generation.

2.2 Networking rate: bandwidth and duty cycle

The bandwidth Γ_t and duty cycle $D = T_{\text{on}}/T_{\text{rep}}$ also impact the communication rate of the transducer, where T_{on} is the transducers pulsed “on-time” and T_{rep} is the repetition time (Fig. 2.3). Neglecting noise, the average communication rate is proportional to the product $\eta_t \Gamma_t D$. There is




Figure 2.2: **Efficiency and added noise.** A given signal photon is transduced with probability η_t . Noise photons contaminate the transduced signal at the output at rate N_{out} . The input-referred added noise $N_{\text{add}} = N_{\text{out}} / \eta_t$ compares the noise to the input signal.

additional nuance when considering these individual contributions to the overall rate.

One crucial criterion is that Γ_t must be sufficiently large to be compatible with the lifetimes of the superconducting qubit, such that the transduction efficiency does not suffer due to a bandwidth mismatch with the signal. Bandwidth mismatch was the leading source of inefficiency when reading out the state of a transmon qubit via our electro-optic transducer in Ref. [Delaney et al. 2022] (see Table 5.2). When the signal is a short pulse, the transducer bandwidth will limit the efficiency of the transduced readout pulse as

$$\eta_{\text{bw}} \approx 1 - 2 \frac{1 - e^{-\Gamma_t T_p / 2}}{\Gamma_t T_p}, \quad (2.1)$$

where the signal pulse is approximated as a square wave of width T_p [Delaney et al. 2022]. State-of-the-art superconducting qubit energy decay lifetimes are in the hundreds of μs to 1 ms range [Pop et al. 2014; Place et al. 2021; Somoroff et al. 2023; Ganjam et al. 2023], requiring transduction bandwidths of at least $\Gamma_t = 2\pi \cdot 6 \text{ kHz}$ for efficient transduction.

The impact of a transducer's duty cycle also depends on intricacies of the protocol being implemented. For example, a transducer with duty cycle $D = 10^{-3}$ defined by $T_{\text{on}} = 100 \mu\text{s}$ of continuous operation repeated every $T_{\text{rep}} = 100 \text{ ms}$ could transduce multiple times sequentially during the course of a computation. On the other hand, a transducer with the same value of D resulting from a single $T_{\text{on}} = 1 \mu\text{s}$ pulse repeated every $T_{\text{rep}} = 1 \text{ ms}$ would not get a second attempt within state-of-the-art qubit lifetimes.

Additionally, there is ongoing theoretical work that explores the interplay of these figures of




Figure 2.3: **Bandwidth and duty cycle.** The transducer bandwidth Γ_t limits the rate at which a signal can be transduced: a signal pulse of width $T_p < 2\pi/\Gamma_t$ will be temporally smeared out, and the efficiency with which it is transduced will be decreased (left). The duty cycle $D = T_{\text{on}}/T_{\text{rep}}$ quantifies the amount of time a transducer is limited to running T_{on} per repetition of length T_{rep} (right).

merit, for example, the impacts of finite efficiency and bandwidth on the quantum capacity of a bosonic channel in pure-loss and thermal-loss models [Wang, Li, and Jiang 2022].

2.3 Comparison between several transducer platforms

I'll now detail the relative performances of different transducer platforms, comparing their values of $N_{\text{add,up}}$, η_t , Γ_t , and D and discuss their relative strengths. The comparisons are summarized in Table 2.3.

2.3.1 Transduction via low-frequency mechanical mode

First, I'll discuss the platform that is the focus of this dissertation, a transducer based on the coupling of the same mechanical mode of a tensioned silicon nitride membrane to optical and microwave resonators. This platform leverages the expertise of the cavity optomechanics community, and its primary strengths and weaknesses are the consequence of the MHz-frequency mechanical modes, and the benefits of being able to use high finesse optical cavities. The membranes can be engineered to have low mechanical loss rates compared to the electromechanical and optomechanical rates, so that efficiency is only limited by the loss rates of the microwave and optical resonators and a factor for matching the optical cavity mode to the optical control and measurement setup. As a result, in our collaboration, we regularly achieve efficiencies of approximately $\eta_t = 0.5$.

Table 2.1: Comparison of transducer platforms.

Platform	$N_{\text{add,up}}$	η_t	$\Gamma_t/2\pi$	D	$\frac{\eta_t \Gamma_t D}{2\pi}$	Ref.
Membrane-based	1600	0.09	5.6 kHz	1	500 Hz	[Andrews et al. 2014]
	34	0.47	3.5 kHz	1	1.6 kHz	[Higginbotham et al. 2018]
	3.2	0.44	200 Hz	1	90 Hz	[Brubaker et al. 2022]
Resonant piezo-opto-mechanical	0.57	7×10^{-4}	20 kHz	$\frac{38 \text{ ns}}{10 \text{ ms}} = 4 \times 10^{-6}$	60 μ Hz	[Mirhosseini et al. 2020]
	0.15-0.2	0.08	1 kHz	$\frac{160 \text{ ns}}{20 \text{ } \mu\text{s}} = 8 \times 10^{-3}$	160 mHz	[Meesala et al. 2024]
				$\frac{96 \text{ ns}}{20 \text{ } \mu\text{s}} = 5 \times 10^{-3}$	400 mHz	[Meesala et al. 2023]
Direct electro-optic	1.8×10^4	3×10^{-4}	10 MHz	1	3 kHz	[Hease et al. 2020]
	0.16	0.09	18 MHz	$\frac{300 \text{ ns}}{100 \text{ ms}} = 3 \times 10^{-6}$	50 Hz	[Sahu et al. 2022]
mmwave-to-optical with cold atoms	0.6	0.025	360 kHz	$\frac{200 \text{ } \mu\text{s}}{14 \text{ ms}} = 1.4 \times 10^{-2}$	130 Hz	[Kumar et al. 2023]

A drawback of the MHz-scale mechanical resonance frequency is that membrane-based transducers require a microwave pump to bridge the mechanical and microwave frequencies, in addition to the optical pump required by all electro-optic transducers. Furthermore, the low mode frequency means that it remains highly thermally occupied, even in the ~ 50 mK environment of a dilution refrigerator. As a consequence, the transducer must be driven continuously and strongly by the optical and microwave pumps in order to avoid transmitting a quantum signal through a noisy channel (see Ch. 4.1.3 for more details).

On the optical side, the Fabry-Pérot cavity architecture and the associated low amounts of scattered and absorbed light ensures isolation of the high amounts of optical circulating power from the fragile superconducting resonator. This isolation is a strength unique to membrane-based transducers when compared with other microwave-to-optical transducers, and permits them to operate continuously, such that $D = 1$. However, the power required of the microwave pump does deleteriously impact the loss and noise of the the microwave resonator, which limits the performance of the transducer. The added noise for our platform is most recently limited at the threshold of quantum performance, to a value of $N_{\text{add}} = 3.2$ photons/s/Hz.

Finally, the sharp dependence of the noise on pump power also currently limits the bandwidth of of our transducers to kHz-scales, at the threshold of compatibility with superconducting qubit lifetimes. The mechanical frequency itself and few-hundreds-of-kHz proximity of additional vibrational membrane modes place additional constraints on the achievable bandwidth with this platform.

In Table 2.3, I have compiled the performance of our transducers in Refs. [Andrews et al. 2014; Higginbotham et al. 2018], and [Brubaker et al. 2022], tabulating values that were measured at a consistent set of operating parameters when possible. For example, though we measured a maximum transduction efficiency of $\eta_t = 0.47$ in Ref. [Brubaker et al. 2022], that was not the efficiency at the operating point of lowest $N_{\text{add}} = 3.2$ photons/s/Hz.

2.3.2 Microwave-frequency resonant piezo-optomechanical transducers

A transducer based on a microwave-frequency optomechanical resonator avoids the problems associated with the low mechanical mode-frequency of the membrane-based devices. Such devices do not require a microwave pump, and because the mechanical mode's thermal environment is close to its quantum ground state at dilution-refrigerator temperatures, continuous pumping is not required to prevent thermal phonons from contaminating the transduced signal. A number of experimental efforts have applied nanomechanical resonators to reach these mechanical frequencies for microwave to optical transduction [Bochmann et al. 2013; Balram et al. 2016; Forsch et al. 2020; Jiang et al. 2020; Han et al. 2020]. It was used for the first successful demonstration of quantum-enabled microwave-to-optical transduction, with $N_{\text{add}} = 0.57$ [Mirhosseini et al. 2020]. That demonstration coupled a superconducting transmon qubit to an optomechanical crystal cavity (OMC) [Eichenfield et al. 2009; Chan et al. 2011], via a resonant piezoelectric coupler. An optical pulse then read out the phononic occupation of the OMC with bandwidth $2\pi \cdot 20 \text{ kHz} = \Gamma_t$, upconverting the microwave phonon transduced from the superconducting qubit to optical frequencies.

However, absorption of the optical pulse led to appreciable heating of the device, destroying the state of the qubit and requiring a repetition rate of 100 Hz to allow the device to rethermalize. Combined with an intrinsic transduction efficiency of 10^{-3} and an optical collection efficiency of 10^{-2} , in a low rate of successfully detected transduced photons of 1 mHz, or about 100 counts per day. The value for η_t I include in Table 2.3 is the product of their quoted intrinsic efficiency and the independently measured optical mode matching of the readout pulse to the device's readout port.

In a more recent experiment from the Painter group, the duty cycle was dramatically improved by a number of technical improvements including using a superconductor with a faster quasiparticle recombination rate [Meesala et al. 2024]. In this experiment, rather than upconversion, the transducer was used to generate pairs of microwave and optical photons using spontaneous parametric down-conversion (SPDC). By conditioning on the measurement of a single optical pho-

ton from the transducer, nonclassical correlations were observed in the field emanating from the transducer’s microwave port. With the conservative assumption that the other upconversion metrics remain the same as the group’s earlier work, the improvement to the duty cycle alone would result in an over three orders of magnitude improvement in the rate metric $\eta_t \Gamma_t D / 2\pi$ from 60 μHz to 160 mHz.

A few months later, the group used the same SPDC process combined with a time-bin encoding in order to measure entangled pairs of photons at the outputs of the device, an impressive result that demonstrates the building blocks needed to entangle two superconducting quantum computers [Meesala et al. 2023].

A very promising variation on the piezo-optomechanical transduction scheme uses high-overtone bulk acoustic wave resonators (HBARs), simultaneously leveraging microwave-frequency acoustic frequencies and a larger optical mode volume compatible with Fabry-Pérot cavities to mitigate optical heating [Yoon et al. 2023; Doeleman et al. 2023; Bild et al. 2023].

2.3.3 Direct electro-optic transduction with nonlinear materials

It is also possible to directly transduce microwave photons to optical frequencies via a three-wave mixing process with a nonlinear material such as lithium niobate or aluminum nitride. These devices can operate with MHz-scale bandwidths [Fan et al. 2018; Xu et al. 2021]. The most successful demonstration of direct microwave-to-optical transduction uses a millimeter-scale disk of lithium niobate embedded in a tunable 3D superconducting microwave cavity [Rueda et al. 2016], recently employed very successfully by the Fink group [Hease et al. 2020; Sahu et al. 2022]. In that device, the disk guides optical whispering gallery modes, and the microwave resonator is tuned to be resonant with the optical free-spectral-range. By taking advantage of avoided crossings between optical polarization modes to frustrate the regular frequency spacing between adjacent optical modes, a single optical mode can be isolated for efficient transduction. Optical absorption leads to increased loss and noise in the microwave cavity, and so continuous operation of this device has efficiency of $\eta_t = 3 \times 10^{-4}$ and tens of thousands of input-referred added noise photons [Hease et al.

2020]. By pulsing the device’s operation, the performance of the device improves dramatically, demonstrating efficiencies of close to 0.1. And by operating at a duty cycle of approximately 3×10^{-6} , the device has demonstrated $N_{\text{add,up}} = 0.16$ photons [Sahu et al. 2022].

2.3.4 Millimeter-wave-to-optical transduction with trapped atoms

Trapped atoms also offer a mechanism for bridging microwave and optical frequencies [Verdú et al. 2009; Hafezi et al. 2012; Covey, Sipahigil, and Saffman 2019; Vogt et al. 2019]. The final platform I’ll discuss uses trapped atoms to mediate transduction of excitations of a millimeter wave (mmwave) resonator to optical frequencies, focusing on the results of the collaboration between the Simon and Schuster groups [Kumar et al. 2023]. Though the mmwave resonator frequency of 99 GHz is roughly a factor of 10 higher than the microwave frequencies used by more mature superconducting qubit devices, this platform has demonstrated high efficiency, high bandwidth, and low added noise.

In this experiment, mm-scale waveguides intersect at right angles in superconducting niobium to form a mm-wave resonance at the intersection whose mode evanescently decays in the waveguides, isolating it from the environment. Rubidium-85 atoms are trapped and loaded into the mmwave resonator through one of the waveguides. An optical cavity whose axis is defined along the second waveguide enhances the interactions of the atoms with the optical field of wavelength 780 nm. The final waveguide is used to couple in mmwave signals. Two additional pump lasers applied to the atoms enhance a four-wave mixing process linking frequencies near resonance of the optical and mmwave cavities via transition through a Rydberg state. Scaling the size and reducing the temperature of the superconducting resonator, and selecting a different Rydberg state offers a path towards modifying the platform to be compatible with microwave frequencies.

The platform demonstrated transduction efficiency of $\eta_t = 0.025\%$ at a $\Gamma_t = 2\pi \cdot 360$ kHz bandwidth, and with $N_{\text{add}} = 0.6$ photons of added noise. The dominant source of noise and inefficiency was due to operating the mmwave cavity at a temperature of 5 K, leading thermal occupation and its being undercoupled by a factor of approximately 15. Reducing the operating

temperature of the mmwave cavity to 1 K would reduce the thermal contribution to the added noise to 0.01 photons, and reduce its loss rate by a factor of 100, which would increase the efficiency to approximately $\eta_t = 0.3$ with that change alone. Unlike transducer platforms using guided optical modes, the duty cycle of the experiment is not limited by optical heating disrupting the superconductor and adding noise, but by the 14 ms loading time of the atoms, after which the transducer can operate continuously for hundreds of microseconds. One potential consequence of the longer timescale is that it could allow for multiple states being transduced in the computation time of a superconducting processor, assuming the computation time is limited to the lifetime of its constituent qubits.

Chapter 3

Primitive networking demonstrations

As a transducer's N_{add} approaches thresholds that permit quantum state transfer, it can conceivably be used to distribute entanglement between two superconducting registers, the most primitive network. There are several different protocols that can realize entanglement distribution. In this chapter, I'll present three illustrative examples of such protocols, and weigh their relative strengths and weaknesses in the context of linking superconducting circuits using electro-optic transducers. I'll then detail demonstrations to show progress towards a simple network that do not require the full topology of two superconducting nodes and two transducers, and the experimental requirements for these demonstrations.




Figure 3.1: **Direct state transfer.** An arbitrary quantum state is prepared in the left node, pitched to the electro-optic transducer where it is upconverted, propagates optically, and is down-converted and absorbed in the right node.

3.1 Entanglement protocols

3.1.1 Direct state transfer

Perhaps the most conceptually straightforward family of protocols is direct quantum state transfer, also known as “pitch-and-catch”-style communication, as illustrated in Fig. 3.1. Direct state transfer has been successfully demonstrated in experiments linking superconducting nodes via microwave-frequency transmission lines [Axline et al. 2018; Grebel et al. 2024]. Here, a quantum state is prepared at one superconducting node, and is then swapped to a transmission line. The quantum state propagates as a flying qubit towards the receiving register, which absorbs the transmitted state. To generate entanglement, the protocol can be modified to release only half of the energy from the sending node (a “half pitch”), resulting in the creation of a Bell state, $(|10\rangle + |01\rangle)/\sqrt{2}$, on absorption by the receiving node.

I’ll note here that the time-dependent control over the coupling between the superconducting register and the transmission line would be achieved using four-wave mixing and a setup similar to that described in Refs. [Pfaff et al. 2017; Axline et al. 2018], for the direct state transfer protocol as well as all subsequent protocols described in this chapter. According to this experimental realization, each superconducting node is composed of two microwave cavities, a long-lifetime storage cavity I’ll refer to as the register and a second cavity, the communication cavity, strongly coupled to a transmission line. A pumped transmon dispersively coupled to both cavity modes permits time-dependent coupling between the cavity modes.

Introducing transducers along the transmission line as shown in Fig. 3.1 would permit communication between nodes in separate cryostats. However, the transducer inefficiencies and transmission losses incurred directly affect the fidelity of the resulting state for direct transfer protocols. Furthermore, direct state transfer requires operating the first transducer in upconversion from microwave to optical frequencies, and the second in downconversion, back to microwave frequency. The performance of the transducer need not be the same in upconversion and downconversion, as detailed further in Ch. 5, so a pitch-and-catch protocol could require separate optimizations for the




Figure 3.2: **Heralded entanglement generation with the DLCZ protocol.** The state ϕ_{DLCZ} is prepared on both left and right nodes, resulting in the entanglement of the superconducting register with a microwave photon propagating towards the transducer, with low probability p_e . The upconverted signals are combined on a beamsplitter and detected with single-photon detectors. A click from either detector heralds the creation of the Bell state $|\Psi^\pm\rangle$, and the two superconducting registers are entangled.

two transducers, and is potentially a more complex path forward. Moreover, downconversion would require additional technical care in shielding the receiving qubit, which could prove a demanding requirement.

3.1.2 Singly-heralded protocols with weak excitation

To prevent transmission loss from negatively affecting communication fidelity, it is possible “herald” and post-select on successful transmission events, at the cost of a reduced communication rate. There are a number of protocols that make use of a single detection event to overcome transmission loss [Cabrillo et al. 1999; Duan et al. 2001; Moehring et al. 2007; Campbell and Benjamin 2008]. Compared to doubly-heralded protocols that use two detection events, such as the one described in the following section, singly-heralded protocols are advantageous for communication rates when transmission efficiency is very low, as the probability of detecting a single click is greater than detecting two clicks. A network topology to be used in a singly-heralded protocol is shown in Fig. 3.2. Though the storage cavities are bosonic modes, we are selecting the Fock states $|0\rangle$ and $|1\rangle$ to be our computational basis, and I’ll illustrate the register as a qubit. Initially, each node is prepared by a weak preparation pulse, resulting in a low probability $p_e \ll 1$ of creating an

entangled pair of excitations in the register and propagating down the transmission line as a flying qubit,

$$|\phi_{\text{weak}}\rangle = |0_{\text{N}}\rangle \otimes |0_{\text{F}}\rangle + \sqrt{p_{\text{e}}} |1_{\text{N}}\rangle \otimes |1_{\text{F}}\rangle + \mathcal{O}(p_{\text{e}}), \quad (3.1)$$

where $|0_{\text{N}}\rangle$ and $|0_{\text{F}}\rangle$ denote the ground state of the node register and transmission line, respectively, both of which I'm treating as bosonic modes, $|1_{\text{N}}\rangle$ and $|1_{\text{F}}\rangle$ analogously represent single excitations in these channels, and $\mathcal{O}(p_{\text{e}})$ represents terms with probabilities $\leq p_{\text{e}}^2$. The state $|\phi_{\text{weak}}\rangle_{\text{L}} \otimes |\phi_{\text{weak}}\rangle_{\text{R}}$ then describes the two nodes together, where the L and R subscripts denote the left and right nodes, respectively. A click on one of the detectors then heralds the Bell state

$$|\Psi\rangle^{\pm} = \frac{|1_{\text{N}}\rangle_{\text{L}} \otimes |0_{\text{N}}\rangle_{\text{R}} \pm |0_{\text{N}}\rangle_{\text{L}} \otimes |1_{\text{N}}\rangle_{\text{R}}}{\sqrt{2}} = \frac{|10\rangle \pm |01\rangle}{\sqrt{2}} \quad (3.2)$$

up to a relative phase shift dependent on the path lengths, with the sign dependent on which detector clicks. Success will occur probabilistically, approximately once every $1/2\eta p_{\text{e}}$ attempts, where η is the efficiency of each channel (assuming symmetric loss). A given click prepares the states $|11\rangle$, $|02\rangle$, and $|20\rangle$ each with probability p_{e} , so the infidelity due to such events, $1 - F \propto p_{\text{e}}$, can be made small by decreasing p_{e} , of course also decreasing the success rate.

This protocol uses the low probability of creating multiple excitations, along with conditioning on a click to herald out any contribution of photon loss to the infidelity of the prepared entangled state, and so is naturally compatible with systems that have low coupling. It has been used successfully to entangle purely optomechanical systems, notably vibrational modes of room-temperature diamonds, [Lee et al. 2011], and more recently, two optomechanical crystals [Riedinger et al. 2018]. However, it's probabilistic nature and dependence on p_{e} being small dramatically limits the rate that entanglement can be generated.

It should also be noted that the detectors have a background noise count rate (resulting in what are called dark counts), and can also click due to technical noise (e.g. from a transducer's N_{add}) that accompanies the signal. Including these auxiliary detection events in the model results in preparation of the mixed state described by the density operator

$$\rho = \frac{1}{1 + c_0} (c_0 |00\rangle \langle 00| + |\Psi\rangle^{\pm} \langle \Psi|^{\pm}) \quad (3.3)$$

where c_0 is determined by the rate of dark and noise counts. Assuming that the rate of dark counts is much smaller than the rate of noise counts, $c_0 = p_N/p_e$ is the ratio of counts due to noise to counts from successfully transmitted heralding photons. Here, p_N is the probability of detecting a noise photon in one experimental repetition, referred to the qubit output, i.e. the ratio of qubit-referred noise photons to signal photons, if the qubit is successfully pitching the state $|1\rangle$ once every repetition. I'll also define $\text{SNR} = 1/p_N$ here, defined such that the value of SNR is not modified by the state being transduced, in the case that state contains less than one photon. Thus, in order to prepare a Bell state with fidelity greater than $1/2$, the entanglement purification threshold, a system must have a signal-to-noise ratio greater than $1/p_e$, or equivalently $p_N < p_e$. So although this protocol is robust to loss, it is very sensitive to noise because we demand p_e be small.

For concreteness, to put p_N in the context of an acceptable amount of laser technical noise, a value of $p_N = 1$ requires a technical noise level of -141 dBc/Hz for our system, assuming a noiseless transducer and otherwise reasonable transducer and microwave loss parameters (see Sec. 6.2.1). The contribution to p_N from transducer added noise is the qubit-referred added noise, $N_{\text{add}}/\eta_{\text{MW}}$, where η_{MW} is the microwave propagation efficiency between the qubit and the transducer. So for an ideal setup in which the only noise source is from the transducer, $p_N = N_{\text{add}}/\eta_{\text{MW}}$. For a real setup with both transducer and laser technical noise, p_N is simply the sum of the two individual contributions.

It would also be possible to realize a similar singly-heralded experiment with two-level systems at each node, instead of the bosonic mode of the storage cavity acting as the register [Cabrillo et al. 1999; Campbell and Benjamin 2008]. The node would then be composed of a transmon qubit dispersively coupled to a single readout cavity, as Rob Delaney described in his thesis [Delaney 2022]. Such protocols have been demonstrated experimentally using nitrogen-vacancy centers in diamond [Kalb et al. 2017; Pompili et al. 2021].




Figure 3.3: **Doubly heralded entanglement protocol.** At the beginning of the first heralding round, each superconducting register is first prepared in the $|+\rangle$ state, and the microwave transmission line input to the transducer is initially in its ground state. A CNOT gate controlled by the register and targeting the transmission line entangles each node with a flying qubit propagating towards the transducer. The optical outputs are interfered on single-photon detectors, and a click heralds the mixed state $\rho_{1\text{click}}$, which includes support of the Bell state $|\Psi^\pm\rangle$ and the unentangled state $|11\rangle$. A π -pulse is applied to both superconducting registers, rotating the $|11\rangle$ component to $|00\rangle$, before the second round of CNOT gates and detection. A second click then heralds $|\Psi^\pm\rangle$, where the sign depends on whether the same or alternate detectors clicked.

3.1.3 Doubly-heralded protocol

Compared with singly-heralded protocols, doubly-heralded protocols [Duan and Kimble 2003] can bring success probability close to unity assuming high detection efficiency. The Barret and Kok protocol is a heralding scheme that uses what amounts to a dual-rail encoding with a second round of heralding, shown in Fig. 3.3, to eliminate the possibility of preparing the $|11\rangle$ state, in the absence of dark counts [Barrett and Kok 2005]. Removing the infidelity from preparing the $|11\rangle$ state allows for deterministic entanglement of the nodes with flying qubits in each attempt, modifying the probability of success to $\eta^2/2$, compared with $2\eta p_e$ in the singly-heralded case. This protocol has been demonstrated using superconducting qubits as the nodes and an additional qubit as a microwave-frequency single-photon detector [Narla et al. 2016].

With our bosonic registers, we can prepare Fock states with high fidelity using projective measurement, so we can treat them as qubits for this protocol, in line with the original proposal which envisioned platforms like nitrogen-vacancy centers, quantum dots, or trapped ions. The protocol begins with each qubit node initialized in the $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ state. A controlled-NOT (CNOT) gate then entangles the node with a flying qubit, resulting in the state $(|0_N\rangle |0_F\rangle)/\sqrt{2}$ on each side. With perfect efficiency and number-resolving detectors, a single click on one detector would result in the creation of desired Bell state, $|\Psi\rangle^\pm$, with the sign again depending on which detector clicks. However, given finite efficiency and non-number-resolving detection, the mixed state

$$\rho_{\text{1click}} = \mathcal{N} |\Psi\rangle^\pm \langle\Psi|^\pm + (1 - \mathcal{N}) |11\rangle \langle 11| \quad (3.4)$$

is prepared, where the normalization constant \mathcal{N} depends on the loss and photon number-resolving capability of the detector. A π pulse is then applied to both nodes, and the entangling gate and heralding steps are repeated, such that they herald out the second term, and the nodes are left in the entangled state $|\Psi\rangle^+$ if the same detector clicks both times, and $|\Psi\rangle^-$ if alternate detectors click with each round.

Again, the fidelity of the prepared entangled state is affected by dark counts and noise counts

in the detectors. Following the analysis in the supplement of Ref. [Narla et al. 2016], and assuming symmetric loss and a single detector with no dark counts, a Bell state is created with fidelity

$$\mathcal{F} = \frac{3p_N^2 + p_N + 4}{11p_N^2 + 17p_N + 4}. \quad (3.5)$$

Thus, for a fidelity of greater than $1/2$, a value of p_N of less than approximately 0.25 is required, a significantly relaxed constraint compared with that of the singly-heralded protocol. Ref. [Zeuthen et al. 2020] explores the comparison between these protocols in detail, though in the more optimistic limit $p_N \ll 1$.

3.2 Single-node demonstrations

With these entanglement protocols in mind as goals for demonstration of a two-node network, I'll now consider experiments to evaluate a quantum channel that can be done with a single node, a single transducer and optical single-photon detection. In the very first single-node demonstration of a quantum-enabled qubit-transducer system, the Painter group upconverted the state of a transmon qubit prepared with a Rabi drive of variable length, and found that the probability of their optical detector clicking depended on length of the Rabi drive [Mirhosseini et al. 2020]. We have also performed a single-node demonstration in which we upconverted the classical signal from dispersive readout using the device described in Ch. 5, measuring the state of the qubit using optical heterodyne detection. This experiment is described in detail in Rob Delaney's thesis [Delaney 2022] and Ref. [Delaney et al. 2022]. In this chapter, I'll consider additional demonstrations that we could perform with our transducer to demonstrate improving performance. I'll detail a number of intensity correlation measurements and a demonstration of a dual-rail encoding using frequency-bin qubits, and evaluate their noise requirements and the resulting integration times needed to perform the measurement.

3.2.1 Intensity correlation measurement with a deterministic quantum source

The first demonstration I'll consider is an experiment to measure the second-order coherence, $g^{(2)}$, of the transduced signal, also referred to as the intensity correlation function. Such a demonstration would verify the ability to resolve a quantum source at the input of a channel, but places strong demands on the channel efficiency to accomplish the measurement in a reasonable averaging time, as pointed out in Ref. [Mirhosseini et al. 2020], for whom low efficiency prohibited a measurement of $g^{(2)}$. A light source's $g^{(2)}$ quantifies correlations in the arrival time of photons, and gives information about the statistics of the source to reveal if it is e.g. coherent, thermal or a single emitter. The intensity correlation function is measured using single-photon detectors, in a setup somewhat similar to that required for the heralding protocols discussed in the previous section, and is therefore a good intermediate demonstration.

Here, I'd like to briefly make a parenthetical distinction that an intensity correlation measurement is different from another possible correlation measurement, one quantifying the correlations between fields. Two-mode-squeezing, for example, is characterized by strong correlations between the fields at the output of the microwave and optical ports, and is measured by heterodyne detection, as in Ref. [Sahu et al. 2023]. Recent theoretical work suggests that operating our membrane-based transducer in a two-mode-squeezing configuration is advantageous compared with operating in the beamsplitter configuration, assuming that both electromagnetic resonators are not thermally occupied [Rau et al. 2022]. Because of noise associated with our microwave circuit, the assumption of a cold microwave resonator is not currently valid for our transducer, but once we can operate a transducer with lower microwave circuit noise, measuring two-mode-squeezing below vacuum is another exciting experimental direction.

A light source's intensity correlation function $g^{(2)}(\tau)$ gives the comparative likelihood of detecting a second click a time τ after an initial detection event, compared to the probability of that first detection event

$$g^{(2)}(\tau) = \frac{\langle \hat{a}^\dagger(t) \hat{a}^\dagger(t+\tau) \hat{a}(t+\tau) \hat{a}(t) \rangle}{\langle \hat{a}^\dagger(t) \hat{a}(t) \rangle^2}, \quad (3.6)$$

where \hat{a} is the annihilation operator of the optical mode incident on a single-photon detector [Gerry and Knight 2004]. For large delay times (compared to a characteristic coherence time of the light source), $g^{(2)}(\tau)$ is 1, as the photons arrival time is uncorrelated. As τ approaches 0, $g_0^{(2)} = \frac{\langle \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a} \rangle}{\langle \hat{a}^\dagger \hat{a} \rangle^2} = \frac{\langle \hat{n}(\hat{n}-1) \rangle}{\langle \hat{n} \rangle^2}$ reveals information about the statistics of the source, where I am defining $g_0^{(2)} = g^{(2)}(0)$ as a convenient shorthand. A source emitting single photons, for example, has $g_0^{(2)} = 0$, because only a single photon will arrive at a time. For coherent sources, $g_0^{(2)} = 1$, as detection events obey Poissonian statistics, and their probability is independent of any previous click. For thermal sources exhibiting bunching, $g_0^{(2)} = 2$, and for Fock states, $g_0^{(2)} < 1$, so $g_0^{(2)} = 1$ is considered a threshold for demonstrating statistics of a quantum mechanical source.

When characterizing a communication channel, a $g^{(2)}$ measurement quantifies the ability to resolve a quantum emitter at the input. We can therefore evaluate one half of the setup required in the entanglement protocols above, as in Fig. 3.4, where a superconducting node pitches single excitations to a transducer, and the optical output is routed to a pair of single-photon detectors that record the number of detection events after each pitch. Though the canonical Hanbury Brown and Twiss setup uses a beamsplitter and two single-photon detectors to measure the intensity correlation function, if the detectors are fast enough to be considered number resolving, the experiment can be performed with a single detector, in contrast with what is shown in Fig. 3.4. The detection speed is limited by the finite dead time of after each detection event during which the detectors are not sensitive to additional photons. For superconducting nanowire single-photon detectors (SNSPDs), the dead time is < 100 ns [Marsili et al. 2013]. Provided that the dead time is small compared to the flying qubit's pulse width (which can be further broadened by the finite bandwidth of the transducer), a single-photon detector is sufficient to perform the correlation experiment. For a transducer operating with a bandwidth of 10 kHz, the detector reset time is at most a factor of 10^3 smaller than the pulse width of the excitation, and it would be safe to treat such single-photon detectors as number-resolving (see Ref. [Kindem 2019] for more discussion of this approximation).

Assuming all noise is from a single thermal source, measuring $g_0^{(2)} < 1$ implies an SNR of greater than one, as the thermal source contributes $g_0^{(2)} = 2$. However, multiple uncorrelated ther-




Figure 3.4: **Intensity correlation measurement setup.** The superconducting node deterministically sends single microwave photons (green line) to the transducer, where they are upconverted and detected on a Hanbury Brown and Twiss setup to measure the correlation between detection events. For an ideal setup with only a source of single-photon Fock states, the correlation measurement $g^{(2)}$ would be zero as both detectors would never click simultaneously. Noise counts introduced by transducer added noise with thermal statistics (gray line), and optical pump light with coherent statistics (red line) that can leak through the imperfect suppression of the filter cavities, increase the measured $g^{(2)}$ from this ideal value.

mal sources or additional coherent sources will decrease the value of $g_0^{(2)}$ contributed from sources other than the quantum signal, so violating the nonclassical bound $g_0^{(2)} > 1$ would correspond to a relaxed constraint on the SNR.

The 1.5 MHz frequency of the mechanical mode used in our transducer defines the separation in frequency between the transduced signal excitations and optical pump, which is approximately 100 dB larger in magnitude. To isolate the signal from the pump, we will use a system of three cascaded filter cavities to isolate the signal and reject the pump, modeled after the setup used in Eugene Polzik's group in Ref. [Galinskiy et al. 2020], which will be described in more detail in Luca Talamo's thesis. Each cavity has linewidth of approximately 50 kHz, and the three cavities taken together have an effective bandwidth of 30 kHz, i.e. would pass the same amount of power as an infinitely sharp 30-kHz band-pass filter. Any leakage of the pump will contribute to additional counts in a photon counting experiment, with coherent statistics. Other sources of noise tend to be thermal. In Fig. 3.4, I've represented three contributions to detector count rate with different statistics, with the coherent optical pump contribution denoted in red, the single microwave photon signal in green, and the thermal contribution from transducer added noise in gray. Technical noise on the optical pump would also generally have thermal statistics, though I have not depicted that contribution explicitly in Fig. 3.4. Care must be taken in accounting for the photon statistics when combining independent thermal sources of noise, as they will remain uncorrelated at $\tau = 0$.

More quantitatively, when two sources with average powers P_1 and P_2 are combined on a detector, their joint intensity correlation function is given by

$$g^{(2)} = \frac{P_1^2 g_1^{(2)} + P_2^2 g_2^{(2)} + 2P_1 P_2}{(P_1 + P_2)^2} \quad (3.7)$$

where $g_1^{(2)}$ and $g_2^{(2)}$ denote their respective independent intensity correlation functions [Loudon 1973]. Fig. 3.5(a) shows the expected results of a $g^{(2)}$ experiment combining a transduced qubit signal with a noise field as a function of the noise-to-signal ratio $p_N = 1/\text{SNR}$ (introduced in Sec. 3.1.2), using Eq. 3.7, assuming noise from a the combination of a single thermal and a single coherent source. The relative ratio of the two noise sources is characterized by the ratio of the

coherent contribution to the total noise,

$$R_{\text{coherent}} = P_{\text{coherent}} / (P_{\text{coherent}} + P_{\text{thermal}}), \quad (3.8)$$

where P_{coherent} and P_{thermal} are the noise powers incident on the detectors due to a single coherent and a single thermal source, respectively.

Intensity interferometry experiments rely on two detection events, and so when using dim sources, the integration time T_{int} required to resolve a $g_0^{(2)}$ dip below 1 can be prohibitive. The low probability of coincident counts places demands on the efficiency of the system, which was the limitation that prevented such a demonstration from being done in Ref. [Mirhosseini et al. 2020]. I'll now calculate the dependence of T_{int} on the total efficiency from the qubit output to the optical detection η , the rate at which the qubit pitches single photons Φ_{pitch} , and the desired size of the error bars, i.e. the number of standard deviations E below one we require for the expected value of $g_0^{(2)}$. In an intensity correlation experiment, we'll define a time interval δt within which we'll consider multiple counts to be coincident. For an experiment in which we've been measuring for an integration time of T_{int} , we can then express Eq. 3.6 for $\tau = 0$, as

$$g^{(2)}(0) = \frac{N_{\text{coin}}/N_{\text{int}}}{(\Phi\delta t)^2} = \frac{N_{\text{coin}}}{\Phi^2 T_{\text{int}} \delta t}, \quad (3.9)$$

where N_{coin} is the total number of two-photon coincidences, $N_{\text{int}} = T_{\text{int}}/\delta t$ is the number of intervals we've averaged over, and Φ is average photon detection rate, so $N_{\text{coin}}/N_{\text{int}}$ is the average rate of coincidence counts, and $\Phi\delta t$ is the average count rate. The standard deviation of the number of coincidence counts is then $\Delta N_{\text{coin}} = \Phi \sqrt{g_0^{(2)} T_{\text{int}} \delta t}$, and so the uncertainty in $g_0^{(2)}$ is given by

$$\Delta g_0^{(2)} = \sqrt{\frac{g_0^{(2)}}{\Phi^2 T_{\text{int}} \delta t}}. \quad (3.10)$$

To resolve a dip in $g_0^{(2)}$ below zero by an amount $E\Delta g_0^{(2)}$, it then requires an integration time

$$T_{\text{int}} = \frac{E^2 g_0^{(2)}}{(1 - g_0^{(2)})^2 \Phi^2 \delta t} = \frac{E^2 g_0^{(2)}}{(1 - g_0^{(2)})^2 \eta^2 \Phi_{\text{pitch}}}, \quad (3.11)$$

where in the second equality I have made the substitutions $\Phi = \eta\Phi_{\text{pitch}}$ and $\delta t = 1/\Phi_{\text{pitch}}$, and Φ_{pitch} is the rate at which excitations are pitched from the superconducting qubit. Fig. 3.5(b) shows

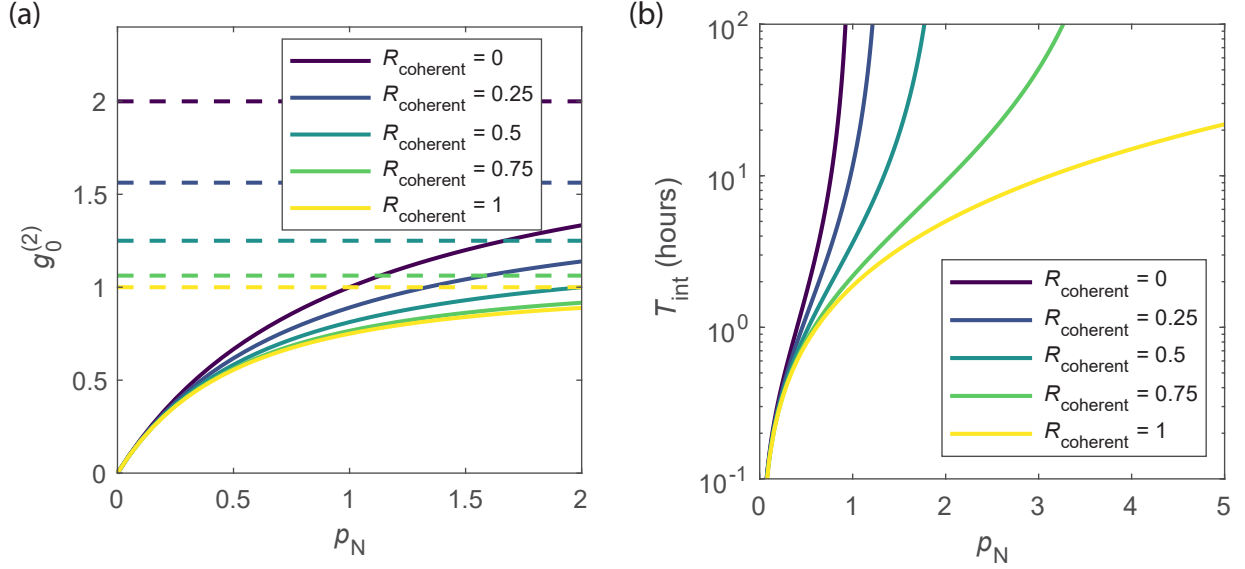


Figure 3.5: **Measurement of $g_0^{(2)}$ with a single-photon source and coherent and thermal noise contributions.** (a) Measured value of $g_0^{(2)}$ vs. the noise-to-signal ratio $p_N = 1/\text{SNR}$ for a single-photon source, assuming noise contributions from a single thermal source and single coherent source, parameterized by R_{coherent} . In this experiment with a single-photon source, p_N is simply the ratio of the total power of the noise sources to the power of the single-photon signal, during the time over which the signal photon is expected to be detected. Vacuum fluctuations do not contribute to p_N because they do not result in counts in direct single photon detection. The dashed lines give the asymptotic value that would be measured in the absence of a signal. (b) Integration time to resolve $g_0^{(2)} < 1$ by 1.5σ vs. p_N , assuming a total efficiency $\eta = 10^{-3}$ and a qubit pitch rate $\Phi_{\text{pitch}} = 1$ kHz, for different statistics of the noise background parameterized by R_{coherent} .

the integration times required to resolve a dip below 1 by $E = 1.5$ standard deviations (1.5σ), for $\Phi_{\text{pitch}} = 1$ kHz and an overall detection efficiency, from qubit output to the optical detection, of $\eta = 10^{-3}$.

3.2.2 Dual-rail interference using a frequency-bin encoding

In preparation for a doubly-heralded two-node protocol, it is possible to test the qubit preparation of a single node and single-photon detection chain with an interference experiment. Fig. 3.6(a) is a schematic depicting the generic protocol, in which the two rails encoding a qubit state are combined on a beamsplitter and detected with a single-photon detector. The figure represents the two rails as two spatial modes, though it would be experimentally more feasible to realize by encoding the qubit state using two time bins or two frequency bins, i.e. photonic modes identifiable by their time or frequency, respectively. Sweeping the phase ϕ of the prepared qubit state $(|0\rangle + e^{i\phi}|1\rangle)/\sqrt{2}$ would lead to interference fringes in the probability of the detector clicking, as shown in Fig. 3.6(b). As this experiment measures a power, it is sensitive to the stability of the chain efficiency. However there is not a firm noise requirement for this demonstration, as the interference fringes are not sensitive to a constant level of additional noise as in the correlation experiments, as it would only contribute to a background offset, leading to increased averaging times to resolve the fringes.

The Barrett and Kok scheme uses time bin qubits as the two rails in the flying qubit encoding. Fig. 3.7(a) shows an interference experiment interfering a superposition of time bins. Here, a single excitation is prepared, and its energy is released equally at two times t_1 and $t_2 = t_1 + \Delta t$. A beamsplitter and delay line allows the two time bins to interfere on the detector at the same time, as shown in Fig. 3.7(a).

However, using time bins for a single-node interference demonstration using our setup is not practical. For a transducer with a bandwidth of ~ 10 kHz, we would need to separate the time bins by at least $\Delta t = 100$ μs for them to remain distinguishable after passing through the transducer. However, in order to recombine these separate rails for interference on the single-photon detector,

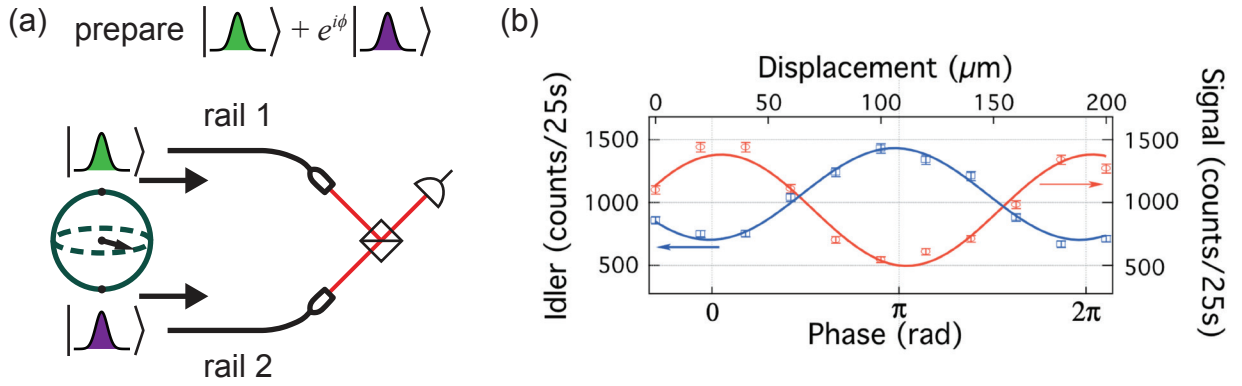


Figure 3.6: **Generic dual-rail interference experiment** (a) Experimental protocol. A single excitation is prepared to occupy both of two rails, here depicted as two spatial modes, with equal probability. The rails are then combined with a beamsplitter operation, and a detected with a single-photon detector. (b) Results from such an experiment using frequency bins [Clemmen et al. 2016]. Note the interference fringes are visible even with a background contribution due to noise counts. Reprinted figure with permission from Stéphane Clemmen, Alessandro Farsi, Sven Ramelow, and Alexander L. Gaeta, *Phys. Rev. Lett.*, **117**, 223601-4, 2016. Copyright 2016 by the American Physical Society.

a delay line of length $l_{\text{delay}} = c\Delta t = 30$ km is required. With fiber losses of 1 dB/km for 1 microm light, a fiber delay would be both impractically long and increase integration times 1000-fold. A free space delay would also be impractical, and require an etalon of LIGO-scale dimensions, both in length and mirror diameter to support the 20 cm beam diameter required to maintain collimation over the delay length.

One alternative would be simply to run at a faster rep rate, such that the two time bins are broadened into one another by the lower transducer bandwidth, but then the two rails would be combined right out of the transducer, rather than propagating independently as they would in the full entanglement generation protocol. A second possibility would be to use frequency bins instead of time bins for our experiment [Clemmen et al. 2016]. Here, a single microwave excitation would be simultaneously mapped onto two propagating modes of distinct frequency, using two concurrent four-wave-mixing processes to create the state $(|10\rangle + e^{i\phi}|01\rangle)/\sqrt{2}$, where here the left and right positions in the ket denote the separate frequency bins [Pfaff et al. 2017; Axline et al. 2018]. The phase of these pumps determine ϕ . For a sense of the frequency scales, the two rails would have bandwidths of approximately $\Gamma_{\text{pitch}} = 2\pi \cdot 1$ kHz, as permitted by the qubit-cavity system, and would be separated by approximately $\Delta\omega = 2\pi \cdot 5$ kHz for simultaneous upconversion within the transducer's bandwidth Γ_t (Fig. 3.7(b)). Because the single-photon detectors are not frequency-resolving, no additional technical preparation is required to realize the beamsplitter operation, such as a frequency shifter to cause the pulses to spectrally overlap. The detector itself is sufficient to observe the interference experiment using frequency bins, as illustrated in Fig. 3.7(c). That the detectors cannot distinguish the excitations is necessary for interference, and is a consequence of the speed of the detectors. We can place a conservative bound on their ability to resolve frequencies by taking the slowest timescale of the detector, its decay time of 100 ns, and calculating the associated frequency uncertainty required to localize an absorption event to a time with that precision, $\Delta f = 10$ MHz. Using the timing jitter of < 100 ps reported in Ref. [Marsili et al. 2013] instead as the uncertainty in the detection time increases Δf to 10 GHz.

In comparison with a $g^{(2)}$ experiment, the integration time required for a dual-rail interference

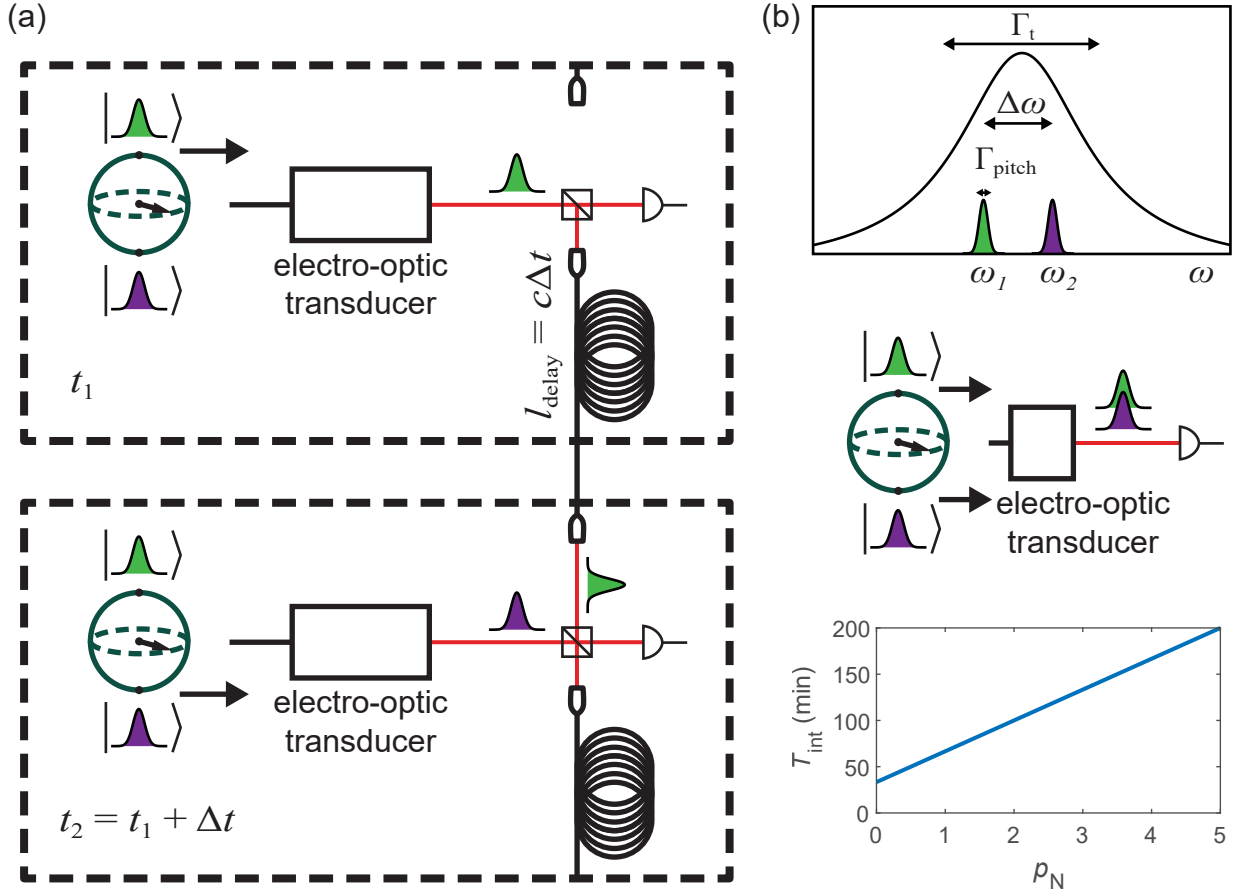


Figure 3.7: **Time- and frequency-bin interference experiments** (a) Schematic for experimental setup to observe interference of a single photon in a superposition of time bins. The superconducting node is used to release a single excitation with equal probability in each of two time bins, t_1 (green) and $t_2 = t_1 + \Delta t$ (purple). The spatial mode is split on a beamsplitter, sent through a delay line, and recombined before detection, such that the two time bins overlap and interfere as illustrated in the bottom panel. (b) Frequency-bin experiment. Top: Frequency diagram for frequency-bin interference experiment. Frequency bins ω_1 and ω_2 with bandwidth Γ_{pitch} , are separated by $\Delta\omega$ such that both can be upconverted with a transducer of bandwidth Γ_t . Center: Schematic of experimental setup for to observe interference between frequency bins. Both frequency bins are simultaneously populated with equal probability of holding the same single excitation, are upconverted, and collected on a single-photon detector. Bottom: Integration time for a frequency-bin interference experiment in order to resolve an interference pattern with 10% error bars, given $\eta = 10^{-3}$, $T_{\text{rep}} = 1$ ms, and $N_\phi = 20$.

experiment is favorable. Whereas coincidence counts in an intensity correlation experiment scale with $1/\eta^2$, the direct detection of the dual-rail interference experiment scales with $1/\eta \cdot n_\phi$, where n_ϕ is the number of relative phase values desired for the demonstration. The linear scaling of integration time is shown in Fig. 3.7(d), again using $T_{\text{rep}} = 1 \text{ ms}$ and $\eta = 10^{-3}$, as well as $n_\phi = 20$. The integration times shown are those required to resolve interference with 1σ error bars $1/10$ the size of the interference fringe.

Though a frequency-bin encoding is desirable to use for a single-qubit interference demonstration, it is not the encoding of choice for a doubly-heralded protocol. The motivation for using frequency bins is the fact that excellent separation achievable with time binning makes an interference experiment infeasible. Conversely, it would be very technically challenging to resolve the two rails in a Barrett and Kok-style entanglement protocol using frequency bins. A sharp frequency resolving filter would be required to route the two rails to separate detectors on which to perform the doubly-heralded protocol, and crosstalk between the rails would lead to sharply reduced fidelity. Therefore, a frequency-bin encoding is primarily of use for demonstrating the necessary state control in the superconducting node and the preservation of the phase of the superposed state during transduction and optical propagation.

3.2.3 Testing single-optical-photon counting without a qubit source.

It is important to test the single-photon counting apparatus, even before a superconducting qubit capable of pitching single microwave photons is available. The transducer and a coherent microwave signal could be used to evaluate how close this portion of the network is to being ready for a quantum demonstration like a nonclassical $g^{(2)}$ measurement with a quantum source at the transducer input, or a heralding experiment. It also provides direct information about the absolute level and the statistics of the noise that would limit such quantum demonstrations, which would be an interesting measurement [Kumar et al. 2023].

Methods of calibrating p_N

In this chapter I have been evaluating potential protocols based on p_N , the noise count rate

relative to the count rate of potential signal photons, achievable by a network. In past experiments, we've used heterodyne detection to benchmark the relatively narrow-band noise performance of our transducer across its transduction bandwidth. Single-photon detectors, however, have different considerations to account for when incorporating them into a networking protocol. For example, they won't click due to vacuum fluctuations of the light field, but are sensitive to photons across a very wide bandwidth. This means that even a small amount of white noise above the vacuum fluctuations as measured by heterodyne detection, could integrate up across the bandwidth of the single-photon detectors to generate a non-negligible noise count rate. The relevant bandwidth should be limited to the 30 kHz effective bandwidth of the filter cavity system. Over this bandwidth, an absolute noise level of approximately -141 dBc/Hz corresponds to $p_N = 1$, assuming similar optomechanical coupling rate to the device in Ref. [Brubaker et al. 2022], using the convention of one-sided double-sideband spectral densities (see Sec. 6.2.1). For example, this means that a technical noise level of -144 dBc/Hz with an otherwise noiseless setup and transducer would allow for the averaging times associated with $p_N = 0.5$ in Figs. 3.5 and 3.7, and that a technical noise level of -147 dBc/Hz is required for the $p_N = 0.25$ required for entanglement fidelity of greater than 1/2 for the Barrett and Kok protocol. Though we have an expectation of what levels of noise are acceptable, it is important to measure the noise count rates directly with the single-photon detection apparatus while operating the transducer and confirm that we get the count rates that we expect from the noise sources we have calibrated in our past work.

To translate a count rate to a value of p_N , i.e. to compare the noise to a quantum signal at the channel's input, a calibration of the overall efficiency η is needed. The most direct calibration is a single-photon source to which the noise can be compared, as done in Ref. [Mirhosseini et al. 2020], illustrated in Fig 3.8(a).

Heterodyne detection allows for an additional possibility: the known SNR from dispersive readout of the state of a qubit can be used to calibrate the total quantum efficiency of the chain, as we used in Ref. [Delaney et al. 2022] (Fig. 3.8(b)). This has the advantage of allowing for amplification by increasing the photon number of the readout pulse to overcome noise and decrease




Figure 3.8: **Methods for calibrating p_N** (a) Direct measurement of a single-photon source can calibrate the noise-to-signal ratio p_n , as done in Ref. [Mirhosseini et al. 2020]. (b) The SNR of dispersive readout of a qubit can calibrate p_N , as in Refs. [Delaney et al. 2022; Hatridge et al. 2013]. (c) Sideband asymmetry thermometry can also calibrate the noise of a single quantum at the plane of the mechanical resonator, in order to calibrate p_N , as in [Brubaker et al. 2022].

averaging times.

Sideband asymmetry thermometry (see Ch. 4) can also be used to calibrate the size of a single excitation of the mechanical mode in the transducer, provided that the technical noise on the light reading out the thermomechanical motion is low enough, or that it is correctly accounted for [Safavi-Naeini et al. 2013; Jayich et al. 2012]. The amplitude equivalent to a single mechanical excitation can then be used to calibrate the power of the laser technical noise, and referred to the input of the transducer, as we did in Ref. [Brubaker et al. 2022], as illustrated in Fig. 3.8(c). We used heterodyne detection for that calibration, and sideband asymmetry thermometry can also be used with single-photon detection, as in Ref. [Galinskiy et al. 2020], which would provide the calibration of p_N we are looking for, i.e. including all the contributions to the noise counts. In principle, sideband asymmetry could also be used to calibrate an incident coherent microwave signal to compare the noise count rate with, provided it results in a mechanical amplitude of approximately 1 phonon^{1/2} or smaller.

Intensity correlations with a coherent microwave input

Once we have calibrated the noise that the single-photon detection apparatus is sensitive to, we could perform intensity correlation experiments to tell us about the statistics of that noise, and therefore about its origin. In Ref. [Kumar et al. 2023], Jonathan Simon and David Schuster’s collaboration measured the $g^{(2)}$ of their transducer noise to verify its thermal statistics, rather than being the result of contamination from an off resonant coherent drive. They then swept the power of a coherent millimeter wave input and observed $g_0^{(2)}$ decrease as it contributed coherent statistics to the detection events.

Measuring the intensity correlations of the noise from our membrane-based transducer could interrogate a more nuanced noise model, as there are several somewhat comparable contributions to the noise. Fig. 3.9(a) shows the setup for this experiment with an optional coherent signal injected into the microwave port of the transducer. As in Fig. 3.4, the contributions to the single-photon detection events are indicated by color. The contributions are labeled by their count rate referred to the transducer input, with N_{sig} corresponding to the input microwave coherent

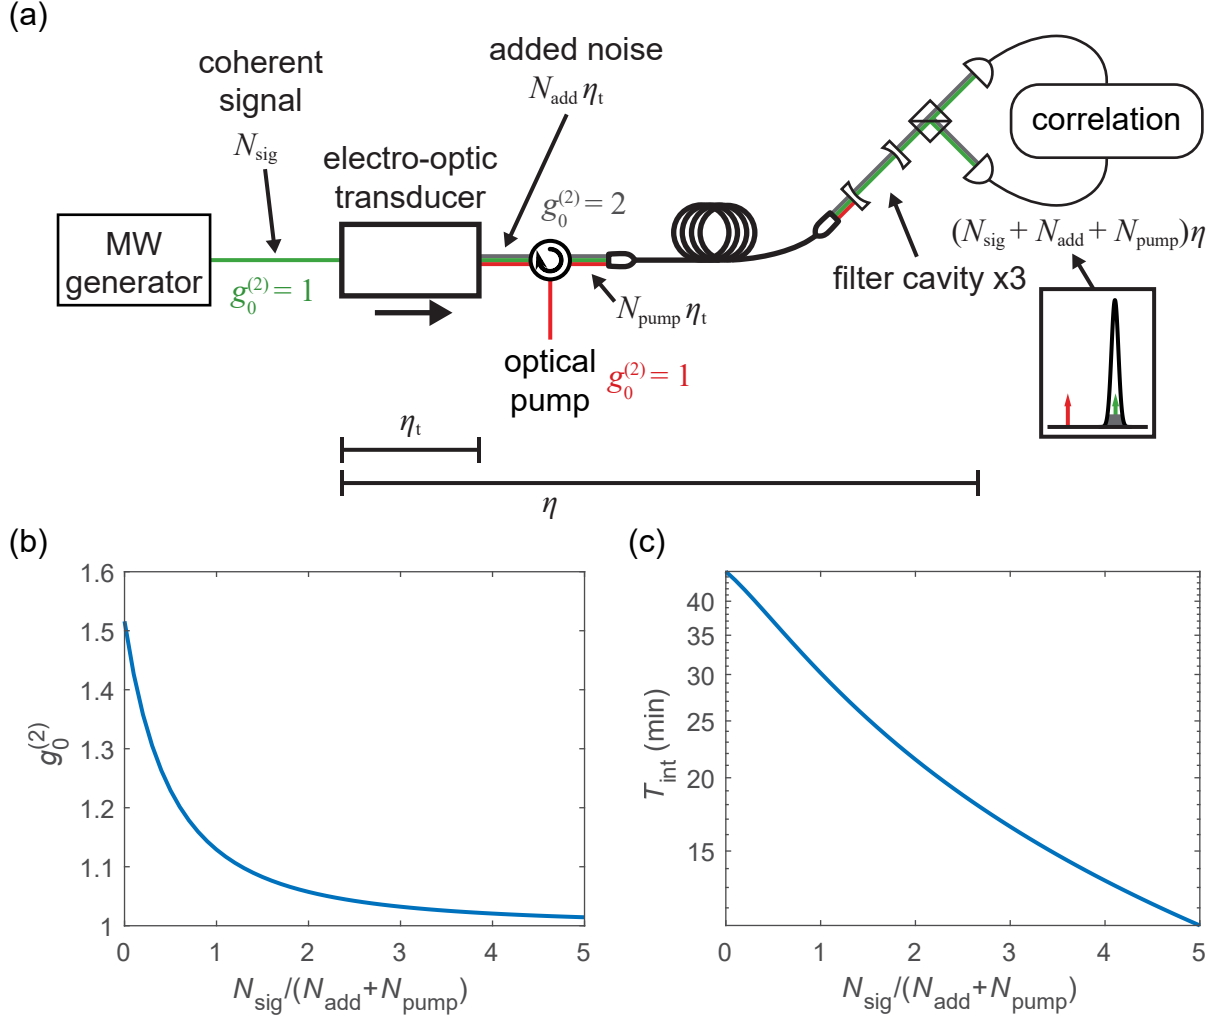


Figure 3.9: **Single-photon counting with a coherent microwave source.** (a) Measurement setup. The intensity correlation setup shown in Fig. 3.4 is replicated with a coherent microwave tone at the input to the transducer. A total count rates $(N_{\text{sig}} + N_{\text{add}} + N_{\text{pump}})\eta$ is detected on the single-photon counting apparatus, where N_{sig} , N_{add} , and N_{pump} are the count rates due to the signal, transducer added noise, and coherent leakage of the optical pump, respectively, referred to the input of the transducer. Inset illustrates the frequency of the count sources: the signal and noise are passed by the filter cavities (black Lorentzian) and the optical pump is suppressed. (b) Model of $g_0^{(2)}$ vs. the coherent input signal power. Please see Table 3.1 and Sec. 3.2.3 for the model assumptions. (c) Integration times required for the measurements in (b) with 5% error bars, assuming an optomechanical measurement efficiency of $\eta_{\text{OM}} = 10^{-2}$.

signal, N_{add} corresponding to the total contribution from the transducer added noise (including technical noise from the optical pump, as in Fig. 3.4), and N_{pump} corresponding to the leakage of the coherent pump through the filter cavity system. The average count rate of the detector is then $(N_{\text{sig}} + N_{\text{add}} + N_{\text{pump}})\eta$, with η as the total efficiency, and output-referred added noise is $N_{\text{add}}\eta_t$, where η_t is the transducer efficiency. Similarly, $N_{\text{pump}}\eta_t$ is the output-referred count rate due to the imperfectly suppressed coherent pump, though it is worth noting that this does not correspond to the number of photons in the pump at this point in the chain, because of the frequency-dependent response of the filter cavity system.

In the absence of an input signal, the measured $g_0^{(2)}$ would be 2 for a single source of thermal noise counts, as seen in Ref. [Kumar et al. 2023]. This would be the case if, for example, our noise budget were dominated by $n_{\text{th}}\gamma_m$, rate at which thermal phonons enter the mechanical mode. Coherent pump leakage would bring the value of $g_0^{(2)}$ towards 1. Additional uncorrelated sources of thermal photons, such as technical noise on the laser or backaction due to an imperfect transducer optical cavity (see Secs. 3.3 and 4.1.3), will also bring $g_0^{(2)}$ down from a value of 2. Therefore, a measurement of $g_0^{(2)}$ alone would be useful in tracking down the various sources of additional counts that would limit a heralding experiment or a violation of $g_0^{(2)} \geq 1$, in addition to simply knowing the average rate of these counts from the calibration of p_N .

Sweeping the power of the coherent signal N_{sig} would bring the measured $g_0^{(2)}$ towards 1, according to Eq. 3.7, as shown in Fig. 3.9(b). This model contains contributions from coherent pump leakage, technical laser noise, and contributions from the mechanical mode's thermal environment and optomechanical backaction at rate $n_{\text{min}}\Gamma_o$ (see Sec. 4.1.3), with operating parameters given in Table 3.1. Here the pump power is calculated from Γ_o and Eq. 6.3 using the transducer parameters in Table 6.2. For these operating parameters, the contribution from laser noise and coherent pump leakage taken together are a factor of 10 lower the contribution to counts from the rate of thermal phonons entering the mechanical mode $n_{\text{th}}\gamma_m/2$, where I divide by two assuming that half of the thermal photons are routed out the microwave port of the transducer. The contribution from optomechanical backaction $n_{\text{min}}\Gamma_o$ is twice the thermal contribution—here I don't

divide by two assuming an equal contribution from backaction due to the microwave pump. Observing this dependence of $g_0^{(2)}$ on the microwave signal power would be an additional check on the calibration of p_N and the accuracy of the sideband-asymmetry calibration of the amplitude of the coherent drive. Fig. 3.9(c) shows the integration times T_{int} calculated for a desired fractional uncertainty $\Delta N_{\text{coin}}/N_{\text{coin}} = 0.05$ from $T_{\text{int}} = \left(\frac{\Delta N_{\text{coin}}}{N_{\text{coin}}}\right)^2 \Phi^2 \eta_{\text{OM}}^2 g^{(2)} \delta t$, where I've introduced η_{OM} , the optomechanical measurement efficiency. It is worth emphasizing that this demonstration has a more lenient requirement on the total system noise counts compared with demonstrations of non-classicality using a single-photon source, and only requires that it be close to the count rate due to a microwave tone weak enough to be calibrated by the sideband asymmetry of the mechanical modulation, such that the microwave tone can affect the statistics of the total noise.

technical noise	$\Gamma_o/2\pi$	R_{coherent}	n_{th}	$\gamma_m/2\pi$	$n_{\text{min,o}}$	η_{OM}	$\Delta N/N$
-154 dBc/Hz	1 kHz	0.6	1000	0.1 Hz	0.1	10^{-2}	0.05

Table 3.1: Parameters for calculation of $g^{(2)}$ and T_{in} when sweeping coherent microwave input.

3.3 Nonclassical intensity and higher-order correlations with pure optomechanical system

It is also in principle possible to evaluate the part of the chain including the optomechanical system and the single-photon counting apparatus by performing a $g^{(2)}$ experiment similar to that done by the Groblacher group in Ref. [Hong et al. 2017], in which they herald a single-phonon state using single-photon detection. They then read out that heralded Fock state optomechanically, and perform a $g^{(2)}$ experiment to verify its quantum statistics. This experiment does not require the microwave portion of the experiment, but would demonstrate the ability to unambiguously measure a quantum signature with the optical measurement chain, which places stringent requirements on the noise introduced on the optical side of the experiment. Furthermore, because the single-photon source is the mechanical mode, the constraints on technical noise will be relaxed, as the approximately 8 dB of loss due to propagation of the microwave single photon to the transducer

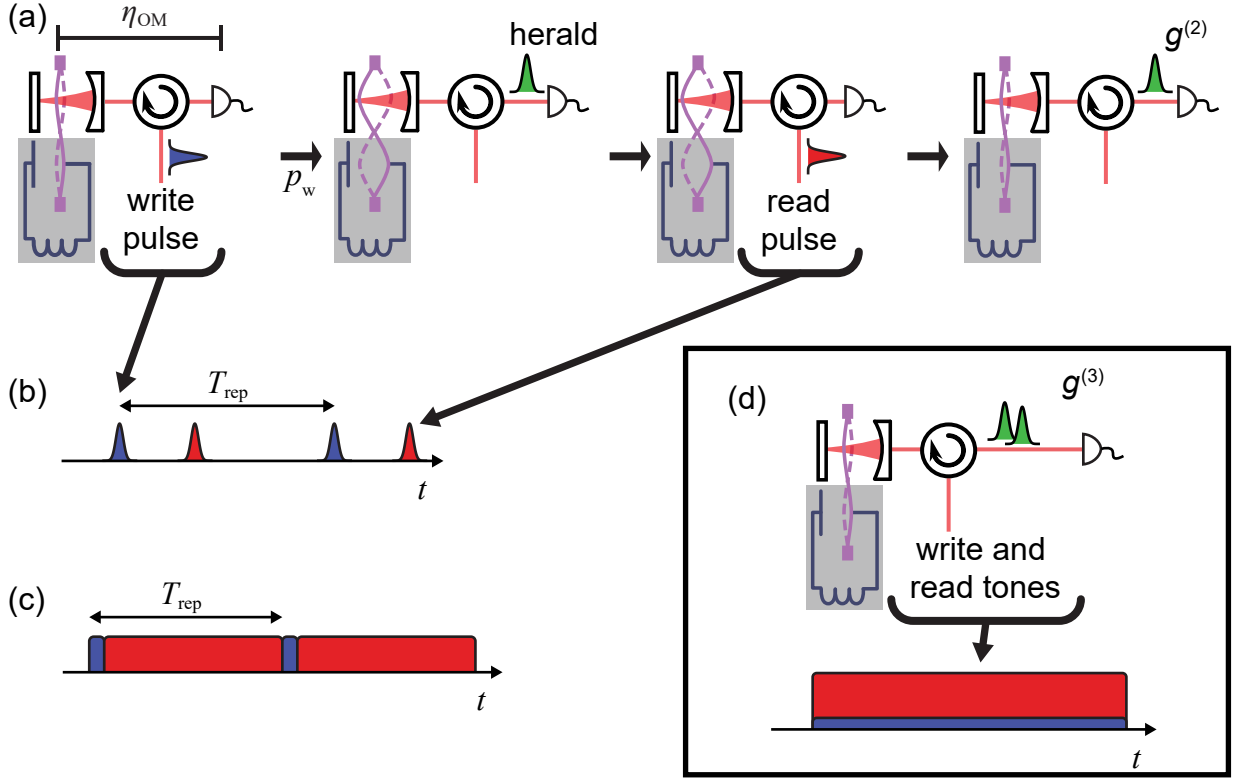


Figure 3.10: **Optomechanical intensity correlation experiments** (a) Pulsed protocol for heralding and detection of single-phonon Fock state in mechanical mode [Hong et al. 2017]. First a write pulse (blue) is routed towards the optomechanical cavity. Then, with probability p_w , this creates an entangled pair of excitations: a phonon in the mechanical mode (purple) and a propagating optical photon (green), and detection of the optical photon heralds the created phonon in the mechanical mode. Next, a read pulse (red) is directed to the optomechanical system, converting the phonon to an optical photon (green), directed to single-photon detection for a $g^{(2)}$ measurement. (b) Timing diagram for the pulse sequence. (c) Timing diagram for membrane-based system, illustrating a duty cycle close to 100%. (d) The continuous experiment proposed in Ref. [Børkje, Massel, and Harris 2021]. Two continuous write and read pumps drive the optomechanical device simultaneously. When the read tone has greater optomechanical damping rate than the write tones antidamping rate, photons arrive in pairs, leading to an observed value of $g^{(3)} < 1$, a nonclassical signature allowed by a mechanical mode near its ground state.

and the transducer efficiency will not attenuate the phononic signal (see Sec. 6.2.1).

Their experimental protocol is represented in Fig. 3.10(a). First, a weak two-mode-squeezing “write” pulse is applied to an optomechanical system to prepare an entangled pair of excitations in the mechanical mode and the propagating optical field with low probability p_w , preparing the state $|00\rangle + \sqrt{p_w}|11\rangle + p_w|22\rangle + \mathcal{O}(p_w^{3/2})$. Detection of the optical excitation by single-photon detection then heralds a single phonon in the mechanical mode with a fidelity of approximately $1 - 2p_w$. Next, a stronger optical readout pulse swaps the mechanical state onto the optical field, again routed towards the single-photon detection apparatus. By post-selecting on the heralded events, a $g^{(2)}$ measurement of the optical field due to the read pulses then revealed a value below the classical bound of 1.

Though it seems plausible to recreate this experiment as-is with a membrane-based transducer in our setup, there are technical details that make another proposal, which can be viewed as the continuous-wave analogue, as a more exciting prospect [Børkje, Massel, and Harris 2021]. Nonetheless, it is illuminating to compare a membrane-based platform with that in Ref. [Hong et al. 2017] in the context of the pulsed experiment just described. The platform used in their work is an optomechanical crystal (OMC) with mechanical resonant frequency 5.25 GHz. Though OMCs permit higher optomechanical damping rates than a membrane optomechanical system, optical absorption increases the mechanical occupation with each pulse, so, although the write and read pulses had 32 ns widths, the repetition rate was limited to a long repetition time of $T_{\text{rep}} = 50 \mu\text{s}$ to allow the resonator to rethermalize to the base plate of the dilution refrigerator, as illustrated in the pulse timing diagram, Fig. 3.10(b). In contrast, a membrane-based transducer in such an experiment would benefit from operating at a duty cycle close to 100% (Fig. 3.10(c)), as optical absorption is not a limitation, and the sideband cooling from the readout pump would be required to maintain the mechanical mode close to its ground state (see Sec. 4.1.3). As a consequence of using the readout pump to cool the mechanical mode, the experimental signal competes with a source of thermal photons that arrive at a rate given by the product of the environmental occupation n_{th} and the intrinsic loss rate of the mechanical mode γ_m . For example, in Ref. [Brubaker et al. 2022]

our transducer's mechanical mode operated at $n_{\text{th}} = 1000$, and with $\gamma_{\text{m}} = 2\pi \cdot 100$ mHz, so these thermal photons arrive at an average rate of approximately 600 cps. This source of noise limits the minimum p_{w} that can be used for a given rep rate, in order for the signal to dominate. Assuming $T_{\text{rep}} = 16$ μs , permitted by an optomechanical damping rate of $2\pi \cdot 10$ kHz for the readout pulse, p_{w} would need to be set to 0.1 in order for the transduced signal phonons to dominate over this thermal noise source by a factor of 10. Thus, the infidelity from preparing the $|2\rangle$ state in the mechanical mode rather than the $|1\rangle$ state would be 20%, rather than 1.2% as in [Hong et al. 2017]. Another difference to note between these mechanical platforms is that because the read pulse widths in [Hong et al. 2017] are comparable to the recovery time of single-photon detectors, they required a traditional Hanbury Brown and Twiss setup including a beamsplitter and two detectors, whereas a membrane with lower optomechanical damping rates would be able to use a single detector.

In comparison with the $g^{(2)}$ measurement with a microwave source capable of deterministically sending single excitations, the round of heralding in this experiment increases the number of measurements by $1/\eta p_{\text{w}}$, the expected number of attempts before detecting a photon scattered by the write pulse. However, the total efficiency $\eta = \eta_{\text{om}}$ is greater in this purely optomechanical experiment because it does not include microwave losses during propagation or transduction. For our system, this improvement in η would be roughly a factor of 10, depending on the device. Furthermore, as the transducer bandwidth is limited by the electromechanical circuit performance, the optomechanical experiment could be operated at a faster repetition rate of 10 kHz, rather than 1 kHz. Taking these three effects together with a value of $p_{\text{w}} = 0.1$, this purely optomechanical experiment would require ten times the number of measurements as a $g^{(2)}$ experiment with a superconducting qubit, but an equal integration time to that shown in Fig. 3.5(b), due to the increased repetition rate.

Finally, there are two remaining complications that preclude the the pulsed heralding and detection of single-phonons demonstration from being performed with devices we have optimized for recent transducer demonstrations. The first complication is from additional optomechanical backaction due to increasing the optomechanical damping rate to 10 kHz, which was required

for the signal to dominate over the thermal noise. This backaction arises from imperfect sideband resolution (see Sec. 4.1.3), which modifies the perfect read and write pulses described above. Rather than purely swapping the mechanical state to the light in the read pulse, it would also contribute an additional source of photons at a rate $n_{\min,o}\Gamma_o$, which would be a rate of greater than 3 keps for our system for optical cavity loss rates of $\kappa_o/2\pi > 1.3$ MHz. We could reduce the impact of backaction by reducing the cavity loss rate, e.g. by increasing the cavity length.

The second complication arises from the narrow bandwidth of the filter cavities used to filter the pump. Their effective bandwidth is 30 kHz, corresponding to a timescale of $(2\pi \cdot 30 \text{ kHz})^{-1} \approx 5 \text{ } \mu\text{s}$, a significant fraction of $T_{\text{rep}} = 16 \text{ } \mu\text{s}$, so an experiment running at this rate would result in the heralding and heralded photons broadening their temporal envelope, and overlapping with each other.

While there are difficulties in applying the protocol of Ref. [Hong et al. 2017] to our membrane-based system, Ref. [Galinskiy et al. 2023] details a different pulsed experiment that uses a membrane-optomechanical system to demonstrate violation of a Cauchy-Schwarz inequality indicating non-classical correlations between the mechanical mode and light.

The protocol proposed in Ref. [Børkje, Massel, and Harris 2021] is the continuous analogue of the pulsed protocol of Ref. [Hong et al. 2017], and more naturally lends itself to our setup than that protocol. It instead uses two continuous pumps that are unbalanced, and measures the third-order correlation function,

$$g^{(3)}(\tau) = \frac{\langle \hat{a}^{\dagger 2}(t) \hat{a}^{\dagger}(t+\tau) \hat{a}(t+\tau) \hat{a}^2(t) \rangle}{\langle \hat{a}^{\dagger}(t) \hat{a}(t) \rangle^3}, \quad (3.12)$$

as illustrated in Fig. 3.10(d). The analogous nonclassical bound is a measurement of $g^{(3)}(0) < 1$, i.e. that given two clicks, it is much less likely to measure a third. By driving the optomechanical system with simultaneous write and read tones, the two processes described above happen simultaneously; the write tone tends to create phonon-photon pairs, and the read tone swaps phonons to signal photons. When the read tone dominates, photons will then tend to be created in pairs, leading to a violation of the classical bound $g^{(3)}(0) \geq 1$, provided sufficient sideband resolution. Given a

mechanical frequency of $\omega_m/2\pi = 1.5$ MHz, as for our transduction membrane mode, an optical cavity loss rate of approximately less than $2\pi \cdot 800$ kHz is required to violate this bound. The integration times of this experiment would scale similarly to those in the pulsed experiment, as measurement of $g^{(3)}$ also relies on three detection events, happening at a rate limited by the reduced power of the write tone. There may also be advantages in averaging times by introducing a small additional detuning to the read and write tones, such that the signals from the two tones are detuned by 2δ , and looking for a violation of $g^{(3)}(\pi/2\delta) \geq 1$ [Børkje, Massel, and Harris 2021].

To summarize this chapter, I have presented a number of experimental demonstrations for a transducer device, and explored their relative demands on transducer performance and additional requirements on the experimental setup, such that a demonstration can be chosen to align with state of the experiment at a given time. In the next chapter, I discuss in greater detail how the membrane-based transducer works, and various technical improvements towards reaching the requirements for a networking demonstration.

Chapter 4

Fabry-Pérot cavities in membrane-based transducers

My secret, the secret of my uninterrupted financial victories in a period that has witnessed so many crises and market crashes and bankruptcies, has always been this: that I never thought directly of money, business, profits, but only of the angles of refraction established among shining surfaces variously inclined.

—Italo Calvino, *If on a winter's night a traveler*

In this chapter I lay out the operating principles of the transducer design that the Regal and Lehnert groups have opted to use, in which a mechanical mode mediates an interaction between a superconducting LC circuit and a Fabry-Pérot optical cavity. I then detail design choices and modifications pertaining to the optical cavity that have improved transducer performance and reliability, with an emphasis on how I think about the design space. The presentation of these improvements is roughly chronological. I conclude the chapter detailing two aspects of mechanical design: a consideration of membrane thermalization in a cryogenic environment due to the high circulating powers in a cavity optomechanical system, and the serendipitous observation, permitted by optical readout, of a low-loss mechanical resonance in a collapsed electromechanical device. Details of the procedure I have developed for assembling optical cavities can be found in Appendix A.

4.1 Operating theory of a mechanical transducer

4.1.1 The optomechanical and electromechanical interactions

The electro-optomechanical transducer is comprised of a superconducting microwave-frequency LC circuit and a Fabry-Pérot optical cavity simultaneously coupled to the same mechanical

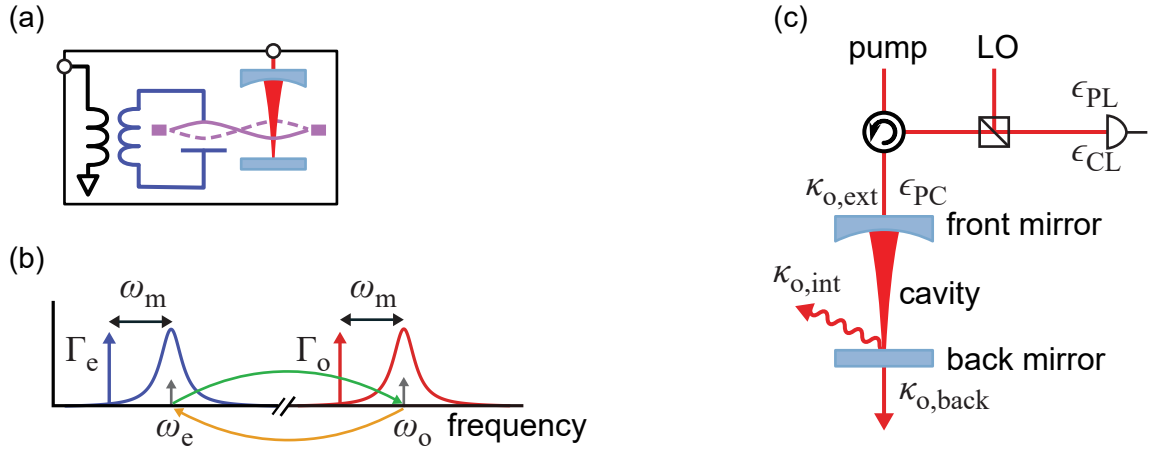


Figure 4.1: **Principles of transducer operation** (a) Schematic diagram of the transducer illustrating the vibrational mode of the membrane (purple) coupling to the optical (red) and microwave (blue) resonator modes. (b) Frequency diagram. The microwave (blue) and optical (red) pumps are red-detuned from their respective resonator frequency, ω_e and ω_o , by the mechanical frequency ω_m . The electromechanical and optomechanical damping rates Γ_e and Γ_o scale with the pump powers. With matched damping rates that dominate the intrinsic mechanical damping rate γ_m , efficient upconversion (green arrow) and downconversion (yellow arrow) result. (c) Cavity loss rates $\kappa_{o,ext}$ from front mirror mirror transmission, $\kappa_{o,back}$ from the back mirror transmission, $\kappa_{o,int}$ from scattering/absorption, and modematchings between the pump (P), cavity (C), and LO (L) modes.

mode of a silicon nitride membrane, shown schematically in Fig. 4.1(a). The system is described by the Hamiltonian

$$\hat{H} = \hbar\omega_o\hat{a}^\dagger\hat{a} + \hbar\omega_e\hat{b}^\dagger\hat{b} + \hbar\omega_m\hat{c}^\dagger\hat{c} + \hat{H}_{\text{OM}} + \hat{H}_{\text{EM}}, \quad (4.1)$$

where ω_o , ω_e and ω_m are the frequencies of the optical Fabry-Pérot cavity, the microwave LC resonator, and the intermediary mechanical mode, respectively, and \hat{a} , \hat{b} and \hat{c} are the corresponding annihilation operators. The optical and microwave electromagnetic modes interact with the mechanical mode respectively via the electromechanical and optomechanical interactions, in which a change in the position of the mechanical degree of freedom results in a frequency shift of the electromagnetic resonator:

$$\hat{H}_{\text{OM}}/\hbar = g_o\hat{a}^\dagger\hat{a}(\hat{c} + \hat{c}^\dagger) \text{ and } \hat{H}_{\text{EM}}/\hbar = g_e\hat{b}^\dagger\hat{b}(\hat{c} + \hat{c}^\dagger), \quad (4.2)$$

where g_o and g_e are respectively the single-photon optomechanical and electromechanical coupling rates, and must be real-valued for hermiticity. Applying optical and microwave pumps detuned from the respective resonator by Δ_o and Δ_e displaces the associated mode by \bar{a} and \bar{b} . Linearizing about these displacements by moving to the rotating frame of the pumps, the Hamiltonian becomes

$$\hat{H}_{\text{lin}}/\hbar = -\Delta_o\hat{a}^\dagger\hat{a} - \Delta_e\hat{b}^\dagger\hat{b} + \omega_m\hat{c}^\dagger\hat{c} + g_o\bar{a}(\hat{a}^\dagger + \hat{a})(\hat{c} + \hat{c}^\dagger) + g_e\bar{b}(\hat{b}^\dagger + \hat{b})(\hat{c} + \hat{c}^\dagger), \quad (4.3)$$

where I have assumed \bar{a} and \bar{b} to be real without loss of generality. By red-detuning the pumps by $-\omega_m$, as shown in Fig. 4.1(b), this Hamiltonian then describes three resonators of equal energy, that can exchange excitations resonantly at rates $g_o\bar{a}$ and $g_e\bar{b}$ via the beamsplitter interaction terms $g_o\bar{a}\hat{a}^\dagger\hat{c} + g_o\bar{a}\hat{a}\hat{c}^\dagger$ and $g_e\bar{b}\hat{b}^\dagger\hat{c} + g_e\bar{b}\hat{b}\hat{c}^\dagger$.

4.1.2 Efficient transduction through a mechanical mode

To evaluate the efficiency of this exchange of excitations, I'll now introduce decay rates for each of the resonators. The mechanical mode has an internal energy dissipation rate, denoted γ_m . The microwave mode has total dissipation rate κ_e , comprised of the rate of external coupling to a transmission line $\kappa_{e,\text{ext}}$ and internal dissipation rate $\kappa_{e,\text{int}}$. Similarly, the optical mode has

total decay rate κ_o , comprised of the external coupling rate through the more transmissive front mirror $\kappa_{o,\text{ext}}$, and loss rate due to the sum of scattering/absorption and transmission through the back mirror, $\kappa_{o,\text{int}} + \kappa_{o,\text{back}}$, as indicated in Fig 4.1(c). In the weak coupling regime ($g_o \ll \kappa_o = \kappa_{o,\text{ext}} + \kappa_{o,\text{int}} + \kappa_{o,\text{back}}$), the optomechanical interaction then contributes net damping to the mechanical mode given by

$$\Gamma_o = g_o^2 |\bar{a}|^2 \left(\frac{\kappa_o}{\kappa_o^2/4 + (\omega_m + \Delta_o)^2} - \frac{\kappa_o}{\kappa_o^2/4 + (\omega_m - \Delta_o)^2} \right), \quad (4.4)$$

where $|\bar{a}| = \sqrt{(P_o/\hbar\omega_{p,o}) \epsilon_{\text{PC}} \kappa_{o,\text{ext}} / (\kappa_o^2/4 + \Delta_o^2)}$, $\omega_{p,o} = \omega_o - \Delta_o$ is the optical pump frequency, and ϵ_{PC} is the modematching of the incident pump beam to the optical cavity mode. The second term in Eq. (4.4) represents anti-damping from imperfect suppression of the two-mode squeezing terms $\hat{a}^\dagger \hat{c}^\dagger$ and $\hat{a} \hat{c}$ in Eq. (4.3); it is suppressed relative to the first term at optimal detuning $\Delta_o = -\omega_m$ in the resolved sideband limit $\kappa_o \ll 4\omega_m$. Likewise, the electromechanical damping rate for $g_e \ll \kappa_e = \kappa_{e,\text{ext}} + \kappa_{e,\text{int}}$ is

$$\Gamma_e = g_e^2 |\bar{b}|^2 \left(\frac{\kappa_e}{\kappa_e^2/4 + (\omega_m + \Delta_e)^2} - \frac{\kappa_e}{\kappa_e^2/4 + (\omega_m - \Delta_e)^2} \right), \quad (4.5)$$

where $|\bar{b}| = \sqrt{(P_e/\hbar\omega_{p,e}) \kappa_{e,\text{ext}} / (\kappa_e^2/4 + \Delta_e^2)}$ and $\omega_{p,e} = \omega_e - \Delta_e$ is the microwave pump frequency.

The transmission of a narrowband signal from the external port of the microwave resonator through the transducer to the external port of the optical cavity is governed by the scattering parameter

$$\mathcal{S}_{oe}(\omega) = \sqrt{\mathcal{A}_e \mathcal{A}_o \frac{\kappa_{e,\text{ext}}}{\kappa_e} \frac{\kappa_{o,\text{ext}}}{\kappa_o}} \frac{\sqrt{\Gamma_e \Gamma_o}}{\Gamma_T/2 - i(\omega - \omega_m)}, \quad (4.6)$$

where ω is the detuning of both the input signal from the microwave pump and the detuning of the output signal from the optical pump, and $\Gamma_T = \Gamma_e + \Gamma_o + \gamma_m$ defines the transduction bandwidth. Gain factors \mathcal{A}_e and \mathcal{A}_o arise from the optomechanical two-mode squeezing terms in Eq. (4.3), with $\mathcal{A}_o = -((\kappa_o/2)^2 + (\Delta_o - \omega_m)^2) / 4\Delta_o\omega_m$ and \mathcal{A}_e defined analogously [Andrews et al. 2014]. These factors are undesirable because any transducer gain is necessarily accompanied by added noise [Caves 1982]. With optimally red-detuned pumps, the gain factors reduce to $\mathcal{A}_o = 1 + (\kappa_o/4\omega_m)^2$ and $\mathcal{A}_e = 1 + (\kappa_e/4\omega_m)^2$, and the transducer gain approaches unity in the resolved

sideband limit. The transmission of a narrowband signal modematched to the optical cavity through the transducer to the external port of the microwave resonator, $\mathcal{S}_{\text{eo}}(\omega)$, is formally identical to $\mathcal{S}_{\text{oe}}(\omega)$.

As in Ref. [Andrews et al. 2014], I use a conservative definition of the transducer efficiency that accounts for the energy lost in each of the three modes involved in the transduction process as well as imperfect optical cavity modematching. This modematching factor is not included in Eq. 4.6, and in general is not the same in upconversion and downconversion. In downconversion the relevant modematching factor is the modematching ϵ_{PC} of the incident pump beam to the cavity, while in upconversion the relevant factor is the modematching ϵ_{CL} of the light emerging from the cavity to the LO beam, as indicated in Fig 4.1(c). Following Ref. [Brubaker et al. 2022], I redefine the optical input and output ports to absorb modematching factors, so that $\mathcal{S}_{\text{oe}}(\omega_{\text{m}}) \rightarrow \sqrt{\epsilon_{\text{CL}}}\mathcal{S}_{\text{oe}}(\omega_{\text{m}})$ and $\mathcal{S}_{\text{eo}}(\omega_{\text{m}}) \rightarrow \sqrt{\epsilon_{\text{PC}}}\mathcal{S}_{\text{eo}}(\omega_{\text{m}})$. The upconversion efficiency for an incident signal on resonance is then defined as $\eta_{\text{t,up}} = |\mathcal{S}_{\text{oe}}(\omega_{\text{m}})|^2 / \mathcal{A}_{\text{o}}\mathcal{A}_{\text{e}}$, and $\eta_{\text{t,down}} = |\mathcal{S}_{\text{eo}}(\omega_{\text{m}})|^2 / \mathcal{A}_{\text{o}}\mathcal{A}_{\text{e}}$ as the downconversion efficiency.

The bidirectional efficiency can then be defined as

$$\eta_{\text{t}} = \sqrt{\eta_{\text{t,up}}\eta_{\text{t,down}}} = \eta_{\text{M}} \frac{4\Gamma_{\text{e}}\Gamma_{\text{o}}}{\Gamma_{\text{T}}^2}. \quad (4.7)$$

Because Γ_{e} and Γ_{o} can be made large compared to γ_{m} in this platform, high transduction efficiencies are possible. For matched electromechanical and optomechanical damping rates $\Gamma_{\text{e}} = \Gamma_{\text{o}} \gg \gamma_{\text{m}}$, η_{t} is limited only by the loss rates of the microwave and optical resonators and optical mode matching, and achieves its maximum value $\eta_{\text{M}} = \epsilon \frac{\kappa_{\text{o,ext}}}{\kappa_{\text{o}}} \frac{\kappa_{\text{e,ext}}}{\kappa_{\text{e}}}$, where I have introduced the bidirectional modematching $\epsilon = \sqrt{\epsilon_{\text{PC}}\epsilon_{\text{CL}}}$.

4.1.3 Optomechanical sideband cooling

Though this platform can be made to be highly efficient, the MHz-scale resonant frequencies of membranes means that the mechanical transduction mode will still be highly occupied when precooled to dilution-refrigerator temperatures. For a membrane-based transducer to faithfully

transduce a quantum state, transduction must occur at a rate faster than environmental phonons enter the mechanical intermediary mode. Because the same beamsplitter processes that mediate transduction cool the membrane mode in the absence of a signal to transduce, this requirement is the same as that for cooling the mechanical mode to its ground state. For definiteness, I consider the optomechanical interaction in isolation. Optomechanical cooling arises from the imbalance between the rates of Stokes and anti-Stokes scattering processes when the pump is red-detuned from the cavity mode. The Stokes process, wherein a pump photon is scattered to lower energy by emitting a phonon into the membrane mode, is suppressed relative to the complementary anti-Stokes process by the reduced density of states away from cavity resonance. The rates of the anti-Stokes and Stokes processes are given respectively by $n_m \Gamma_{o,+}$ and $(n_m + 1) \Gamma_{o,-}$, where n_m is the phonon occupancy of the membrane mode, and $\Gamma_{o,+}$ and $\Gamma_{o,-}$ are the anti-Stokes and Stokes contributions to the net optomechanical damping rate Γ_o (i.e., the positive and negative terms in Eq. (4.4)) [Aspelmeyer, Kippenberg, and Marquardt 2014].

If the mechanical mode were completely decoupled from its thermal environment, it would reach equilibrium when the Stokes and anti-Stokes scattering rates become equal, resulting in a final occupancy

$$n_{\min,o} = \frac{\bar{r}_o}{1 - \bar{r}_o} = \frac{(\kappa_o/2)^2 + (\Delta_o + \omega_m)^2}{-4\Delta_o\omega_m}, \quad (4.8)$$

where $\bar{r}_o = \Gamma_{o,-}/\Gamma_{o,+}$. To generalize to the case where the mechanical mode is coupled to more than one bath, note that $\Gamma_{o,-} = \Gamma_o n_{\min,o}$: thus heating by Stokes scattering can be described as a coupling to a bath $n_{\min,o}$ at rate Γ_o , to which the membrane mode equilibrates when Γ_o overwhelms all other contributions to the total mechanical mode damping Γ_T . This lower bound on the mode occupancy arising from residual Stokes scattering is called the quantum backaction limit. The derivation of the electromechanical backaction limit is entirely analogous to the optomechanical derivation. The transducer gain factors introduced in Sec. 4.1.2 are related to the backaction bath occupancies by $\mathcal{A}_o = 1 + n_{\min,o}$ and $\mathcal{A}_e = 1 + n_{\min,e}$.

Technical noise leading to increased effective occupancies $n_{\text{eff},o}$ and $n_{\text{eff},e}$ of the optical and

microwave resonator modes will modify these baths to $n_o = n_{\min,o} + n_{\text{eff},o}$ and $n_e = n_{\min,e} + n_{\text{eff},e}$, respectively. The competition between the membrane mode's coupling to its thermal environment with occupancy n_{th} through its intrinsic dissipation rate γ_m and its coupling to optomechanical and electromechanical baths results in a weighted average for the final occupancy of the membrane mode:

$$n_m = \frac{\gamma_m n_{\text{th}} + \Gamma_e n_e + \Gamma_o n_o}{\Gamma_T}. \quad (4.9)$$

With sufficiently low technical noise and high enough sideband resolution, it is possible to cool the mechanical mode to an occupation of less than one phonon, and close to its ground state.

4.1.4 Optical output spectrum

Any residual thermal occupation of the mechanical mode will contribute to N_{add} . To infer the membrane mode occupancy and study the transducer's noise performance, we make heterodyne measurements of the noise density at the optical cavity output port. Here I describe the relevant features of the optical output spectrum in the absence of technical noise, and in Sec. 4.1.5 I generalize the expressions to include the effects of technical noise. The scattering processes described above generate Lorentzian peaks in the optical output spectral density at detunings $\pm\omega_m$ from the pump frequency. In units of photons/s/Hz referred to the transducer's optical output port, the spectrum around the upper (anti-Stokes) sideband has the form

$$S_{o,+}(\omega) = 1 + \epsilon_{\text{CL}} \frac{\kappa_{o,\text{ext}}}{\kappa_o} \frac{\Gamma_o \Gamma_T (1 + n_{\min,o}) n_m}{\Gamma_T^2/4 + (\omega - \omega_m)^2}, \quad (4.10)$$

and the spectrum around the lower (Stokes) sideband is

$$S_{o,-}(\omega) = 1 + \epsilon_{\text{CL}} \frac{\kappa_{o,\text{ext}}}{\kappa_o} \frac{\Gamma_o \Gamma_T n_{\min,o} (n_m + 1)}{\Gamma_T^2/4 + (\omega + \omega_m)^2}, \quad (4.11)$$

where the background terms correspond to the sum of vacuum noise and the added noise of an ideal heterodyne detector. The upper and lower sideband amplitudes $N_{\text{out},o,\pm} = S_{o,\pm}(\pm\omega_m) - 1$ are proportional to the optomechanical anti-Stokes and Stokes scattering rates, respectively, and thus

the ratio of measured sideband amplitudes $r_o = N_{\text{out},o,-}/N_{\text{out},o,+}$ can be used to infer the mode occupancy by inverting

$$r_o = \frac{n_{\text{min},o}}{n_{\text{min},o} + 1} \frac{n_m + 1}{n_m}. \quad (4.12)$$

At high temperatures, $n_m \approx n_m + 1$, and comparing Eq. (4.12) to Eq. (4.8) reveals that $r_o \rightarrow \bar{r}_o$. As n_m decreases, r_o increases, asymptoting toward 1 as n_m approaches its minimum value $n_{\text{min},o}$.

When both pumps are present, the beamsplitter interactions transduce incident microwave signals detuned by $+\omega_m$ from the microwave pump to outgoing signals at detuning $+\omega_m$ from the optical pump, and vice versa. Thus the transducer noise that competes with an upconverted signal at the optical output port is $N_{\text{out},o,+}$. The input-referred upconversion added noise due to membrane motion is then given by

$$N_{\text{add},\text{up}} = \frac{N_{\text{out},o,+}}{\mathcal{A}_e \mathcal{A}_o \eta_{t,\text{up}}} = \frac{n_m}{\mathcal{A}_e \frac{\kappa_{e,\text{ext}}}{\kappa_e}} \frac{\Gamma_T}{\Gamma_e}, \quad (4.13)$$

This equation indicates that, up to small effects from imperfect sideband resolution and microwave mode internal loss, the ability to ground-state cool the mechanical mode with purely electromechanical damping is a necessary and sufficient condition for $N_{\text{add},\text{up}} < 1$. Assuming comparable electromechanical and optomechanical sideband resolution, n_m will remain approximately constant at constant total damping Γ_T , but trading off electromechanical damping for optomechanical damping while holding Γ_T constant can only increase the added noise. Moreover, $N_{\text{add},\text{up}}$ is independent of all other parameters of the optomechanical system. Qualitatively, this is because after the signal to be upconverted enters the mechanical mode, it copropagates with any noise routed to the optical output port, so the effects of inefficiency due to the optical cavity parameters impact the signal and noise equally.

The form of the upper and lower electromechanical sideband amplitudes $N_{\text{out},e,+}$ and $N_{\text{out},e,-}$ in the microwave output spectrum is analogous to the optomechanical case described above. In the absence of technical noise, the input-referred downconversion added noise is given by

$$N_{\text{add},\text{down}} = \frac{N_{\text{out},e,+}}{\mathcal{A}_e \mathcal{A}_o \eta_{t,\text{down}}} = \frac{n_m}{\epsilon_{\text{PC}} \mathcal{A}_o \frac{\kappa_{o,\text{ext}}}{\kappa_o}} \frac{\Gamma_T}{\Gamma_o}, \quad (4.14)$$

independent of electromechanical system parameters. Without technical noise, the ability to ground-state with the purely optomechanical damping is a necessary and sufficient condition for $N_{\text{add,down}} < 1$.

4.1.5 Effects of technical noise

I define technical noise broadly, to encompass all pump power-dependent effects on the transducer noise performance. Noise sources that fall into this category include amplitude and phase fluctuations of the incident microwave or optical pumps at detuning ω_m , fluctuations of resonator parameters at frequency ω_m (e.g., fluctuations of the superconducting circuit resonant frequency due to two-level systems [Gao et al. 2007]), and heating of the superconducting circuit by either of the pumps. All of these noise sources present either as an effective microwave mode thermal occupancy $n_{\text{eff,e}}$ coupled to the membrane mode at rate Γ_e or as an effective optical mode occupancy $n_{\text{eff,o}}$ coupled to the membrane mode at rate Γ_o [Aspelmeyer, Kippenberg, and Marquardt 2014; Jayich et al. 2012]. Therefore, their contributions to n_m and N_{add} are already included in Eqs. 4.9, 4.13, and 4.14.

Noise sources that are excluded from this definition include heating of the mechanical mode due to its coupling to its thermal environment with occupation n_{th} , the electromechanical and optomechanical backaction contributions described in Sec. 4.1.3, and an additional term due to the backaction of an auxiliary locking beam used to stabilize the frequency of the laser to that of the optical cavity. Though this last contribution is arguably quite technical in nature, it doesn't scale with either pump power, so does not enter into the formalism in the same way. It enters as an additional backaction term $\gamma_{\text{lock}} n_{\text{min,lock}}$ in the numerator of Eq. 4.9. Though the locking beam's power is small compared to that of the pump beam, and its damping rate γ_{lock} is negligible compared to Γ_o , because it is detuned close to the optical cavity's resonance where $n_{\text{min,lock}}$ diverges, this additional source of backaction is not negligible. This contribution to noise is considered in Sec. 4.2.3, and characterizing it is discussed in Section 5.6.4.

When readout of mechanical motion is performed using a pump coupled to a bath with

nonvanishing effective occupancy n_{eff} , two other effects of technical noise on the output spectrum must be considered. First, technical noise on the output side will generate an additional contribution to N_{add} distinct from its direct effect on n_m . Second, technical noise from pump or parameter fluctuations can interferometrically modify the amplitudes of the Stokes and anti-Stokes peaks in the output spectrum, and to correctly infer n_m we must correct for this effect.

In the remainder of this section, I restrict my focus to technical noise arising from amplitude and phase fluctuations of the incident optical pump as it is the noise relevant for upconversion, following the formalism of Ref. [Jayich et al. 2012]. The optical mode effective thermal occupancy due to amplitude and phase fluctuations is given by

$$n_{\text{eff,o}} = \frac{1}{4} \epsilon_{\text{PC}} \frac{\kappa_{\text{o,ext}}}{\kappa_{\text{o}}} \mathcal{A}_{\text{o}} \frac{\kappa_{\text{o}}^2}{4} \left(|B_{\text{x}}(\omega_{\text{m}})|^2 C_{\text{xx}} + |B_{\text{y}}(\omega_{\text{m}})|^2 C_{\text{yy}} + 2\text{Im} [B_{\text{x}}(\omega_{\text{m}}) B_{\text{y}}^*(\omega_{\text{m}})] C_{\text{xy}} \right), \quad (4.15)$$

where $B_{\text{x}}(\omega) = e^{-i\phi} \chi_{\text{o}}(\omega) + e^{i\phi} \chi_{\text{o}}^*(-\omega)$ and $B_{\text{y}}(\omega) = e^{-i\phi} \chi_{\text{o}}(\omega) - e^{i\phi} \chi_{\text{o}}^*(-\omega)$ are amplitude and phase noise susceptibilities, $\phi = \tan^{-1}(2\Delta_{\text{o}}/\kappa_{\text{o}})$ is the optomechanical phase shift, $\chi_{\text{o}}(\omega) = 1/(\kappa_{\text{o}}/2 - i(\omega + \Delta_{\text{o}}))$ is the optical cavity susceptibility, and C_{xx} , C_{yy} , and C_{xy} are technical noise spectral densities. Specifically, C_{xx} and C_{yy} are two-sided amplitude and phase noise spectral densities, respectively, while C_{xy} is the two-sided amplitude/phase cross-correlation spectral density, subject to the constraint $C_{\text{xy}}^2 \leq C_{\text{xx}} C_{\text{yy}}$. All three spectral densities are evaluated at detuning ω_{m} from the pump, and expressed in units of photons/s/Hz normalized to the shot noise of the beam incident on the optical cavity. They are thus all proportional to pump power, motivating the parameterization $n_{\text{eff,o}} = a_{\text{o}} \Gamma_{\text{o}}$ when fitting n_m vs. Γ_{o} in Sec. 5.2 (see Sec. 5.7.3 for further discussion).

C_{xx} , C_{yy} , and C_{xy} can all be modeled as white noise spectral densities in the narrow spectral window of interest around $\pm\omega_{\text{m}}$. They thus generate frequency-independent excess noise $\tilde{S}_{\text{o},+}$ ($\tilde{S}_{\text{o},-}$) over the background in the optical output spectrum around the upper (lower) sideband. These excess noise spectral densities are given by

$$\begin{aligned} \tilde{S}_{\text{o},\pm} = & \frac{1}{4} \epsilon_{\text{CL}} \epsilon_{\text{PC}} \left(\left[|\rho|^2 + |\kappa_{\text{o,ext}} \chi_{\text{o}}(\pm\omega_{\text{m}}) - 1|^2 \right] (C_{\text{xx}} + C_{\text{yy}}) \right. \\ & \left. - 2\text{Re}[\rho^* (\kappa_{\text{o,ext}} \chi_{\text{o}}(\pm\omega_{\text{m}}) - 1) (C_{\text{xx}} + 2iC_{\text{xy}} - C_{\text{yy}})] \right) \end{aligned} \quad (4.16)$$




Figure 4.2: **Technical pump noise** The amplitude and phase fluctuations on the pump $\tilde{S}_{o,+}$ (red) impart correlated motion onto the mechanical mode. The optomechanical signal (purple) then interferes with the reflected pump noise with which it is correlated.

where $\rho = 1 - \kappa_{o,\text{ext}} / (\kappa_o/2 - i\Delta_o)$ is the attenuation of the reflected pump due to absorption in the cavity. Amplitude noise C_{xx} would contribute to $\tilde{S}_{o,\pm}$ even in the absence of a cavity, while C_{yy} and C_{xy} contribute because the cavity rotates phase fluctuations into the amplitude quadrature. For an overcoupled cavity at optimal detuning in the resolved sideband limit, phase noise contributes primarily to $\tilde{S}_{o,+}$ and amplitude noise contributes primarily to $\tilde{S}_{o,-}$. In the discussion of transducer added noise in Sec. 5.3, for which only noise in the upper sideband is relevant, we introduce the symbol $\tilde{N}_{\text{out},o} = \tilde{S}_{o,+}(+\omega_m) = \tilde{S}_{o,+}$, to parallel the notation used for the Lorentzian contribution to the noise (see Fig. 4.2).

The final effect we must account for is a modification of the Stokes and anti-Stokes peaks in the optical output spectrum due to correlations between membrane motion and amplitude and phase fluctuations of the incident light. Eqs. (4.10) and (4.11) become

$$S_{o,+}(\omega) = 1 + \tilde{S}_{o,+} + \epsilon_{\text{CL}} \frac{\kappa_{o,\text{ext}}}{\kappa_o} \mathcal{A}_o \frac{\Gamma_o}{\Gamma_T^2/4 + (\omega - \omega_m)^2} \times \left[\Gamma_T \left(n_m - \epsilon_{\text{PC}} \frac{\kappa_o^2}{4} \text{Re}[\tilde{B}_+] \right) - 2\epsilon_{\text{PC}} (\omega - \omega_m) \frac{\kappa_o^2}{4} \text{Im}[\tilde{B}_+] \right] \quad (4.17)$$

and

$$S_{o,-}(\omega) = 1 + \tilde{S}_{o,-} + \epsilon_{\text{CL}} \frac{\kappa_{o,\text{ext}}}{\kappa_o} \mathcal{A}_o \frac{\Gamma_o}{\Gamma_T^2/4 + (\omega + \omega_m)^2} \times \left[\Gamma_T \left(\frac{n_{\text{min},o}}{\mathcal{A}_o} (n_m + 1) + \epsilon_{\text{PC}} \frac{\kappa_o^2}{4} \text{Re}[\tilde{B}_-] \right) - 2\epsilon_{\text{PC}} (\omega + \omega_m) \frac{\kappa_o^2}{4} \text{Im}[\tilde{B}_-] \right], \quad (4.18)$$

where

$$\begin{aligned} \tilde{B}_{\pm} = \frac{e^{-i\phi}}{4} & \left(\kappa_{\text{o,ext}} |\chi_{\text{o}}(\pm\omega_{\text{m}})|^2 [B_{\text{x}}(\mp\omega_{\text{m}})(C_{\text{xx}} + iC_{\text{xy}}) + B_{\text{y}}(\mp\omega_{\text{m}})(iC_{\text{xy}} - C_{\text{yy}})] \right. \\ & - \chi_{\text{o}}^*(\pm\omega_{\text{m}}) [(B_{\text{x}}(\mp\omega_{\text{m}})C_{\text{xx}} + iB_{\text{y}}(\mp\omega_{\text{m}})C_{\text{xy}})(1 + \rho) \\ & \left. + (iB_{\text{x}}(\mp\omega_{\text{m}})C_{\text{xy}} - B_{\text{y}}(\mp\omega_{\text{m}})C_{\text{yy}})(1 - \rho)] \right). \quad (4.19) \end{aligned}$$

Eqs. (4.17) and (4.18) reveal that these correlations both change the peak amplitudes $N_{\text{out,o},\pm}$ and add an anti-symmetric contribution to the lineshape of each peak, where the sign of these extra terms depends on the relative values of C_{xx} , C_{yy} , and C_{xy} . Eq. (4.13) for the transducer's upconversion added noise is then modified to account for amplitude and phase noise by the substitution $N_{\text{out,o},+} \rightarrow N_{\text{out,o},+} + \tilde{S}_{\text{o},+} = N_{\text{out,o},+} + \tilde{N}_{\text{out,o}}$.

4.2 Transducer design

From the first demonstration of a membrane-based electro-optic transducer [Andrews et al. 2014], there have been significant refinements that have brought this platform to the cusp of quantum performance. This section will delineate the more impactful refinements that I have contributed, as well as the motivation for these changes, with an eye toward how these might change with both other device improvements and developing experimental goals.

4.2.1 Transducer architecture

I'll first lay out the general architecture for the transducer devices we construct and measure, and establish a vocabulary for their various parts and substructures. The beating heart of the transducer is the silicon nitride (SiN) membrane whose (2,2) mechanical mode is used to mediate between the electromechanical LC circuit and the optomechanical Fabry-Pérot cavity, as depicted in Fig. 4.3(c). One corner of the membrane is metalized with superconducting niobium-titanium-nitride (NbTiN) to form one pad of a mechanically compliant capacitor. The silicon chip which holds the membrane (the top chip) is mounted to a second silicon chip (the bottom chip) in a flip chip assembly. The bottom chip holds the remainder of the electromechanical circuit—the split-pad




Figure 4.3: **Transducer architecture.** Detailed illustration depicting spatial location of the membrane (purple) relative to the superconducting LC resonator (blue), and optical resonator mode (red). Arrows indicate microwave (blue) and optical (red) pumps, and signals in upconversion (green) and downconversion incident on the microwave and optical ports. Inset shows a photograph of the phononically patterned silicon top chip that supports the membrane, false colored purple in the center. The illustration is zoomed-in to focus on the membrane, and shows the metalized corner that defines the mechanically compliant capacitor pad of the LC circuit (blue). The top chip (light gray) is exploded upward in the illustration to show the location of the split-pad capacitor plates and inductive loop comprising the remainder of the LC circuit on the bottom chip (dark gray). The optical cavity mode forms between the curved mirror (blue cylinder) and the planar mirror chip (blue) bonded beneath the bottom chip. The inset spatial mode profile illustrates the capacitor-pad-perturbed (2,2) mechanical mode of the membrane.




Figure 4.4: **Stable cavity geometries.** (a) Parallel planar mirrors. (b) Concentric spherical mirrors ($L = 2 \cdot \text{ROC}$). (c) Spherical mirrors with $L < 2 \cdot \text{ROC}$.

of the capacitor and inductive loop—as well as an inductive loop for external coupling, and a planar LC “reference resonator” for diagnostics. The opposite corner of the SiN membrane participates in the fundamental mode of the optical cavity, defined by a pair of mirrors with high-reflectivity dielectric coatings, whose geometry and design evolution will be detailed further in the following subsections.

4.2.2 Membrane misalignment

The performance of the first transducer was limited by misalignment of the membrane relative to the optical cavity mode [Andrews et al. 2014]. This misalignment perturbed the cavity mode, making the optical cavity the predominant source of inefficiency, by introducing loss by coupling the fundamental mode of the optical cavity to lossier higher-order modes, and increasing the difficulty of matching the fundamental Gaussian mode of our propagating laser to the perturbed cavity mode. Misalignment is also likely to result in greater amounts of optical power being scattered toward the superconductor, though this is difficult to model because of the strong spatial dependences involved. I will detail two changes that have mitigated the problem of misalignment: compactifying and rigidifying the manner that the electromechanical chip is mounted within the optical cavity, dramatically reducing misalignment, and designing a cavity that is less sensitive to inevitable misalignment in the first place.




Figure 4.5: **Optomechanical cavity geometry.** Curved (left) and flat (right) mirror surfaces (black lines) define the optical cavity mode that forms along the axis (dashed black line) that is perpendicular to the flat mirror and intersects the center of curvature of the right mirror. A cavity length of $L = 2 \text{ mm}$ is chosen for illustrative purposes, and the resulting cavity mode waist has $1/e^2$ radius of 37 um that expands slightly to a radius of 41 um at the curved mirror surface. The membrane (gray line) is a distance $t = 700 \text{ um}$ from the flat mirror, as in earlier devices, and misaligned from the cavity mode by θ . All geometries are to scale.

4.2.2.1 Stable cavity geometry

An optical cavity can only have high finesse if it is geometrically stable, i.e. if the mirror geometry directs a reflected beam toward the axis that defines the cavity mode, such that the distance from this axis does not grow after many reflections [Siegman 1986]. There are various stable geometries, but two illustrative extremes on the boundary of stability are a cavity made of parallel, planar end mirrors, and one made of concentric spherical mirrors (see Fig. 4.4). Any slight deviation from parallel renders the former unstable, and similarly an infinitesimal increase in separation, L , between the spherical mirrors causes the latter to be unstable. However, bringing spherical mirrors closer together than concentric, such that $L < 2 \cdot \text{ROC}$, where ROC is the mirrors' radius of curvature, results in geometries whose stabilities are robust to small perturbations. Equivalently, a cavity formed by a curved mirror and a flat mirror is stable if their separation is less than the radius of curvature of the curved mirror, $L < \text{ROC}$. The optical cavities used in our transducer are made up of a planar mirror and a spherical mirror with $\text{ROC} = 1 \text{ cm}$,¹ separated by a distance $L = 1 - 2 \text{ mm}$. The resulting cavity mode forms between the mirrors, along the line that intersects the center of curvature of the spherical mirror and is normal to the planar mirror, with waist located at the flat mirror with $1/e^2$ radius of 30-40 μm . See Fig. 4.5 for a scale diagram of this cavity geometry.

This choice of geometry lends itself to a helpful framework for thinking about the cavity mode in which there is a “division of labor” between the mirrors—the flat mirror determines the direction of propagation of the mode, and the curved mirror position determines its location. Additionally, the flat mirror can be placed closer to the flat electromechanical chip that houses the membrane, without interfering with it. Because of the flat mirror's proximity to the membrane, it is the optical field that on that side of the membrane that is responsible for a greater share of the optomechanical coupling.

For this reason, it's clear that referencing the membrane to the flat mirror would ensure that

¹ This is the shortest radius of curvature that commercial suppliers are willing to produce with the low RMS surface roughness that we require of $\sim 1 \text{ \AA}$ using traditional superpolishing techniques, though there is currently exciting progress toward mm-scale mirror curvatures using lithographic techniques [Jin et al. 2022].

the membrane is always perpendicular to the cavity mode, an important consideration for stable optomechanical cavities. I attempted many ways of mounting the silicon electromechanical flip-chip to the fused silica substrate of the flat mirror using cryogenic epoxy in an attempt to make an “etalon” with parallel surfaces, and developed a set of tools to measure the resulting misalignment, θ in Fig. 4.5, using a temperature-tunable distributed-feedback laser (QPhotonics QDFBLD-1064) at room temperature and at 4 K using an optical-access Montana Instruments Fusion F2 cryostat. Visual inspection often gave sufficient information, as any visible light interference fringes over the 500 μm width of the membrane could bound the misalignment to the desired precision of approximately 1 milliradian. By far the two most successful methods to make well-aligned etalons were accomplished in the cleanroom by Pete Burns [Burns 2019]: optical contacting, in which two surfaces are pressed together, and are sufficiently clean and flat that intermolecular electrostatic forces cause adhesion, and a chemical hydrophilic Si-SiO₂ bonding process which took advantage of silica (SiO₂) being one of the dielectric materials composing the dielectric mirror stack. Our mirror coating supplier Ramin Lalezari at FiveNine Optics could add 100 Å of SiO₂ as a capping layer, with minimal impact on mirror performance. The greater strength of the chemical bonding process resulted in a bond that was more robust to differences in the coefficients of thermal contraction of the silicon chip and the mirror coating.

4.2.2.2 Choice of input mirror

To maximize transduction efficiency it is necessary to use asymmetric mirror transmissions and measure in reflection on the side of the cavity with greater transmission, which I’ll refer to as the front mirror. One option is to choose the flat mirror to be the front mirror. This way, the transducer could be operated in a mode in which the optical field in the cavity preferentially builds up on the portion of the cavity between the flat mirror and the membrane, and the optomechanical coupling is maximized simultaneously with the external coupling, in theory offering the greatest transduction efficiency with the lowest optical power circulating in the cavity. However, when operating transducer cavities designed in this way, we observed that we could not achieve the



Figure 4.6: **Input mirror choice** Choosing the flat mirror to be in front (left configuration) is susceptible to misalignment of the membrane as the optical field builds up in the unstable portion of the cavity when most overcoupled. Instead, choosing the curved mirror to be the input mirror (right configuration) optimizes stability when external coupling is greatest.

optomechanical couplings we predicted, and that the overall cavity loss we measured was consistent with the optical field preferentially building up in the region between the curved mirror and the membrane. Apparently, in the presence of inevitable misalignment, the portion of the cavity between the flat mirror and the membrane, itself an unstable geometry, contributed enough loss to make the cavity difficult to operate in the desired configuration, with the light building up on the shorter side of the cavity.

Instead, choosing the curved mirror to be the front mirror provides two benefits. The first is that it allows the efficiency to be maximized when the cavity will perform most ideally in the presence of misalignment, as depicted in Fig. 4.6. This comes at the cost at a moderately reduced optomechanical coupling from the light field building up in the longer section of the cavity. For our geometries this tends to be a reduction of about a factor of two in g_o , and depends on the length ratio between the two sections of the cavity. The second is that it allows experimental flexibility when operating an *in-situ* device. With this choice of front mirror, the cavity is made up of a longer section with a higher throughput mirror, and a shorter section with a more reflective mirror, so the overall loss rate of the cavity can tend to remain the same when changing between the higher-coupling, and higher-stability configurations. This allows the transducer to be operated in either configuration by changing the wavelength of the laser by a fraction of a nanometer while maintaining consistently high sideband resolution, and therefore lower added noise. We took advantage of this freedom while measuring the transducer that demonstrated ground-state cooling and qubit readout [Brubaker et

al. 2022; Delaney et al. 2022]—even though the cavity was more robust because of its bonded etalon, we achieved optimal performance when we operated it in the more conservative, and lower-coupling configuration.

4.2.2.3 1-D model for the transducer cavity and cryogenic cavity measurements

To better understand the changes to the cavity design that have made it more insensitive to misalignment, it is helpful to understand the effect of a perfectly-aligned membrane in an optical cavity, which is encapsulated by a one-dimensional model of the cavity. The insertion of a dielectric membrane into an optical cavity will apply a position-dependent perturbation to its resonance frequency [Thompson et al. 2008]. Additionally, the total cavity energy loss rate κ_o , individual mirror loss rates of the front mirror $\kappa_{o,\text{ext}}$ and back mirror κ_{back} , and optomechanical coupling rate g_o depend on the membrane position, resulting from the membrane reflection causing constructive interference on one side of the membrane relative to the other. By using a transfer matrix formalism, it is relatively straightforward to model a cavity with a dielectric membrane [Wipfli 2015], and extract the optomechanical coupling rate, along with the cavity loss rate due to the total transmission of the mirrors by fitting the Lorentzian response. And by slightly modifying the technique laid out in [Wipfli 2015], the individual mirror coupling rates can be solved for. This model does not include additional optical loss as it depends on the cavity misalignment, the membrane-perturbed frequencies of higher-order modes, and the position of lossy scatterers.

Fig. 4.7(b) shows the results of this transfer matrix model to extract the cavity frequency shift per membrane displacement G_o , the external coupling rate $\kappa_{o,\text{ext}}$, and the total cavity linewidth in the absence of loss $\kappa_{o,\text{ext}} + \kappa_{\text{back}}$, using the mirror reflectivities, cavity length L , and membrane position t of an earlier device (that measured in [Higginbotham et al. 2018]), in which we used the flat mirror as the externally-coupled front mirror, as depicted schematically in Fig. 4.7(a). In this device, both mirrors were mounted on piezoelectric actuators that allowed *in-situ* control of the cavity resonance frequency and membrane position relative to the optical standing wave of the cavity. These quantities are displayed as a function of the flat mirror position, while the

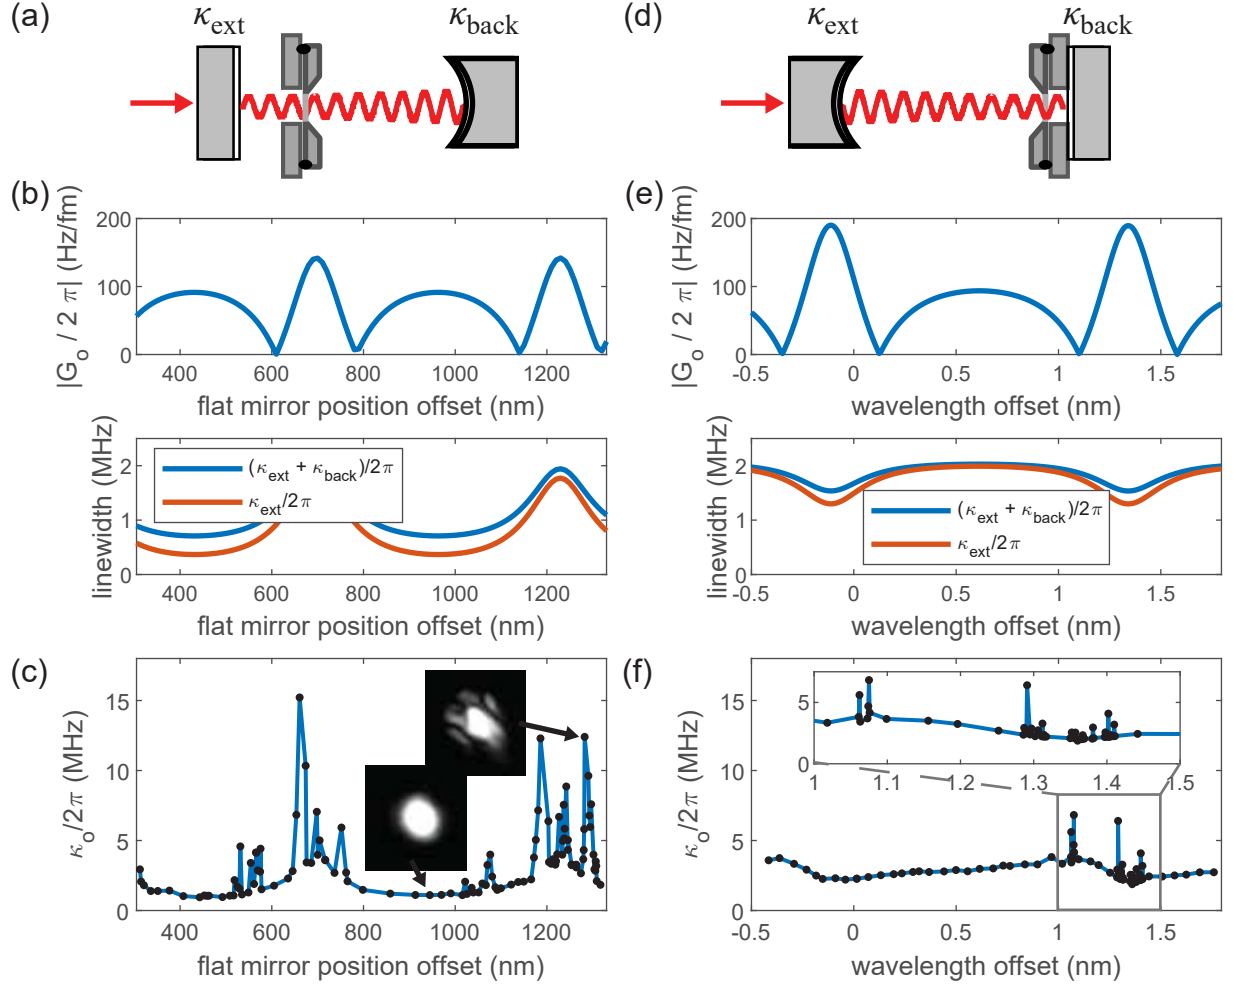


Figure 4.7: Optical cavity architecture improvements. (a) Schematic representation of a transducer cavity using the flat mirror as the externally-coupled port, with the membrane mounted independently of the flat mirror, as in Ref. [Higginbotham et al. 2018]. (b) Results from transfer-matrix simulation calculating G_o , $\kappa_{o,\text{ext}}$, and $\kappa_{o,\text{ext}} + \kappa_{\text{back}}$ vs. flat mirror position, for the device represented in (a). In this architecture, greater overcoupling ratio is chosen to coincide with maximum G_o and the intended operating κ_o . (c) Measured κ_o vs. flat mirror position. The loss rate is continuous over some offset intervals, and peaks sharply in flat mirror positions associated with hybridization between the fundamental Gaussian mode and higher-order optical modes. Inset photographs of the cavity mode imaged in transmission are associated with the indicated operating points. The high-loss peaks are concentrated in high- G_o configurations, and prohibit transducer operation. (d) Schematic representation of a transducer cavity architecture using the curved mirror as the externally-coupled port and a bonded membrane-mirror etalon, as simulated in (e) and measured in (f). (e) Transfer-matrix simulation of cavity parameters vs. laser wavelength. In this architecture, in contrast, the cavity operates with the intended κ_o for all wavelengths, allowing *in-situ* experimental choice between either local maximum of G_o . (f) Measurement of κ_o vs. wavelength. This device demonstrates the enhanced stability of this cavity geometry and bonded etalon, as shown by the fewer, narrower, and lower peaks associated with mode hybridization in comparison with (c). The inset expands the x-axis in the indicated wavelength region of high- G_o to show that the cavity could be operated at the global coupling maximum, near an offset of 1.32 nm. Technical detail: in the measurements shown in (f), I did not repeat the more involved process of searching out and measuring the hybridized modes for wavelength offsets between -0.5 nm and 0 nm, corresponding to similar operating parameters as the region detailed in the inset, one free-spectral-range of the etalon away.

curved mirror's position moves continuously in order to maintain resonance of the cavity with a fixed-frequency laser. In the course of moving the mirror by $\lambda/2$, half the laser wavelength, $|G_o|$ experiences two nulls, when the membrane sits at the node and antinode of the optical field intensity, as well as two local maxima when the gradient of the intensity field is maximized. The greater local maximum corresponds to the light preferentially building up on the shorter cavity subsection between the membrane and flat mirror, and the lower maximum corresponds to light building up on the curved-mirror side of the membrane. Because the flat mirror is chosen to be the front mirror, $\kappa_{o,\text{ext}}$ is maximized at the global maximum of G_o . There is an asymmetry in the widths of the features stemming from the choice of which mirror position is chosen for the x-axis, because of how the membrane perturbs the standing wave as that mirror shifts the membrane from an antinode to a node of the optical field. The result is that the narrower feature corresponds to when the light builds up on the side of the mirror whose position is on the x-axis, or the mirror whose position is actuated experimentally for this sweep, and the wider feature corresponds to light building up on the side of the cavity whose mirror position is constrained by maintaining resonance with the laser.

Experimentally measured values of κ_o of the optical cavity operating at the base temperature of a dilution refrigerator are shown in Fig. 4.7(c), with a position offset chosen to align the experimentally obtained data with the simulation, as it is difficult to measure the cavity and etalon lengths to sufficient precision for absolute congruence of the x-axis. Though there are continuous regions that qualitatively match the simulated position dependence of $\kappa_{o,\text{ext}} + \kappa_{\text{back}}$, there are also regions where the experimentally measured κ_o is much higher, corresponding with observation of visible hybridization with higher order modes (see inset photographs). The transducer had degraded performance at these lossier operating points, which cluster about the intended operating points of high optomechanical coupling and have greater optical fields building up in the region of the cavity between the flat mirror and the membrane. The instability of this portion of the cavity leads to greater coupling to higher-order modes of the un-membrane-perturbed cavity, because of the delocalizing effect of misaligned planar surfaces.

Fig. 4.7(e) and (f) respectively show simulation and measurement at base temperature for a later cavity, with a chemically bonded etalon, and with the curved mirror as the strongly-coupled port, as depicted in Fig. 4.7(d). The cavity is shorter overall, and the bonded etalon allows for a closer spacing between the flat mirror and the membrane chip, so G_o is generally higher. The bonded etalon also precludes changing the spacing between the etalon and the flat mirror, fixing it to a distance of $t \approx 390$ μm (the thickness of the bottom chip), so to vary G_o *in situ*, the laser frequency was changed. The curved mirror was actuated with a piezo, and its position was again changed to maintain cavity resonance with the laser, in both simulation and experiment. Again we observe behavior that is periodic in the free spectral range of the $t = 390$ μm etalon, $\Delta\lambda_{\text{fsr}} = \lambda^2/2t = 1.5$ nm . The device is designed so $\kappa_{o,\text{ext}}$ is maximized at the lower local maximum of G_o , and the variation in $\kappa_{o,\text{ext}} + \kappa_{\text{back}}$ is less. The measured values of κ_o exhibit fewer and narrower peaks associated with hybridization of the fundamental optical mode with higher-order modes, indicating the marked improvement in alignment with the bonded etalon. These peaks are again concentrated at wavelengths associated with more light in the side of the cavity with the flat mirror, however in this device, the features are sufficiently narrow that the cavity could be operated between these features close to the global maximum of G_o , as can be seen from the inset showing greater detail in this region. This resulted in a cavity that could be operated at dramatically higher optomechanical coupling, with $g_o \sim 100$ Hz increased by an order of magnitude from 7 Hz [Higginbotham et al. 2018], and with higher efficiency, with $\kappa_{o,\text{ext}}/\kappa_o$ improved to 70% from 50%.

4.2.3 Selecting values for mirror transmissions

Having developed the machinery in the previous section to simulate $\kappa_{o,\text{ext}} + \kappa_{\text{back}}$, ideally the primary contributions to κ_o , I'll now briefly relate the considerations for choosing transmission values toward a given experimental goal. The primary tension in this choice is between achieving the highest level of sideband resolution by making $\kappa_{o,\text{ext}} + \kappa_{\text{back}}$ as small as possible to ensure $\kappa_o \ll 4\omega_m$, and ensuring that efficiency remains high by maximizing $\kappa_{o,\text{ext}}/\kappa_o$, recognizing that κ_o

will inevitably have nonideal contributions due to scattering and absorption from the mirror and membrane surfaces.

The sideband resolution determines the minimum mechanical occupation achievable by laser cooling with a noiseless optical pump, as given by Eq. 4.8. This simplifies to $(\kappa_o/4\omega_m)^2$ for optimal pump detuning $\Delta_o = -\omega_m$. In any transducer demonstration, we would like this number to be sufficiently below one so it is possible ground-state cool the mechanical mode.

For upconversion, the contribution to the efficiency from the optical cavity is important in ensuring that the signal is dominant over any technical noise associated with the optical chain, though of course any noise introduced earlier in the transduction process, such as residual thermal noise of the mechanical mode or microwave technical noise, is attenuated as the signal is, so this efficiency has little effect on N_{add} if these are its dominant contributions. Conversely, for downconversion, this contribution to efficiency is paramount. In practice, we have observed scattering and absorption contributions to the optical power loss on the scale of 10 ppm, and so this sets the scale for mirror transmissions in order to dominate over this number.

In recent transducers, we have operated with mirror coatings designed to give values of $\kappa_o/2\pi$ of approximately 2 MHz, to obtain $n_{\text{min,o}} \approx 0.1$ for a 1.5 MHz mechanical mode frequency. This has meant using coatings with power transmissions of ~ 200 ppm for ~ 2 mm long cavities as in Ref. [Brubaker et al. 2022], and more recently ~ 100 ppm for ~ 1 mm long cavities. Assuming additional scattering/absorption losses of 20 ppm and using back mirror transmissions of 10 ppm to allow for a sufficient amount of light in transmission to be useful for diagnostics, this anticipates values of $\kappa_{o,\text{ext}}/\kappa_o$ of 0.87 and 0.77, respectively. The device in [Brubaker et al. 2022] had anomalously high additional optical loss, but only underperformed this slightly, with $\kappa_{o,\text{ext}}/\kappa_o = 0.79$, indicating that this is the correct scale to anticipate for additional contributions to optical internal loss.

With the changing experimental goals associated with our transducers approaching quantum performance, it could be advantageous to tailor mirror coating choices for a given experiment. For instance, in a network distributing entanglement via squeezed light from an optical source in down-




Figure 4.8: **Locking beam backaction.** The locking beam (orange) stabilizes the laser with respect to the optical cavity using the PDH technique. Although the locking beam can be operated with orders of magnitude lower power than the powers we require from the optical pump (red), because it is near resonance it contributes optomechanical backaction as given by Eq. 4.20.

conversion, the optical efficiency is relatively more important than in upconversion. Alternatively, in the context of a photon counting experiment in upconversion through a noisy electrical circuit or mechanical mode with residual thermal occupation, sideband resolution could be more important than efficiency through the optical cavity and so using lower throughput mirrors than our current design could be advantageous.

An additional consideration is the optomechanical backaction on the membrane mode due to the locking beam used to stabilize the frequency of the laser relative to that of the transducer cavity using the Pound-Drever-Hall (PDH) technique (see Fig. 4.8). The locking beam causes backaction

$$\begin{aligned} \gamma_{\text{lock}} n_{\text{min,lock}} &= -g_o^2 \bar{n}_{\text{lock}} \left(\frac{\kappa_o}{(\kappa_o/2)^2 + (\omega_m + \Delta_{\text{lock}})^2} - \frac{\kappa_o}{(\kappa_o/2)^2 + (\omega_m - \Delta_{\text{lock}})^2} \right) \\ &\quad \cdot \frac{(\kappa_o/2)^2 + (\omega_m + \Delta_{\text{lock}})^2}{4\Delta_{\text{lock}}\omega_m} \\ &\approx g_o^2 \bar{n}_{\text{lock}} \frac{\kappa_o}{(\kappa_o/2)^2 + \omega_m^2}, \end{aligned} \quad (4.20)$$

discussed in Secs. 4.1.5 and 5.6.4. This is decreased by decreasing κ_o , as can be seen on the second line, when Taylor expanded about $\Delta_{\text{lock}} = 0$ for a resonant lock beam. Here \bar{n}_{lock} is the number of lock photons circulating in the cavity, and is given by $(P_{\text{lock}}/\hbar\omega_{\text{lock}})\epsilon_{\text{LC}} (\kappa_{o,\text{ext}}/(\kappa^2/4 + \Delta_{\text{lock}}^2))$, where P_{lock} and ω_{lock} are the lock power and detuning, respectively, and ϵ_{LC} is the modematching factor between the lock beam and cavity. This effect contributed 0.4 photons/s/Hz to N_{add} in Ref. [Brubaker et al. 2022].

With this sizable contribution from lock backaction in mind, for the nearer-term goal of demonstrating a quantum-enabled transducer, it could also be advantageous to increase the optical cavity length, increasing sideband resolution by decreasing κ_o , at the expense of g_o rather than efficiency. Lengthening the cavity, in addition to increasing sideband resolution, would decrease g_o , which would have the further benefit of decreasing the lock backaction with the square of the this decrease in g_o . Assuming a sufficiently low-noise laser, this is without much drawback. However, with the more stringent requirements on laser noise of a photon-counting experiment, decreasing g_o becomes a steeper price to pay, as more laser power would be required for the same transduction bandwidth, leading to more noise counts from pump leakage and laser noise (see Sec. 6.2). In principle it is also possible to use a separate laser to lock the cavity at a longitudinal mode with lower optomechanical coupling, but stabilizing separate lasers detuned by hundreds of GHz or more would increase experimental complexity.

4.2.4 Fully fixed optical cavity

In operating the device measured in Fig. 4.7(e), we observed noise on the optical signal due to electrical pickup modulating the piezo used to actuate the curved mirror of the transducer. This noise could be attenuated by shorting the piezo leads at their input to the cryostat, but not removed completely, as required for a low-noise transducer. To eliminate this source of noise, we operated subsequent cavities without this degree of freedom, first by shorting the piezo leads at the transducer as in Ref. [Brubaker et al. 2022], and after by mounting the curved mirror without a piezo attached. Without the ability to tune the cavity length to maintain resonance with the laser at all wavelengths, it is not possible for the transducer to operate at the full continuum of operating points available to the devices in Fig. 4.7, but rather the device is limited to a discrete subset of wavelengths defined by the membrane-perturbed free-spectral-range of the cavity. Fig. 4.9 shows the simulated parameters of an optical cavity with the transmissions and dimensions used in Ref. [Brubaker et al. 2022], and the vertical dashed lines indicate the operating points available to a fixed cavity.

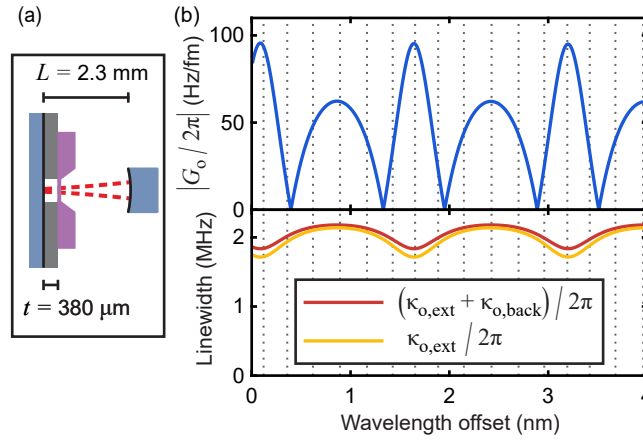


Figure 4.9: **Fully fixed optical cavity** (a) Length scales of cavity used in Ref. [Brubaker et al. 2022]. (b) Simulated cavity operating parameters. The lines indicate the continuum of operating points available if the curved mirror were piezoelectrically actuated, as in Fig. 4.7, and the dotted gray lines indicate the restricted operating parameters permitted by the membrane-perturbed free-spectral-range of a fixed optical cavity.

4.2.5 Thermal environment of the mechanical mode

Cooling the mechanical mode to an occupation of less than one phonon is a requirement for quantum transduction, and this challenge becomes easier by decreasing the rate at which environmental phonons enter the mechanical mode $\gamma_m n_{\text{th}}$. There has been a lot of progress in decreasing γ_m by engineering the density of states of the structure supporting the mechanical resonator of interest, and creating a phononic crystal with a band gap at the desired frequency [Alegre et al. 2011; Yu et al. 2014; Tsaturyan et al. 2014]. For instance, by etching a repeating pattern of holes in the silicon chip surrounding a SiN membrane, the contribution to γ_m from phonons radiating from the membrane mode to the environment through the silicon chip, called “radiation loss,” can be dramatically suppressed [Yu et al. 2014; Tsaturyan et al. 2014]. The next dominant contribution to γ_m becomes “bending loss,” from material loss channels and whose scaling can be calculated from the energy-normalized mode displacement profile, $u(x, y)$, as $\int dx dy (\partial_x^2 u(x, y) + \partial_y^2 u(x, y))^2$ [Reetz et al. 2019]. This loss tends to be concentrated in the region of sharp bending at the edge of the membrane, where it is clamped to the more rigid silicon chip. There has also been a lot of success in reducing this loss contribution [Tsaturyan et al. 2017; Ghadimi et al. 2018], for example by patterning a phononic crystal in the SiN membrane itself, around a “defect” with a modified geometry designed to resonate at a frequency in the band gap of the crystal. The resulting mechanical mode of the defect decays exponentially into the phononic crystal, and exhibits lower overall mode curvature because of this “soft clamping.”

One improvement that facilitated ground-state cooling of our transducer was introducing a single-unit-cell “phononic shield” into the silicon chip supporting the membrane, and that will be detailed further in Ch. 5. A patterned membrane could offer further advantages, including even higher mechanical quality factor through soft clamping, as well as allowing for more unit cells in the same footprint because of the reduced speed of sound in the SiN membrane compared with bulk silicon. Albert Schliesser’s group has successfully operated such soft clamped membranes at mK temperatures [Seis et al. 2022; Planz et al. 2023]. In considering this choice however, thermal

modeling indicated that this would have a steep drawback: though the mechanical mode is less strongly coupled to environmental degrees of freedom, it would see a hotter environment, from reduced thermal conductance because of the modified membrane geometry as well as the spatial distribution of the mechanical mode loss.

Fig. 4.10(a) shows the results of thermal finite-element modeling of an unpatterned membrane similar to that we use in our transducer, with 500 μW of circulating optical power in the cavity, with the boundary of the membrane held to a temperature of $T_{\text{boundary}} = 70 \text{ mK}$ where it would meet the silicon frame. I assume a power absorption coefficient of $A_{\text{SiN}} = 10^{-8}$, and a bulk thermal conductivity of $K_{\text{membrane}} = 14.5 \cdot (T/\text{K})^{1.98} \text{ mW}/(\text{K} \cdot \text{m})$, as obtained by ref. [Leivo and Pekola 1998] for SiN membranes. We measure thermal baths of approximately 70 mK for our transducer when the laser is operating, which (perhaps counterintuitively) is not inconsistent with the center of the membrane rising to a temperature that is significantly higher, because the thermal environment that the mechanical mode sees is dependent on the spatial distribution of its loss, i.e. where the mode curvature is greatest, assuming bending losses dominate. This phenomenon is analogous to the lumped resistances employed at various temperature stages in cryogenic electrical measurement chains to engineer the thermal bath seen by the device being measured. Refs. [Singh and Purdy 2020] and [Shaniv, Reetz, and Regal 2023] explore this more systematically in the context of mechanical modes. The membrane mode is hard-clamped, so has bending loss that is located close to the silicon frame, which because of its thickness being three order of magnitude greater than the membrane, would be better thermalized to the base plate of the fridge.

Fig. 4.10(b) shows similar results for a soft-clamped membrane designed after that in Ref. [Tsaturyan et al. 2017], with the same T_{boundary} , A_{SiN} , and circulating power, and a modified thermal conductivity of $K_{\text{bridges}} = 1.58 \cdot (T/\text{K})^{1.54} \text{ mW}/(\text{K} \cdot \text{m})$, as ref. [Leivo and Pekola 1998] found for SiN bridges with 25 μm width. This geometric dependence remains not well-understood, however, using K_{membrane} with this geometry reduces the highest temperature by roughly a factor of 2, to 340 mK. This geometry not only heats to a higher temperature, but the soft-clamped defect mode would be primarily coupled to the material temperature at the center of the patterned membrane




Figure 4.10: **Thermal modeling of membrane heating from laser absorption** (a) Simulated thermal gradient of an unpatterned, 2 mm square membrane with 500 μW of optical power dissipated at the center, with the boundary held to 70 mK. (b) Left, simulated thermal gradient of a patterned membrane under the same conditions as (a). Right, normalized displacement profile of the defect mode of this structure. Inset shows detail of the central defect. (c) Left, simulated thermal gradient (left) of a 500 μm square membrane, with optical power dissipated in the upper left corner, for optimal readout of the (2,2) mechanical mode. Right, normalized displacement profile of (2,2) mechanical mode.

because of the location of the mode curvature, negating any benefit from the decreased γ_m . Furthermore, because we'd like to operate the transducer at few-kHz-bandwidths to be compatible with superconducting qubit lifetimes, this win in γ_m is less significant, and would still require optical powers comparable to those in the simulation, without an additional improvement to g_o . Albert Schliesser's group is pursuing a narrow-band transducer that leverages the high quality factor of soft-clamped membranes [Catalini, Tsaturyan, and Schliesser 2020; Seis et al. 2022; Planz et al. 2023], and they have made observations consistent with their mode heating to temperatures greater than 1 K at powers of 500 μ W (see supplementary information to Ref. [Page et al. 2021]).

To investigate this the spatial distribution of bending losses more quantitatively, I assign an environmental mode temperature by a bending-weighted average of the simulated thermal distributions:

$$T_{\text{avg}} = \frac{\int T(x, y) \cdot (\partial_x^2 u(x, y) + \partial_y^2 u(x, y))^2 dx dy}{\int (\partial_x^2 u(x, y) + \partial_y^2 u(x, y))^2 dx dy}. \quad (4.21)$$

For the patterned membrane, the environmental temperature seen by the defect mode whose spatial profile is shown in Fig. 4.10(b) is evaluated to be 550 mK. As expected, this is close to the maximum temperature of the center of the membrane, because of the spatial profile of the mode.

Fig. 4.10(c) shows a membrane geometry more similar to that of our transducer, with a width of 500 μ m, and the optical spot located in the corner in order to address the (2,2) mechanical mode, defined as the mode with two antinodes in both directions, whose mode profile is shown. Again the thermal profile is simulated, assuming 500 μ W of optical circulating power, $T_{\text{boundary}} = 70$ mK, and $A_{\text{SiN}} = 10^{-8}$. Though the maximum temperature of 140 mK is located near an antinode of motion, T_{avg} for this mode is calculated to be 71 mK, barely greater than T_{boundary} .

In this thermal model, I've chosen the two parameters A_{SiN} and T_{boundary} . We observe our devices to thermalize to 70-100 mK (depending on the device) once we turn on the locking beam, but any additional optical power from the pump beam does not measurably increase the temperature of the mechanical mode's thermal bath, so I set $T_{\text{boundary}} = 70$ mK. This could be due to some temperature-dependent thermal conductivity linking the silicon frame of the membrane device to

the base plate of the dilution refrigerator, leading to saturation of the mode temperature with increasing dissipated power. The absorption coefficient of silicon nitride membranes at cryogenic temperatures for 1 μm light is still not known, though it is possible to place upper bounds at about the 10^{-6} level. Using $A_{\text{SiN}} = 10^{-8}$ and $T_{\text{boundary}} = 70$ mK reproduces temperatures roughly consistent with those measured in Ref. [Page et al. 2021], though 800 nm light was used in that measurement, and $A_{\text{SiN}}(\lambda)$ need not be the same value at $\lambda = 800$ nm and $\lambda = 1080$ nm, the laser wavelength we use in Ref. [Brubaker et al. 2022]. At the same time, when simulating varying optical powers over the range we used in Ref. [Brubaker et al. 2022], this set of choices replicates the independence of the transducer mode’s thermal environment on the optical power we observed; relative to the 70 mK thermal environment measured at low optical powers, repeating the simulation with the highest optical power measured would increase that temperature by 10%, but because of the high optical damping at this power, this would lead to an increase in the measured mode occupation of merely 1%.

Though there is still uncertainty in the exact value of A_{SiN} , this comparison of the mechanical mode thermal environments due to absorbed optical power in these two designs indicates that the bending loss at the silicon boundary in an unpatterned membrane is a benefit for applications that require high optical power and operate in a dilution refrigerator environment. Transduction goals will be easier to attain with higher bandwidths, and the associated higher laser powers, in a platform with a thermally conductive, phonically shielded silicon frame, as opposed to a phonic crystal engineered directly into the silicon nitride membrane.

4.2.6 Membrane collapse and intentional pinning

In measuring one transducer device with a silicon phononic shield, we observed no microwave resonance at frequencies at which we would expect to see the LC resonance. Optical measurement revealed mechanical resonances whose frequencies were inconsistent with those measured in previous devices and simulation. Finite-element modeling of a “collapsed” membrane, whose capacitor pad was held fixed as a boundary condition in addition to the edge where the membrane is held by the

silicon chip, predicted mode frequencies close to those we observed in measurement (see Fig. 4.11). Having observed qualitative indications that the device had a Hz-scale γ_m , we attempted to measure this purely optomechanically. We generally measure this quantity by measuring the mechanical response function with a swept-sideband measurement, sweeping the microwave pump strength and extrapolating to zero microwave power. The dynamic range of our control over our microwave pump makes it a natural choice for this measurement. To better bound γ_m using only optical measurement, we tuned the laser to a low-coupling wavelength, and swept the laser power, finding a residual damping rate of less than one Hz, including the contribution from the lock beam. This is an interesting and counterintuitive measurement result, as the mechanical quality factor was high, $\sim 10^6$, comparable to only the highest-Q membranes without phononic shields, though the collapse indicated a touch between the membrane and the unpatterned bottom chip. This indicates that contact between the membrane and the bottom chip doesn't preclude having high quality factors—apparently, the touch does not “short out” the phononic shield.

This interesting observation is currently being investigated further to improve the transducer. The transducer used in Ref. [Brubaker et al. 2022] had anomalously low g_e , because the membrane chip bowed away from the bottom chip, and this significantly limited its performance, precluding quantum transduction in the presence of noise from the electrical circuit. By placing a post to pin the membrane to the bottom chip along the node line of the transduction mechanical mode, we hope to achieve better control over smaller capacitor spacings to improve g_e . These spacings are technically difficult to define, because of the aspect ratio of the cm-scale chip to the 100 nm-scale capacitor gap, coupled with thermal contraction associated with cooling to cryogenic temperatures.

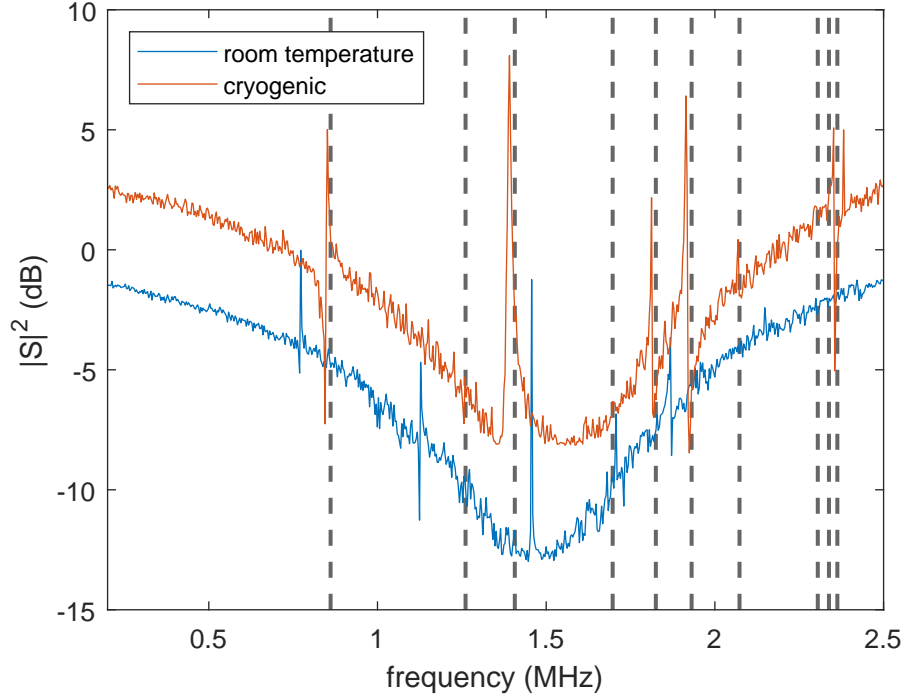


Figure 4.11: **Evidence of a collapsed membrane.** Measured optomechanical amplitude response of the same device operating at room temperature (blue) and cryogenically (orange). The response reveals a MHz-scale reflection dip due to the optical cavity, and kHz-scale optomechanical features exhibiting a Fano lineshape. Whereas the room temperature response aligns with expectations for mechanical frequencies, the frequencies of the mechanical features have changed dramatically at mK temperatures. Simulating a collapsed membrane yields normal mechanical modes at frequencies indicated by the dashed gray lines. Given the position of the optical spot on the membrane and the simulated mode shapes, low optomechanical coupling is expected for the modes at 1.26 MHz, 1.70 MHz, 2.07 MHz, and 2.31 MHz. Note: the large dip in reflection associated with the optical cavity in this measurement could be the result of imperfect modematching between the optical cavity mode and the pump and local oscillator beams, due to an effect discussed in Sec. B.5.

Chapter 5

Optomechanical ground-state cooling in a transducer

In this chapter, I describe the device in which we first optomechanically cooled the mechanical transduction mode to an occupation of less than one phonon, and the measurements that demonstrate its proximity to its quantum ground state [Brubaker et al. 2022]. Under different operating parameters, we achieved a minimum N_{add} of 3.2 photons in upconversion with the same device, which I will also detail in this chapter, with an exploration of why these optimizations required different operating parameters. This device also demonstrated a maximum transduction efficiency $\eta_t = 0.47$. We used the same device under a third suite of operating parameters to realize optical heterodyne readout of the state of a superconducting qubit via the transducer [Delaney et al. 2022]. This result has been described in detail in Rob Delaney’s thesis [Delaney 2022], so in this chapter I will focus on how the qubit-centered transducer measurements in [Delaney et al. 2022] relate to the measurements qualifying the transducer performance in isolation from [Brubaker et al. 2022]. Notably, all the above measurements were accomplished with the transducer running continuously, and the laser both caused minimal increase in noise of the transducer’s superconducting LC circuit 200 μm from the optical cavity mode, as well as negligible backaction on the superconducting qubit system connected to the transducer’s microwave input port.

The device was a fully-fixed optical cavity with a bonded etalon, and measured with a tunable laser. The curved front mirror allowed us to operate this device in the more stable configuration, as the configuration with higher optomechanical coupling proved lossier in this device. We took these measurements in a cooldown that began in December 2020 and ended in April of 2021.

Secs. 5.1 through 5.4 give a description of the experiment.

5.1 Transducer details

This transducer device is depicted in Fig. 4.3(c), and follows the general architecture described in Ch. 4. The mechanical mode used for transduction has resonant frequency $\omega_m/2\pi = 1.45$ MHz. The microwave resonator has resonant frequency $\omega_e/2\pi = 7.938$ GHz, and is coupled to a transmission line at rate $\kappa_{e,\text{ext}}/2\pi = 1.42$ MHz. Its total linewidth κ_e increases from 1.64 MHz to 2.31 MHz as the microwave power incident on the circuit increases. The we operated the optical cavity with a laser of wavelength 1084.4 nm ($\omega_o/2\pi = 277$ THz), where we measure total cavity linewidth $\kappa_o/2\pi = 2.68$ MHz, which is dominated by its external coupling $\kappa_{o,\text{ext}}/2\pi = 2.12$ MHz through the curved mirror. The transducer is operated in an optical-access dilution refrigerator with a 40 mK base-plate temperature. Though the transducer is inherently bidirectional [Andrews et al. 2014], we characterize this device in upconversion, as that what is important for the network architectures we are working to implement, as described in Ch. 3.

5.1.1 Mechanical isolation with phononic shield

The single-unit-cell phononic shield patterned into the silicon substrate chip surrounding the membrane ensures a lower density of substrate modes around the transduction mode's frequency and grants it a low intrinsic loss rate of $\gamma_m/2\pi = 113$ mHz. The phononic shield is made up of $900\text{ }\mu\text{m} \times 900\text{ }\mu\text{m}$ masses joined by $100\text{ }\mu\text{m} \times 550\text{ }\mu\text{m}$ tethers, with the 100 nm-thick silicon nitride membrane supported on the central mass (Fig. 5.1(a)). The transverse dimensions of the membrane ($500\text{ }\mu\text{m} \times 500\text{ }\mu\text{m}$) determine the resonant frequency $\omega_m/2\pi = 1.45$ MHz of the vibrational mode with two antinodes in each direction (Fig. 5.1(b)), which we label the (2,2) mode [Andrews 2015]. The mass and tether dimensions are chosen to center the phononic bandgap on ω_m , and finite-element simulations of the shield with periodic boundary conditions predict a bandgap between 1.10 and 1.98 MHz.

The overall size of the top chip is constrained by the flip-chip electromechanical circuit ar-

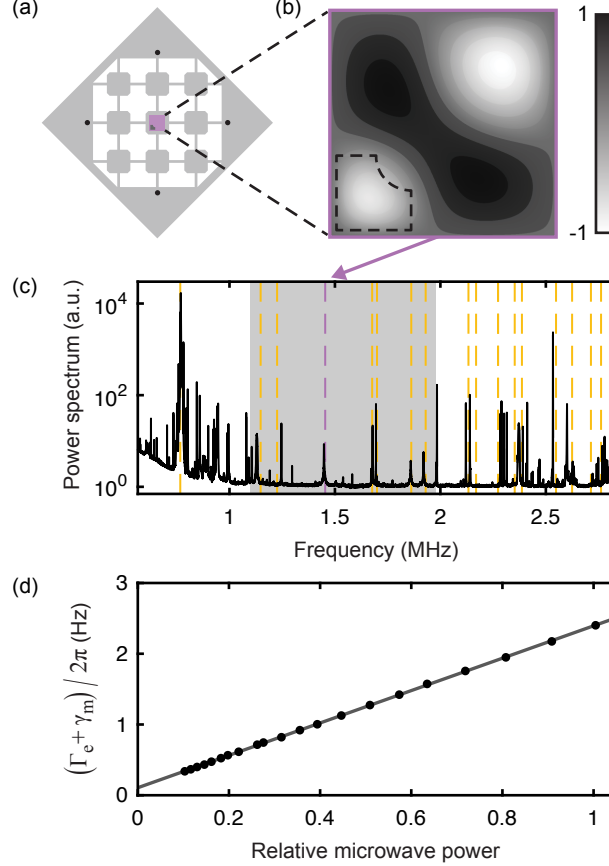


Figure 5.1: **Mechanical design and characterization.** (a) Phononic shield design. Black circles (not to scale) indicate location of 20 μm diameter posts that determine the nominal capacitor pad separation. (b) Normalized displacement of the 1.45 MHz vibrational mode from finite-element simulation of the membrane and capacitor pad (dashed line). (c) Optically measured mechanical spectrum. The shaded region indicates the bandgap obtained from a simulation with periodic boundary conditions. Expected frequencies of membrane modes (dashed lines) are obtained from finite-element simulation. (d) Total damping obtained from electromechanical ring-down measurements vs. microwave pump power. Extrapolation to zero pump power yields an intrinsic mechanical damping rate of $\gamma_m/2\pi = 108$ mHz.

chitecture: larger chips make it difficult to control the capacitor pad separation d that determines G_e , the coefficient relating the microwave mode frequency shift to membrane displacement. This constraint limits the number of unit cells in the phononic shield, but the design nevertheless yields a vibrational spectrum with relatively few substrate modes in the vicinity of the (2,2) mode (Fig. 5.1(c)). Membranes embedded in similar phononic crystal structures with more unit cells have been shown to thermalize to $T \lesssim 100$ mK in the presence of laser light [Kampel et al. 2017].

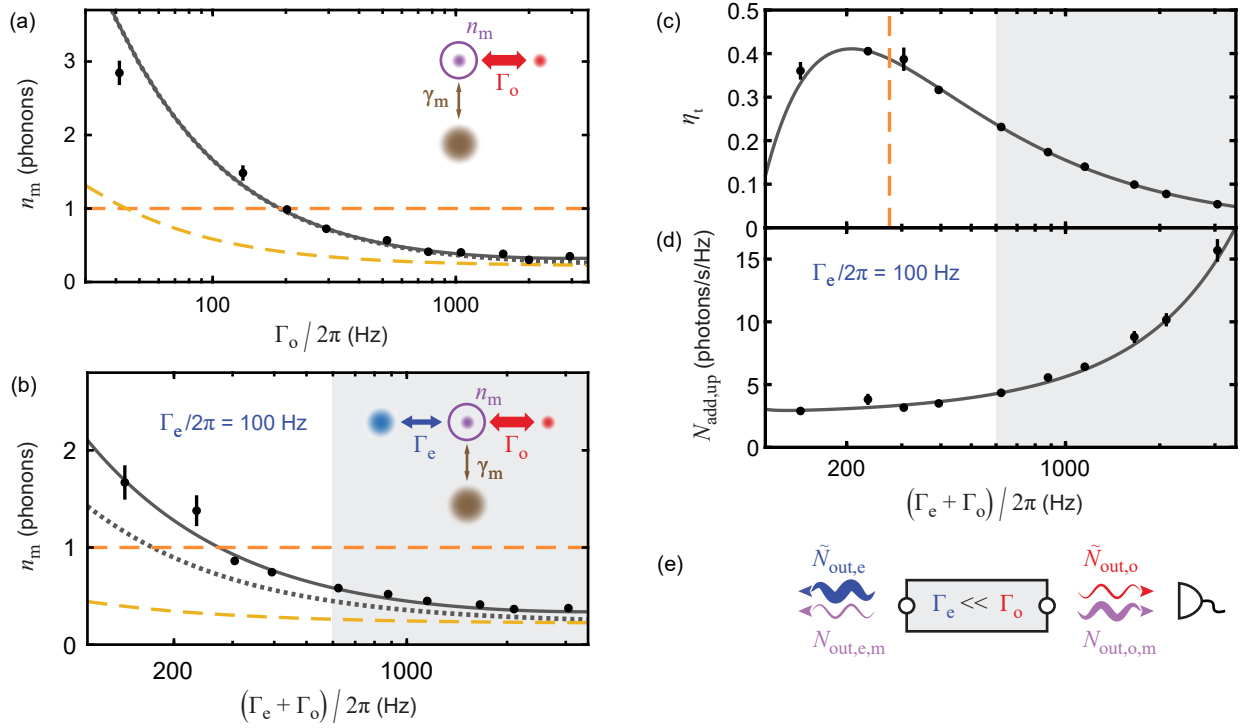


Figure 5.2: Ground-state cooling of electro-optomechanical transducer. (a) Membrane mode occupancy n_m vs. optomechanical damping rate Γ_o with electromechanical damping $\Gamma_e = 0$. Points are data, solid gray line is a fit, and dotted gray line shows the predicted behavior in the absence of technical noise. Orange dashed line indicates $n_m = 1$, and yellow dashed line indicates the backaction limit (increase at low Γ_o due to backaction from auxiliary optical cavity lock beam). Inset illustrates optomechanical sideband cooling, where external blurred discs represent baths (with size indicating occupancy) to which the membrane mode (purple circle) is coupled by the corresponding damping rates (arrows, with width indicating magnitude). The membrane mode occupancy (purple blurred disk) is determined by the weighted average of bath occupancies. (b) n_m vs. $\Gamma_e + \Gamma_o$, sweeping Γ_o at fixed electromechanical damping $\Gamma_e/2\pi = 100$ Hz, with curves color-coded as in panel (a). Inset illustrates relative bath occupancies and damping rates for electro-optomechanical sideband cooling in the gray shaded region where $\Gamma_e \ll \Gamma_o$. (c) Transducer efficiency η_t vs. $\Gamma_e + \Gamma_o$, sweeping Γ_o as in (b), with fit. Orange dashed line marks the damping at which $n_m = 1$. (d) Upconversion added noise $N_{\text{add,up}}$ vs. $\Gamma_e + \Gamma_o$, with theory curve. (e) Schematic representation of transducer output noise in the $\Gamma_e \ll \Gamma_o$ regime shaded gray in panels (b), (c), and (d). The box represents the transducer, and wavy arrows represent contributions to the microwave and optical output noise around mechanical resonance, from membrane mode occupancy (purple) and technical noise (red/blue), with thickness indicating noise density. Photodetector indicates optical readout. All error bars represent one standard deviation.

5.2 Optomechanical ground-state cooling

In thermal equilibrium at temperature T_{eq} , the membrane mode is coupled to a bath with occupancy $n_{\text{th}} = k_{\text{B}}T_{\text{eq}}/\hbar\omega_{\text{m}}$ by its intrinsic decay rate γ_{m} . Optomechanical damping couples the membrane mode to a bath with occupancy $n_{\text{o}} = n_{\text{min,o}} + n_{\text{eff,o}} \ll n_{\text{th}}$, where the first term represents the quantum backaction limit [Peterson et al. 2016] from imperfect sideband resolution and the second term is an effective optical mode thermal occupancy due to pump power-dependent technical noise (see Sec. 4.1). Analogous effects contribute to the electromechanical bath occupancy n_{e} . The membrane mode occupancy n_{m} is given by the weighted average of the baths to which it is coupled,

$$n_{\text{m}} = \frac{\gamma_{\text{m}}n_{\text{th}} + \Gamma_{\text{e}}n_{\text{e}} + \Gamma_{\text{o}}n_{\text{o}}}{\Gamma_{\text{T}}}, \quad (5.1)$$

leading to sideband cooling of the membrane mode (Fig. 5.2(a) and (b), insets). Given sufficiently low n_{o} , Γ_{o} can be increased to overwhelm the rate at which environmental phonons enter the mechanical mode $\gamma_{\text{m}}n_{\text{th}}$, cooling the mechanical mode to close to its ground state.

We first demonstrate optomechanical ground-state cooling of the membrane mode with the microwave pump absent (Fig. 5.2(a)), to show that our transducer realizes the successes of past membrane optomechanical systems [Peterson et al. 2016; Underwood et al. 2015] even with the additional design complexity required for electromechanical coupling to superconducting circuitry. We measure the optomechanical sideband ratio while varying the damping Γ_{o} , and correct for squashing effects due to laser phase noise (see Refs. [Jayich et al. 2012; Safavi-Naeini et al. 2013] and Secs. 4.1 and 5.7) to obtain the membrane mode occupancy n_{m} at each value of Γ_{o} . We then fit this data to Eq. (5.1) with $\Gamma_{\text{e}} = 0$ and the equilibrium occupancy n_{th} as a fit parameter. We also introduce an additional fit parameter $a_{\text{o}} = n_{\text{eff,o}}/\Gamma_{\text{o}}$ to model the optical mode occupancy due to phase noise, and include a fixed term to account for the effects of a weak near-resonant beam used to lock the optical cavity (see Sec. 5.5).

The fit (solid gray line) yields $n_{\text{th}} = 1000 \pm 90$, implying that the membrane mode equilibrates to $T_{\text{eq}} = 70$ mK in the presence of laser light, consistent with an independent calibration based on

sweeping the temperature of the cryostat base plate (see Sec. 5.7). Phase noise parameterized by $a_o = (2.8 \pm 1.6) \times 10^{-6} \text{ Hz}^{-1}$ causes the fit to deviate slightly at high damping from the behavior expected from thermal noise with the best-fit value of n_{th} and imperfect sideband resolution alone (dotted gray line). The good agreement with the fit, without additional power-dependent terms, indicates that the laser-induced heating of the membrane mode responsible for its elevated equilibrium temperature T_{eq} saturates quickly at very low power and does not preclude cooling below $n_m = 1$ (orange dashed line) for $\Gamma_o/2\pi > 190 \text{ Hz}$. With large Γ_o , we reach a minimum mode occupancy of $n_m = 0.32$, within a factor of 1.5 of the backaction limit (yellow dashed line).

To enable transduction, we now introduce a microwave pump with fixed electromechanical damping $\Gamma_e/2\pi = 100 \text{ Hz}$, and again sweep Γ_o in order to ground-state cool the membrane mode (Fig. 5.2(b)). We observe that n_m deviates further from the expected behavior from thermal noise and imperfect sideband resolution alone (dotted gray line) than in the purely optomechanical case, indicating an additional source of technical noise. Fixing n_{th} and a_o at the best-fit values from the pure optomechanical ground-state cooling data, we find good agreement with Eq. (5.1) with effective thermal occupancy $n_{\text{eff,e}} = 0.8 \pm 0.1$, consistent with $n_{\text{eff,e}} = 0.77 \pm 0.01$ obtained from an independent heterodyne measurement of the noise in the transducer's microwave output spectrum (see Sec. 5.8). This good agreement, with $n_{\text{eff,e}}$ independent of laser power, indicates that the additional source of technical noise is coupled in via the microwave pump. With sufficiently high Γ_o , the membrane mode can again be cooled close to the backaction limit, reaching a minimum occupancy of $n_m = 0.34$.

5.3 Transducer characterization

We next consider how the electro-optomechanical device performs as a transducer when the optomechanical damping Γ_o is varied at constant electromechanical damping Γ_e , as in Fig. 5.2(b). We first consider its transduction efficiency, given by

$$\eta_t = \eta_M \frac{4\Gamma_e\Gamma_o}{\Gamma_T^2}, \quad (5.2)$$

where the matched efficiency $\eta_M = \epsilon \frac{\kappa_{o,\text{ext}}}{\kappa_o} \frac{\kappa_{e,\text{ext}}}{\kappa_e}$ is the maximum efficiency achievable in the presence of imperfect optical cavity modematching ϵ and loss in the electromagnetic resonators. The transducer achieves this maximum efficiency when $\Gamma_e = \Gamma_o \gg \gamma_m$.

We measure the transducer efficiency η_t using the method developed in Ref. [Andrews et al. 2014], and fit the data to Eq. (5.2) with the matched efficiency η_M as the only fit parameter (Fig. 5.2(c)). The fit yields $\eta_{M,\text{fit}} = 43 \pm 1\%$, in tension with $\eta_{M,\text{exp}} = \epsilon \frac{\kappa_{o,\text{ext}}}{\kappa_o} \frac{\kappa_{e,\text{ext}}}{\kappa_e} = 55\%$ expected from independent measurements of the optical cavity modematching ϵ and the overcoupling fractions $\kappa_{\text{ext}}/\kappa$ of the two electromagnetic resonators. We suspect the measured values of η_t were artificially suppressed (see Sec. 5.7) and that the true matched efficiency during these measurements was $\eta_{M,\text{exp}}$ rather than $\eta_{M,\text{fit}}$. Nonetheless, adopting the fit value as a conservative measure of transducer performance, we find $\eta_t = 38\%$ at the damping for which the membrane mode occupancy is reduced below $n_m = 1$.

We can infer the transducer's input-referred upconversion added noise $N_{\text{add,up}}$ from the above measurements of the membrane mode occupancy n_m and the transducer efficiency η_t . We also take into account an additional contribution to $N_{\text{add,up}}$ from white noise $\tilde{N}_{\text{out,o}}$ measured at the transducer's optical output port, which arises from the technical noise responsible for the effective optical mode occupancy $n_{\text{eff,o}}$ (see Sec. 4.1), and plot the results in Fig. 5.2(d). To understand the observed increase in added noise when $\Gamma_o \gg \Gamma_e$, it is helpful to consider a schematic representation (Fig. 5.2(e)) of the noise emerging from the transducer in this regime, shaded gray in Figs. 5.2(b), (c), and (d). When the optomechanical damping Γ_o dominates, almost all the noise arising from thermal occupancy of the membrane mode is routed to the transducer's optical output port. This contribution $N_{\text{out,o,m}}$ to the optical output noise density asymptotes at high Γ_o . But to obtain the input-referred added noise $N_{\text{add,up}}$ we must divide by the transducer efficiency, which decreases with increasing Γ_o , resulting in increased added noise. The fact that $\tilde{N}_{\text{out,o}}$ increases with increasing Γ_o further exacerbates the added noise.

More precisely, the transducer's input-referred upconversion added noise is given in units of

photons/s/Hz (or more simply photons) by

$$\begin{aligned} N_{\text{add,up}} &= \frac{N_{\text{out,o,m}} + \tilde{N}_{\text{out,o}}}{\mathcal{A}_e \mathcal{A}_o \eta_t} \\ &= \frac{\gamma_m n_{\text{th}} + \Gamma_e n_e + \Gamma_o n_o}{\mathcal{A}_e \frac{\kappa_{e,\text{ext}}}{\kappa_e} \Gamma_e} + \frac{\tilde{N}_{\text{out,o}}}{\mathcal{A}_e \mathcal{A}_o \eta_t}, \end{aligned} \quad (5.3)$$

where $\mathcal{A}_e = 1 + n_{\text{min,e}}$ is the transducer gain from imperfect electromechanical sideband resolution, and \mathcal{A}_o is defined analogously. This is equivalent to Eq. 4.13, but includes the contribution from technical noise on the laser, $\tilde{N}_{\text{out,o}}$. The electromechanical and optomechanical bath occupancies n_e and n_o depend implicitly on the damping rates Γ_e and Γ_o , respectively, because of technical noise. In the absence of technical noise in the optomechanical system, $n_o \rightarrow n_{\text{min,o}}$ and the last term vanishes. Then it is clear that increasing the optomechanical damping Γ_o can only ever increase $N_{\text{add,up}}$, and that quantum-enabled upconversion would require electromechanical ground-state cooling.

As the added noise is lowest at small Γ_o in Fig. 5.2(d), we now fix $\Gamma_o/2\pi = 85$ Hz and vary Γ_e to minimize $N_{\text{add,up}}$. We begin by measuring transducer efficiency as a function of Γ_e (Fig. 5.3(a)). We observe that the microwave mode linewidth κ_e increases with increasing microwave pump power P_e , an effect likely related to the elevated thermal occupancy $n_{\text{eff,e}}$ observed in the electro-optomechanical experiment in Fig. 5.2(b). We account for this power-dependence in fitting the efficiency measurements (see Sec. 5.7), and obtain a peak efficiency $\eta_{t,\text{max}} = 47 \pm 1\%$ at $\Gamma_e/2\pi = 75$ Hz, close to the expected value of 49%.

We then measure noise in the transducer's optical output spectrum while sweeping the electromechanical damping Γ_e over this same range (Fig 5.3(b)), and combine the results with the efficiency measurements to obtain $N_{\text{add,up}}$ as a function of electromechanical damping. Fitting the data, we find that at $\Gamma_e/2\pi = 135$ Hz the added noise reaches a minimum value of $N_{\text{add,up}} = 3.2 \pm 0.1$ photons/s/Hz, an order of magnitude improvement relative to previously reported measurements on a similar device [Higginbotham et al. 2018]. In this fit, we parameterize the effective microwave mode occupancy as $n_{\text{eff,e}} = a_e \Gamma_e + b_e$ and again including an additional fixed term to account for the optical cavity lock beam (see Sec. 5.7). The fit yields $n_{\text{th}} = 980 \pm 30$

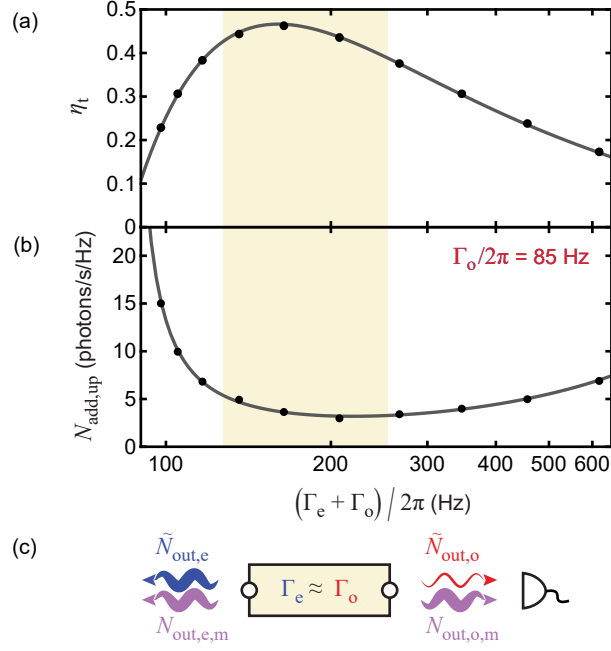


Figure 5.3: Minimizing input-referred added noise. (a) Transducer efficiency η_t vs. $\Gamma_e + \Gamma_o$, sweeping electromechanical damping Γ_e at fixed optomechanical damping $\Gamma_o/2\pi = 85$ Hz, with fit. (b) Upconversion added noise $N_{\text{add,up}}$ vs. $\Gamma_e + \Gamma_o$ at $\Gamma_o/2\pi = 85$ Hz, with fit. (c) Schematic representation of transducer output noise in the $\Gamma_e \approx \Gamma_o$ regime shaded yellow in panels (a) and (b), where the box represents the transducer, wavy arrows represent contributions to the microwave and optical output noise near mechanical resonance, and the photodetector indicates optical readout, as in Fig. 5.2(e). Error bars representing one standard deviation are smaller than the size of the points.

and a_e and b_e consistent with independent microwave noise measurements (see Sec. 5.8).

The dominant contributions to the added noise at the minimum are 1.0 photons/s/Hz from residual thermal motion of the membrane and 1.4 photons/s/Hz from the effective occupancy of the microwave mode, with several smaller sources responsible for the remaining 0.8 photons/s/Hz (see Sec. 5.7). Fig. 5.3(c) shows a schematic representation of transducer noise around the added noise minimum, where $\Gamma_e \approx \Gamma_o$ (yellow shaded region in Figs. 5.3(a) and (b)). The noise arising from membrane mode thermal occupancy is divided roughly equally between the two transducer ports, but the mode is not in its quantum ground state due to the smaller total damping: $n_m = 1.5$ phonons when $N_{\text{add,up}}$ is minimized. Eq. (5.3) indicates that in the absence of technical noise, the added noise would continue to decrease with increasing electromechanical damping even in the $\Gamma_e \gg \Gamma_o$ regime where the efficiency is small. In our experiment, the power-dependent effective microwave mode occupancy $n_{\text{eff,e}}$ causes n_m to increase at large Γ_e , resulting in an optimum value of Γ_e at which $N_{\text{add,up}}$ is minimized.

Finally, we consider how the transducer's downconversion added noise $N_{\text{add,down}}$ (not measured in this work) would scale with microwave and optical pump powers. Added noise is in general not the same in downconversion as in upconversion. In the absence of microwave technical noise, optomechanical ground-state cooling would be a sufficient condition for quantum-enabled downconversion, independent of the strength of the microwave pump. However, the technical noise responsible for the effective microwave mode thermal occupancy $n_{\text{eff,e}}$ also generates a white noise contribution $\tilde{N}_{\text{out,e}}$ to the microwave noise density at the external port of the microwave circuit. As indicated in Figs. 5.2(e) and 5.3(c), this excess noise is significantly larger than the analogous optical technical noise $\tilde{N}_{\text{out,o}}$, and it precludes quantum-enabled downconversion with this transducer.

5.4 Effects of Laser Illumination

As indicated by the data presented in Figs. 5.2 and 5.3, the transducer performance is chiefly limited by the pump power-dependent effective thermal occupancy $n_{\text{eff,e}}$ of the superconducting microwave circuit. The continuously applied optical pump, however, had no discernible effect

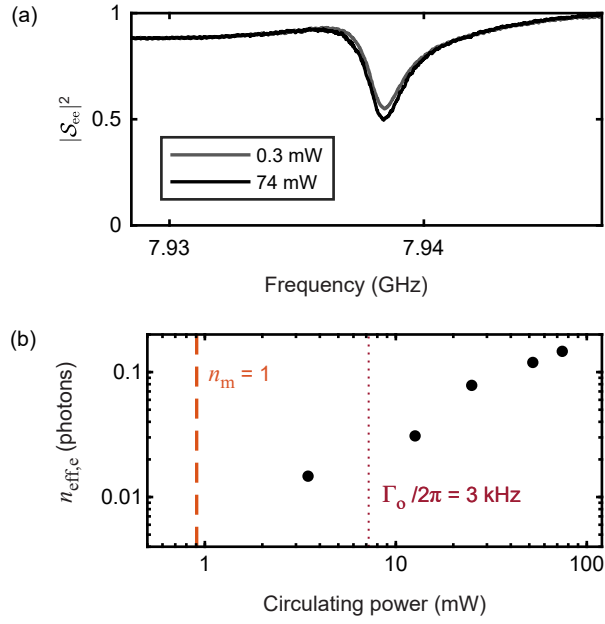


Figure 5.4: **Effects of laser illumination on microwave circuit.** (a) Microwave circuit power reflection measurements with 0.3 mW of circulating lock light (gray) and with 74 mW of circulating pump light (black). This pump power is a factor of 10 greater than that required for the highest damping rate Γ_o shown in Fig. 5.2, and two orders of magnitude greater than that required for optomechanical ground-state cooling. (b) Laser-induced microwave circuit noise vs. circulating power in the optical cavity. The noise remains negligible up to very large circulating power.

on the microwave circuit during the experiments described above, in striking contrast to other platforms in which pulsed operation is required to avoid microwave circuit heating [Mirhosseini et al. 2020; Fu et al. 2021; Stockill et al. 2022; Sahu et al. 2022]. Indeed, laser illumination has very little effect on the microwave circuit even with much larger optical pump power. The gray curve in Fig. 5.4(a) shows the circuit’s power reflection coefficient $|\mathcal{S}_{ee}(\omega)|^2$ when the only light in the optical cavity is 0.3 mW of circulating power from the lock beam, while the black curve shows $|\mathcal{S}_{ee}(\omega)|^2$ in the presence of 74 ± 8 mW of circulating pump power. The laser-induced frequency shift and increase in internal loss are barely perceptible.

By measuring the noise emerging from the external port of the superconducting circuit over a bandwidth much broader than Γ_T , we can directly probe laser-induced heating of the microwave mode as a function of the circulating power in the optical cavity (Fig. 5.4(b)), with no microwave pump present. In the transducer ground-state cooling measurements shown in Fig. 5.2(b), the maximum optomechanical damping of 3 kHz was obtained with 7 mW of circulating power, indicating that damping (and thus transducer bandwidth) could be further increased by more than an order of magnitude while maintaining $n_{\text{eff,e}} < 0.15$. This insensitivity of the superconducting circuit to optical illumination is a consequence of our modular transducer design, and has proven advantageous in work integrating superconducting qubits with electro-optic transducers [Delaney et al. 2022].

5.5 Experimental setup

The transducer is operated in a CryoConcept Horizontal 200 dilution refrigerator with free-space optical access and ^4He precooling to minimize vibrations (previously used in Ref. [Higginbotham et al. 2018]). Optical access to the device is provided by fused silica windows that filter room-temperature thermal radiation to avoid heating of the cryostat base plate [Kuhn et al. 2014], and the cryostat reaches a base temperature of $T_{\text{bp}} = 40$ mK. A back port enables cavity transmission measurements for beam alignment and membrane imaging. The transducer device shares space on the cryostat base plate with a circuit QED module used in Ref. [Delaney et al. 2022] but

with ultralow phase noise (Rohde & Schwartz SMA100B with option SMAB-B711), is injected into the cryostat along with a phase-coherent cancellation tone. The amplitude and phase of the cancellation tone are adjusted to minimize the power routed towards the microwave measurement chain, to prevent saturation of the cryogenic preamplifier (a LNF LNC4.8C HEMT) by the high-power microwave pump. Demodulation of microwave signals is performed by an effective heterodyne detector comprising a homodyne measurement followed by AC-coupled amplification and modulation of the baseband signal up to 10 MHz.

The most significant change relative to the optical control and measurement system of Ref. [Higginbotham et al. 2018] is the use of a widely tunable external cavity diode laser (Toptica Photonics CTL 1050) in place of an Innolight Mephisto Nd:YAG laser to source the optical pump. Wavelength tuning is necessary to compensate for the absence of piezoelectric actuators to adjust the length of the optomechanical cavity. For this experiment we operate the laser at a wavelength of 1084 nm, out of an abundance of caution of any heating that could result from light leaking out the back mirror and being absorbed by the silicon chip. At room temperature, silicon is absorptive at our operating wavelengths of around 1060-1080 nm, though the absorption coefficient depends strongly on both wavelength and temperature [Macfarlane et al. 1958]. At the operating temperatures of our transducer, we have measured a small increase in mechanical occupation when injecting 600 μW into the back side of the mirror, off resonance, but it was much less significant than increasing a resonant beam from 100 nW to 1 μW , which are typical operating powers for our device. Regardless, because our setup is relatively frequency-agnostic, we decided to operate at 1080 nm. The CTL is placed in an acoustically isolated enclosure to reduce its sensitivity to environmental disturbances, and locked to a filter cavity with 80 kHz linewidth [Purdy et al. 2012] via feedback to the diode current and piezo actuators on the laser cavity using the Pound-Drever-Hall (PDH) technique, using Toptica fast analogue feedback (FALC110) and PDH locking (PDD110) modules. Transmission through the filter cavity reduces the laser phase noise at detuning ω_m from the carrier by 37 dB. We use a single-frequency fiber amplifier (Nufern Nuamp NUA-1064-PB-0005-C1) to amplify the CTL output beam before the filter cavity.




Figure 5.6: **Cavity birefringence and beam detunings.** Frequency diagram of the pump (red) and locking beam (yellow), and cavity resonant frequencies. The cavity resonant frequency associated with each beam's respective polarization state is of the corresponding color. Given a cavity birefringence Δ_B , there is experimental freedom to choose which beam will address the higher frequency resonance in order to avoid unintentional antidamping due to crosstalk or to avoid the locking beam from being close to the transduced signal.

The beam transmitted through the filter cavity is then frequency-shifted by a double-pass acousto-optic modulator (AOM) and split three ways: one beam (red in Fig. 5.5) provides the optical pump to operate the transducer, a second (yellow) is used to lock the frequency of the optomechanical cavity to that of the laser, and a third (maroon) provides a local oscillator (LO) for balanced heterodyne detection. The relative detunings of the three beams are controlled by three additional AOMs. The incident pump beam and lock beam are orthogonally polarized to route the beams emerging from the cryostat to separate detectors (see Sec. B.2). The detuning of the pump beam from the cavity mode is thus given by $\Delta_o = (\omega_{\text{pump}} - \omega_{\text{lock}} + \Delta_{\text{lock}}) \pm \Delta_B$, where Δ_{lock} is the detuning of the locking beam from the cavity mode, Δ_B is the cavity birefringence, and the sign depends on whether the lock beam addresses the higher- or lower-frequency cavity mode, as illustrated in Fig. 5.6. The LO beam is detuned from the pump beam by $\Delta_{\text{LO}} = (\omega_{\text{LO}} - \omega_{\text{pump}}) = -2\pi \times 12.9 \text{ MHz}$.

The frequency of the incident lock beam is locked to a TEM_{00} mode of the optical cavity using PDH feedback. (For a very good introduction to this useful technique, please see Ref. [Black 2000]). An electro-optic modulator (EOM) imprints phase modulation sidebands on the lock beam, and the PDH error signal obtained by demodulating the detected lock beam at the sideband frequency

is used to apply feedback in parallel to the double-pass AOM and a piezoelectric actuator on one of the filter cavity mirrors. The AOM channel has high bandwidth but limited range, while the slow feedback to the filter cavity length eliminates drift of the laser frequency relative to the frequency of the cavity mode (see Sec. B.6). The incident lock beam power is $P_{\text{lock}} = 20 \text{ nW}$, and the PDH error signal bias is adjusted to keep the lock beam slightly red-detuned from the cavity, with $\Delta_{\text{lock}}/2\pi$ ranging from -30 kHz to -80 kHz , to avoid optomechanical instability.

For greater detail on the optical measurement setup used to align to and measure the transducer, please see Appendix B, which describes a setup used later, to measure transducer devices in a Bluefors LD400 cryogen-free dilution refrigerator. The measurement setups in that system is very similar to that used in this work with the CryoConcept Horizontal 200. The main difference is that the setup used with the later system required additional considerations in order to compensate for the much larger change in cavity position in the Bluefors LD400 system upon cooling down cryogenically.

A Zurich Instruments HF2LI lock-in amplifier is used for all data acquisition. To characterize the transducer and measure its efficiency, we use the HF2LI to synthesize swept MHz-frequency tones for single-sideband (SSB) modulation of the transducer pumps and demodulate the heterodyne detector output signals to recover these tones (see Sec. 5.6.3). For noise measurements, we use the HF2LI to digitize time traces of the noise at the heterodyne detector outputs, and then compute the noise variance or power spectral density. See Ref. [Higginbotham et al. 2018] for further details of signal processing with the HF2LI.

Fig. 5.7 shows photographs of a transducer device mounted in the CryoConcept Horizontal 200 dilution refrigerator. The qubit module described in Ref. [Delaney et al. 2022] was not yet mounted in the left photograph, but is shown on the right photograph. Fig. 5.8 shows a rendition of the transducer assembly, including positions of the chip, the curved mirror, and the input-coupling lenses.




Figure 5.7: **Photographs of transducer mounted in cryostat** Left: the transducer device is housed in the cylindrical structure made of Invar. Reflection measurements are made through the input coupling lens mounted on the front right side. Right: the transducer is enclosed in its OFHC copper radiation shield. The circuit QED module described in Sec. 5.9 is installed directly to the left of the transducer.




Figure 5.8: **Rendition of transducer assembly**

5.6 Details of transducer characterization

Transducer device parameters are summarized in Table 5.1. Measurements of these parameters are detailed below.

5.6.1 Microwave circuit characterization

We characterize the microwave circuit by measuring the reflection scattering parameter $\mathcal{S}_{ee}(\omega)$ encoding its amplitude and phase response as a function of detuning from resonance. When normalized to the off-resonance level, this scattering parameter is

$$\mathcal{S}_{ee}(\omega) = 1 - \frac{\kappa_{e,\text{ext}}}{\kappa_e/2 - i(\omega - \omega_e)}. \quad (5.4)$$

We use a vector network analyzer to generate a weak swept probe tone and inject it into the fridge together with the strong pump tone at fixed detuning $\Delta_e = -\omega_m$ from ω_e . Sweeping the microwave pump power P_e and fitting $\mathcal{S}_{ee}(\omega)$ then allows us to measure the power-dependence of the microwave circuit's internal loss $\kappa_{e,\text{int}}$ (plotted in Fig. 5.12(a) in Sec. 5.8) as well as the power-independent microwave circuit external coupling rate $\kappa_{e,\text{ext}}/2\pi = 1.42$ MHz.

For the data shown in Fig. 5.3, the incident microwave pump power was swept over the range $0.2 \text{ nW} < P_e < 10 \text{ nW}$, and the circuit's total linewidth $\kappa_e/2\pi = (\kappa_{e,\text{ext}} + \kappa_{e,\text{int}})/2\pi$ varied between 1.64 MHz and 2.31 MHz over this range.

5.6.2 Optical cavity characterization

We also characterize the optical cavity by measuring its reflection response, and the relevant scattering parameter is formally analogous to $\mathcal{S}_{ee}(\omega)$. We tune the laser wavelength to put the lock beam close to resonance with the optomechanical cavity, and then sweep the frequency of the laser relative to that of the optomechanical cavity by ramping the voltage applied to a piezoelectric actuator on the filter cavity to which the laser is locked (see Sec. 5.5). This scheme measures $|\mathcal{S}_{oo}(\omega)|^2$, with the horizontal axis calibrated using the phase modulation sidebands on the lock beam. Fitting the data then yields $\kappa_o/2\pi = 2.68 \pm 0.05$ MHz for the optical cavity linewidth.

Table 5.1: Electro-optic transducer parameters. The range of values for the microwave circuit linewidth κ_e and the electromechanical transducer gain \mathcal{A}_e correspond to the range of pump powers used in Fig. 5.3. For the optical cavity modematchings and transducer efficiency, the first (second) numbers correspond to the data shown in Fig. 5.2 (Fig. 5.3)

Parameter	Symbol	Value
Optical cavity frequency	ω_o	$\omega_o/2\pi = 277$ THz
Optical cavity external coupling	$\kappa_{o,\text{ext}}$	$\kappa_{o,\text{ext}}/2\pi = 2.12$ 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	$\kappa_{e,\text{ext}}$	$\kappa_{e,\text{ext}}/2\pi = 1.42$ 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)	$\eta_{\text{M}} (\eta_{\text{t,max}})$	$\eta_{\text{M}} = 55\%$ ($\eta_{\text{t,max}} = 49\%$)
Optomechanical transducer gain	\mathcal{A}_o	$\mathcal{A}_o = 1.22$
Electromechanical transducer gain	\mathcal{A}_e	$\mathcal{A}_e = 1.08 - 1.16$

To optimize and measure the modematching of the cavity mode to the pump or LO beams, ϵ_{PC} and ϵ_{CL} , respectively (see Secs. 4.1.2 and B.4), we inject an additional probe beam into the cavity through the low-transmission back mirror and steer the transmitted mode to a detector where it interferes with the beam of interest. We then equalize the power of the beams, measure the visibility \mathcal{V} of the interference fringes that result from a relative detuning between the beams, and define the associated modematching as \mathcal{V}^2 .

The modematching factors could change slightly over the course of a cooldown as beam alignments drifted. The data shown in Fig. 5.2 in the main text was acquired immediately after the modematchings were reoptimized and measured to be $\epsilon_{\text{PC}} = 0.86$ and $\epsilon_{\text{CL}} = 0.91$. Both before and after the acquisition of the data shown in Fig. 5.3, the modematchings were measured to be $\epsilon_{\text{PC}} = 0.80$ and $\epsilon_{\text{CL}} = 0.79$. Once ϵ_{PC} and the cavity linewidth κ_{o} are known, the external coupling rate $\kappa_{\text{o,ext}}/2\pi = 2.12$ MHz is inferred from the depth of dip on resonance in a measurement of $|\mathcal{S}_{\text{oo}}(\omega)|^2$ in which the pump beam is used to probe the cavity. The modematching of the lock beam to the cavity mode was not directly measured. We infer $\epsilon_{\text{lock}} = 0.85$ using the depth of dip in the lock beam $|\mathcal{S}_{\text{oo}}(\omega)|^2$ measurement together with the measured value of $\kappa_{\text{o,ext}}$.

One final modematching factor that does not involve the cavity mode is also relevant to the calibration of the efficiency (see Sec. 5.7.1). We define ϵ_{PL} as the modematching of the LO beam to the promptly reflected pump beam with the cavity unlocked. We measured this modematching to be $\epsilon_{\text{PL}} = 0.75$ for the data shown in Fig. 5.2 and $\epsilon_{\text{PL}} = 0.79$ for the data shown in Fig. 5.3.

For the data shown in Fig. 5.2, the optical pump power incident on the cavity was swept over the range $20 \text{ nW} < P_{\text{o}} < 1.7 \text{ }\mu\text{W}$. For the data shown in Fig. 5.4, the relevant quantity is not the incident power P_{o} but rather the intracavity circulating power

$$P_{\text{circ}} = \Delta\nu_{\text{FSR}} \kappa_{\text{o,ext}} \left(\frac{\epsilon_{\text{PC}} P_{\text{o}}}{\kappa_{\text{o}}^2/4 + \Delta_{\text{o}}^2} + \frac{\epsilon_{\text{lock}} P_{\text{lock}}}{\kappa_{\text{o}}^2/4 + \Delta_{\text{lock}}^2} \right), \quad (5.5)$$

where $\Delta\nu_{\text{FSR}} = c/2L$ is the free spectral range of the optical cavity in frequency units. Uncertainty in P_{circ} comes from uncertainty in the cavity length L and uncertainty in the measured value of the optical insertion loss within the cryostat.

5.6.3 Optomechanical and electromechanical characterization

To study the vibrational modes of the membrane, we exploit the phenomenon of optomechanically induced transparency (OMIT), wherein $\mathcal{S}_{oo}(\omega)$ is modified by optomechanical interference effects for a swept probe tone coherent with the pump field [Aspelmeyer, Kippenberg, and Marquardt 2014], as well as the analogous phenomenon of electromechanically induced transparency (EMIT). We generate these coherent probe tones using SSB modulation of either the microwave pump (for EMIT measurements) or the RF drive to the AOM that controls the optical pump beam detuning (for OMIT measurements; see Sec. 5.5), with the pump frequency fixed in both cases. These measurements reveal the vibrational spectrum of the membrane over a wide frequency range (see Fig. 5.1(c)); we then identify the (2,2) mode and sweep over a narrower range for further characterization. The mechanical features in OMIT and EMIT generally have Fano lineshapes, with linewidth Γ_T . We can also obtain Γ_T by transducing a swept signal to measure the mechanical mode's Lorentzian transmission profile (see Sec. 4.1.2), using SSB modulation of the microwave pump and optical heterodyne detection to measure $\mathcal{S}_{oe}(\omega)$ or SSB modulation of the optical pump and microwave heterodyne detection to measure $\mathcal{S}_{eo}(\omega)$.

The membrane mode quality factor $Q_m = \omega_m/\gamma_m$ of the (2,2) mode is sufficiently high that the frequency-domain characterization described above becomes unwieldy, so we interrogate the mechanical mode at low damping using a purely electromechanical time-domain measurement ($\Gamma_o = 0$). We apply a low-power microwave pump tone at fixed red detuning $\Delta_e = -\omega_m$ from ω_e , turn on a temporary SSB probe tone at detuning $+\omega_m$ from the pump to ring up the mechanical mode, then turn off the probe tone and demodulate the microwave heterodyne signal at frequency ω_m to monitor the exponential decay of the membrane oscillations at rate $\Gamma_T = \Gamma_e + \gamma_m$. By repeating this measurement at different microwave pump powers P_e we map out Γ_T vs. P_e (Fig. 5.1(d)). Extrapolating to $P_e = 0$ then yields the bare mechanical linewidth $\gamma_m/2\pi = 108$ mHz, corresponding to a quality factor $Q_m = 1.3 \times 10^7$. Performing these ringdown measurements electromechanically rather than optomechanically allows us to avoid optomechanical damping from the

lock beam (See Sec. 5.6.4). Measurements of the temperature-dependence of γ_m , discussed further in Sec. 5.7.2, indicate that $\gamma_m/2\pi = 113$ mHz is actually the relevant value for the data presented in the main text because the membrane equilibrates to an elevated temperature in the presence of laser light.

We use this same low-power measurement of Γ_T vs. P_e and Eq. (4.5) to infer the vacuum electromechanical coupling g_e (the electromechanical damping Γ_e does not scale linearly with P_e at high power because of power-dependent microwave circuit loss). The slope of the linear fit in Fig. 5.1(d) yields $g_e = 1.6 \pm 0.2$ Hz, with the uncertainty dominated by uncertainty in the cryostat insertion loss. Together with the zero-point motion obtained from simulation (see Sec. ??), this measurement implies that the electromechanical frequency shift per unit membrane displacement is $G_e = 3.2$ Hz/fm. This quantity is given by the relation

$$G_e = p \frac{\omega_e}{2d}, \quad (5.6)$$

where p is the ratio of the circuit's motionally modulated capacitance to its total capacitance, which decreases with increasing pad separation d . Solving Eq. 5.6 self-consistently for d and $p(d)$ then yields $d = 830$ nm and $p = 0.67$. An analogous linear fit to Γ_T vs. P_o at $\Delta_o = -\omega_m$ in a purely optomechanical measurement yields $g_o/2\pi = 60 \pm 6$ Hz, with uncertainty again due to insertion loss. The corresponding frequency shift $G_o/2\pi = 70 \pm 10$ Hz/fm is consistent with operation on the lower local maximum of optomechanical coupling in the top panel of Fig. 4.9(b), and the associated value of $\kappa_{o,\text{ext}}$ from the simulation (bottom panel) matches the measured value $\kappa_{o,\text{ext}}/2\pi = 2.12$ MHz.

The capacitor pad separation inferred from the electromechanical damping sweep is substantially larger than those previously observed in similar flip-chip devices. The device used in Ref. [Higginbotham et al. 2018] had $d = 300$ nm, and purely electromechanical devices tested since then have regularly achieved gaps in the 150 – 200 nm range [Burns 2019]. The large gap in this device was likely a consequence of modifications to the fabrication procedure required to integrate the phononic shield into the flip-chip architecture. In past devices without phononic shielding, the top chip was supported by two inner posts near the membrane as well as four outer posts, but these

inner posts reduced the effectiveness of the phononic shield in simulations.

We are presently exploring the effects of aligning inner posts with nodal lines of the membrane mode's transverse displacement profile to mitigate their impact on phononic shield performance. Another cause of unreliable pad separation is the thermal contraction of the epoxy used to affix the top chip to the middle chip, and we are investigating the Si-SiO₂ bonding technique used to affix the bottom and middle chips as an alternative to epoxy in future devices. Achieving $d = 200$ nm would boost g_e by a factor of 5, corresponding to an enhancement of the damping per incident microwave photon by a factor of 25 at low pump power (the improvement would be larger at higher power because the internal loss of the microwave circuit would be smaller at the power required to yield a given damping). Without any other device improvements, this increased electromechanical coupling would yield a membrane mode occupancy of $n_m = 0.4$ phonons and upconversion added noise $N_{\text{add,up}} = 0.9$ photons/s/Hz for $\Gamma_e = \Gamma_o = 2\pi \times 1$ kHz.

5.6.4 Characterization of other parameters

For all the measurements described above, the microwave pump detuning Δ_e is set to the optimal value $-\omega_m$ by stepping up the pump frequency until it begins to ring up the membrane mode for $\Delta_e > 0$, and then stepping the pump back down by ω_m . We use a similar though slightly more involved procedure to set the pump and lock beam detunings Δ_o and Δ_{lock} to the desired values. First, with the pump beam blocked, we adjust the PDH error signal bias while monitoring ringing of the lock to set $\Delta_{\text{lock}} = 0$. We then unblock the pump beam and adjust the AOM frequency that controls ω_{pump} (see Sec. 5.5) while monitoring the ringing of the lock to set $\Delta_o = -\omega_m$. In doing so we measure the cavity birefringence to be $\Delta_B/2\pi = 2.4$ MHz. Finally, we adjust the PDH error signal bias to slightly red-detune the lock beam. The lock beam detuning can be set to ~ 10 kHz precision by measuring how far ω_{pump} must be adjusted to maintain $\Delta_o = -\omega_m$. The lock beam optomechanical damping γ_{lock} is determined by comparing the mechanical mode linewidth obtained from an \mathcal{S}_{oe} measurement at low damping and $\Delta_{\text{lock}} = 0$ to the value obtained from a measurement with the same Γ_e and Γ_o but $\Delta_{\text{lock}} \neq 0$. The lock beam damping was $\gamma_{\text{lock}}/2\pi = 5$ Hz

during the measurements shown in Fig. 5.2 and $\gamma_{\text{lock}}/2\pi = 2$ Hz during the measurements shown in Fig. 5.3.

The values of γ_{lock} and Δ_{lock} are important because any contribution to the optomechanical or electromechanical damping of the membrane mode also couples the mode to a backaction bath as described in Sec. 4.1.3 above. This coupling cannot be neglected despite the low damping γ_{lock} because the small detuning Δ_{lock} leads to a large backaction bath occupancy $n_{\text{min,lock}}$. Indeed, for $\Delta_{\text{lock}} \ll \omega_{\text{m}}$, the product $\gamma_{\text{lock}}n_{\text{min,lock}}$ is independent of detuning, and can only be reduced by reducing the incident lock power P_{lock} . To fit the data in Figs. 5.2 and 5.3 we thus modify Eqs. (5.1) and (5.3) to include a $\gamma_{\text{lock}}n_{\text{min,lock}}$ term in the numerator, fixed at $\gamma_{\text{lock}}n_{\text{min,lock}} = 2\pi \times 40$ Hz by the independent measurements described above. This term results in the increase in the backaction limit at low damping in Fig. 5.2.

With optimally detuned pumps, the measured optical cavity linewidth κ_{o} above yields $\mathcal{A}_{\text{o}} = 1.22$ for the optomechanical transducer gain (see Sec. 4.1.2), and a backaction-limited membrane mode occupancy $n_{\text{min,o}} = 0.22$ for optomechanical ground-state cooling. Likewise, power-dependence of the microwave mode linewidth implies an electromechanical transducer gain \mathcal{A}_{e} that varies between 1.08 and 1.16 over the range of microwave pump power P_{e} used in this work. The net transducer gain $\mathcal{A}_{\text{e}}\mathcal{A}_{\text{o}}$ thus varies between 1.32 and 1.42.

For the data shown in Fig. 5.2, with fixed electromechanical damping $\Gamma_{\text{e}}/2\pi = 100$ Hz, the microwave mode linewidth was $\kappa_{\text{e}}/2\pi = 1.79$ MHz, yielding $\eta_{\text{M,exp}} = 0.55$ for the expected value of the matched efficiency. For the data shown in Fig. 5.3(a), η_{M} was not constant as a result of the power-dependent microwave loss. In these measurements the transducer efficiency η_{t} was maximized at $\Gamma_{\text{e}}/2\pi = 75$ Hz, where $\kappa_{\text{e}}/2\pi = 1.75$ MHz. The expected maximum efficiency $\eta_{\text{t,max}} = 49\%$ comes from evaluating Eq. (5.2) at $\Gamma_{\text{e}}/2\pi = 75$ Hz and $\Gamma_{\text{o}}/2\pi = 85$ Hz using this value of the microwave circuit linewidth and other independently measured parameters.

5.7 Analysis Details

5.7.1 Transducer efficiency

We measure the transducer efficiency η_t using the method developed in Ref. [Andrews et al. 2014], wherein measurements of upconversion, downconversion, and off-resonance reflection from both microwave and optical resonators can be used to calibrate out unknown path losses and gains. More precisely, a set of network analyzer measurements at fixed Γ_e and Γ_o yield

$$\sqrt{\frac{\left(\alpha |\mathcal{S}_{oe}(\omega_m)|^2 \delta\right) \left(\gamma |\mathcal{S}_{eo}(\omega_m)|^2 \beta\right)}{\left(\alpha |\mathcal{S}_{ee,off}|^2 \beta\right) \left(\epsilon_{PL} \gamma |\mathcal{S}_{oo,off}|^2 \delta\right)}} = \frac{\mathcal{A}_e \mathcal{A}_o}{\sqrt{\epsilon_{PL}}} \eta_t, \quad (5.7)$$

where α is the insertion loss of the microwave path from the HF2LI lock-in amplifier output to the external port of the electromechanical circuit, β is the gain of the microwave path from the electromechanical circuit output to the HF2LI input, γ is the insertion loss of the RF and optical path from the HF2LI output to the external port of the optical cavity, δ is the efficiency of the optical path from the optical cavity output to the HF2LI input (see Fig. 5.5). With path losses and gains defined this way, the “e” and “o” ports coincide with the external ports of the microwave and optical resonators, so $|\mathcal{S}_{ee,off}|^2 = |\mathcal{S}_{oo,off}|^2 = 1$ by construction. In our heterodyne detection scheme, a factor of the pump beam/LO modematching ϵ_{PL} appears in the measurement of the prompt reflection off the optical cavity (see Fig. 5.9, as also noted in Ref. [Higginbotham et al. 2018], and must be calibrated out with an independent measurement (see Sec. 5.6.2).

During the acquisition of the data shown in Fig 5.2(c), a mixer in the RF chain used for single sideband modulation of the optical pump (Fig. 5.5) was underdriven, which may have resulted in variation of γ between measurements of $|\mathcal{S}_{eo}|^2$ and $|\mathcal{S}_{oo,off}|^2$, and thus a miscalibration of the efficiency at each data point. We suspect this was the origin of the discrepancy between the expected matched efficiency $\eta_{M,exp}$ and the value $\eta_{M,fit}$ obtained from fitting the data in Fig. 5.2(c). The alternative explanation is miscalibration of the parameters that determine $\eta_{M,exp}$, but this seems unlikely in light of the good agreement between the measured and expected values of $\eta_{t,max}$ obtained from the data in Fig. 5.3(a).




Figure 5.9: **Losses and modematchings in efficiency calibration** Measurements in upconversion, downconversion and prompt reflection are used to calibrate out the unknown loss factors α , β , γ , and δ , according to Eq. 5.7 [Andrews et al. 2014]. The pump-LO modematching ϵ_{PC} then needs to independently measured to obtain the transducer efficiency.

The fit to efficiency vs. damping must be modified when the electromechanical damping Γ_e is swept, as in Fig. 5.3(a), because of the power-dependent microwave circuit loss. To process this data we define $\zeta_t = \eta_t/(\kappa_{e,\text{ext}}/\kappa_e)$ and fit ζ_t vs. damping to Eq. (5.2), then compare $\zeta_M = 0.59$ obtained from the fit to the expected value $\zeta_M = \epsilon\kappa_{o,\text{ext}}/\kappa_o = 0.63$. The theory curve shown in Fig. 5.3(a) is obtained by multiplying the fit to ζ_t vs. damping data by a polynomial fit to the power-dependent microwave mode overcoupling ratio $\kappa_{e,\text{ext}}/\kappa_e$. Because of the power-dependent loss, the efficiency is maximized at $\Gamma_e = 2\pi \times 75$ Hz rather than at $\Gamma_e = \Gamma_o = 2\pi \times 85$ Hz.

5.7.2 Calibration of spectra and temperature sweeps

The measured optical heterodyne noise spectra are normalized to the sum of LO beam shot noise and heterodyne detector dark noise, obtained by blocking the pump beam incident on the heterodyne detector. These spectra are related to the transducer output-referred spectra (Eqs. (4.17) and (4.18)) by $S_{\text{det},o,\pm}(\omega) = (1 - \xi_o) + \xi_o S_{o,\pm}(\omega)$, where ξ_o is the efficiency of the optical measurement chain, modeled as an effective beamsplitter between the transducer output and the input of an ideal heterodyne detector. Measured microwave heterodyne noise spectra are likewise normalized to the noise measured in the absence of the pump tone, given by the sum of vacuum noise, residual thermal noise at the base plate temperature $T_{\text{bp}} = 40$ mK, and the added noise of the microwave measurement chain. As in the optomechanical case, the spectra normalized this way are related to




Figure 5.10: **Temperature sweeps.** (a) Intrinsic mechanical dissipation rate γ_m vs. cryostat base plate temperature T_{bp} , with linear fit. (b) Area under mechanical peak in optical output spectrum vs. T_{bp} . The quadratic fit (gray) calibrates the optical measurement chain efficiency. Blue squares indicate data excluded from the fit, as the membrane mode thermalizes poorly to the base plate at low temperatures in the presence of the lock beam. (c) Area under mechanical peak in microwave output spectrum vs. T_{bp} in the absence of optical illumination. All error bars represent one standard deviation.

spectra normalized at the transducer's microwave output port by $S_{\text{det,e},\pm}(\omega) = (1 - \xi_e) + \xi_e S_{e,\pm}(\omega)$, where ξ_e is the efficiency of the microwave measurement chain. The measurement chain efficiencies ξ_e and ξ_o appearing in these expressions can equivalently be expressed in terms of the added noise of the respective measurement chains referred to the transducer output, via $\xi_e = 1/(N_{\xi,e} + 1/2)$ and $\xi_o = 1/(N_{\xi,o} + 1/2)$.

As a cross-check on Stokes/anti-Stokes sideband ratio thermometry using Eq. (4.12), we can use Eq. (4.10) to infer the membrane mode occupancy n_m from the upper sideband spectrum $S_{\text{det,o},+}(\omega)$ given independent measurements of ξ_o and parameters in Table 5.1. To determine ξ_o , we sweep the cryostat base temperature T_{bp} while measuring the optical noise spectrum, at low damping such that technical noise is negligible. At sufficiently high temperature, the local thermal environment of the membrane will equilibrate to T_{bp} , and thus we can replace the base-temperature equilibrium occupancy n_{th} in Eq. (5.1) with $n_{bp} = k_B T_{bp}/\hbar\omega_m$. Then, if other parameters are independent of temperature, the measurement chain efficiency ξ_o can be obtained from the slope of $N_{\text{det,o},+}(T_{bp})$, where $N_{\text{det,o},+}$ is the optomechanical upper sideband amplitude normalized at the detector input ($N_{\text{det,o},+} = \xi_o N_{\text{out,o},+}$ in the absence of technical noise). As T_{bp} decreases, the membrane mode may decouple from the fridge base plate and equilibrate instead to a thermal

bath at some elevated temperature $T_{\text{eq}} > T_{\text{bp}}$. The temperature sweep enables us to identify T_{eq} as the temperature below which $N_{\text{det,o,+}}(T_{\text{bp}})$ deviates from the expected linear behavior, and the corresponding thermal occupancy $n_{\text{th}} = k_{\text{B}}T_{\text{eq}}/\hbar\omega_{\text{m}}$ can be compared to the values obtained from fitting the data in Fig. 5.2(a) and Fig. 5.3(b).

While sweeping the base plate temperature T_{bp} , we repeated sets of electromechanical ring-down measurements of the sort shown in Fig. 5.1(d) and observed that the mechanical dissipation rate γ_{m} increased linearly with temperature (Fig. 5.10(a)). Fitting to this data to $\gamma_{\text{m}} = a_{\gamma}T_{\text{bp}} + b_{\gamma}$ yields $a_{\gamma}/2\pi = 176$ mHz/K and $b_{\gamma}/2\pi = 101$ mHz. This behavior will generate quadratic terms in $N_{\text{det,o,+}}(T_{\text{bp}})$ and $N_{\text{det,e,+}}(T_{\text{bp}})$, as the membrane mode's coupling to the temperature-dependent bath is itself temperature-dependent.

The optomechanical temperature sweep data is shown in Fig. 5.10(b). We obtain peak amplitudes $N_{\text{det,o,+}}$ and peak widths Γ_{T} from Lorentzian fits to the spectra at different temperatures, and fit the temperature-dependence of the peak area rather than the peak amplitude to control for fluctuations in damping. We parameterize the temperature-dependence as $N_{\text{det,o,+}}\Gamma_{\text{T}} = a_{\xi}(a_{\gamma}T_{\text{bp}}^2 + b_{\gamma}T_{\text{bp}}) + b_{\xi}$ with a_{γ} and b_{γ} fixed by the fit to the data in Fig. 5.10(a), where the theoretical values of the coefficients are

$$a_{\xi} = 4\xi_{\text{o}}\epsilon_{\text{CL}}\frac{\kappa_{\text{o,ext}}}{\kappa_{\text{o}}}\mathcal{A}_{\text{o}}\frac{\Gamma_{\text{o}}}{\Gamma_{\text{T}}}\frac{k_{\text{B}}}{\hbar\omega_{\text{m}}} \quad (5.8)$$

and

$$b_{\xi} = 4\xi_{\text{o}}\epsilon_{\text{CL}}\frac{\kappa_{\text{o,ext}}}{\kappa_{\text{o}}}\mathcal{A}_{\text{o}}\frac{\Gamma_{\text{o}}}{\Gamma_{\text{T}}}(\gamma_{\text{lock}}n_{\text{min,lock}} + \Gamma_{\text{o}}n_{\text{min,o}}). \quad (5.9)$$

At low temperature, the data begins to deviate from quadratic temperature-dependence, indicating an elevated membrane mode equilibration temperature. We obtain the best fit when excluding the three lowest-temperature data points, marked in blue in Fig. 5.10(b). Then, since all other parameters are independently measured, the best-fit value of a_{ξ} yields $\xi_{\text{o}} = 0.276 \pm 0.007$, while b_{ξ} , whose fractional uncertainty is much larger, yields $\xi_{\text{o}} = 0.27 \pm 0.07$. The mode equilibration temperature T_{eq} , defined as the temperature at which the area predicted by the fit is equal to the measured base-temperature peak area, is inferred to be $T_{\text{eq}} = 70 \pm 10$ mK, with uncertainty again

dominated by the fractional error in b_ξ .

Finally, we explore how the values of the measurement chain efficiency ξ_o and the equilibration temperature T_{eq} depend on the decision to exclude low-temperature points from the fit. Excluding three points yields the best agreement between the values of ξ_o inferred from Eqs. (5.8) and (5.9). When only the lowest point is excluded, Eq. (5.9) yields $\xi_o = 0.37 \pm 0.05$, two standard deviations away from the value obtained from a_ξ , which is very insensitive to the low-temperature data. Excluding two or four points does not change the fit very much, and the inferred equilibration temperature remains within the range quoted above. We thus take $\xi_o = 0.276 \pm 0.007$ and $T_{\text{eq}} = 70 \pm 10$ mK as our final values for the measurement chain efficiency and equilibration temperature. The latter corresponds to $n_{\text{th}} = 1000 \pm 140$, which agrees well with the values obtained from independent fits to the data in Figs. 5.2(a) and 5.3(b).

We follow an analogous procedure to measure the microwave measurement chain efficiency ξ_e via the temperature-dependence of the electromechanical upper sideband amplitude $N_{\text{det,e,+}}$. The analysis of the electromechanical temperature sweep data, shown in Fig. 5.10(c), is very similar to the optomechanical analysis and will not be discussed in detail. The fit yields $\xi_e = 0.029 \pm 0.001$, with no evidence for equilibration at an elevated temperature. We use this value of ξ_e to calibrate broadband measurements of the noise emerging from the external port of the microwave circuit, with which we study the dependence of $n_{\text{eff,e}}$ on microwave and optical power (Figs. 5.12(b) and 5.4(b), respectively).

The electromechanical temperature sweep data indicates that laser light is responsible for the elevated equilibration temperature inferred from Figs. 5.10(b), 5.2(a), and 5.3(b). As noted in the main text, the fact that the optomechanical ground-state cooling data plotted in Fig. 5.2(a) agrees well with Eq. (5.1) without additional power-dependent terms indicates that membrane mode heating saturates at the low power used to lock the cavity, with no further scaling with the power of the pump beam. Similar low-power saturation was observed in membrane optomechanical systems in Refs. [Higginbotham et al. 2018] and [Peterson et al. 2016]. The elevated T_{eq} implies that the appropriate value of the intrinsic mechanical dissipation rate γ_m for analyzing the data in the

main text is the 70 mK value $\gamma_m/2\pi = 113$ mHz rather than the 40 mK value $\gamma_m/2\pi = 108$ mHz.

We can compare the temperature sweep calibration of the optical measurement chain efficiency ξ_o to a more direct but less precise measurement. From power meter measurements we infer the efficiency of the optical path between the optical cavity output and the heterodyne detector input to be $\xi_{\text{path}} = 0.4$. The fractional contribution of LO beam shot noise (at power $P_{\text{LO}} = 1.3$ mW) to the sum of shot noise and heterodyne detector dark noise was measured to be $\xi_{\text{dark}} = 0.79$. The nominal photodetector quantum efficiency of $\sigma_q = 0.87$ then yields $\xi_o = \xi_{\text{path}} \xi_{\text{dark}} \sigma_q = 0.27$, consistent with the temperature sweep measurements.

On the microwave side, we measured the added noise of the microwave chain referred to the input of the HEMT amplifier at the 4 K stage of the cryostat by sweeping T_{bp} and measuring the thermal noise of the effective $50\,\Omega$ load seen by the HEMT. From this measurement we obtained $N_{\text{HEMT}} = 8.5$ photons/s/Hz, equivalent to an efficiency of $1/(N_{\text{HEMT}} + 1/2) = 0.11$ referred to the HEMT input. The discrepancy between this number and the measurement chain efficiency ξ_e referred to the transducer's microwave output port implies a 5.8 dB loss between the transducer's microwave output port and the HEMT input. This is a rather large value, but the value inferred from measurements in Ref. [Higginbotham et al. 2018] using the same cryostat was only slightly smaller, and more microwave connectors were present in the signal path here than in that earlier work.

5.7.3 Ground-state cooling measurements

To process the data shown in Figs. 5.2(a) and (b), we first normalize the optomechanical upper- and lower-sideband spectra as described in Sec. 5.7.2 above. Two such normalized spectra from the pure optomechanical data set are shown in Fig. 5.11(a). We observe that the white noise background around both sidebands increases linearly with optical pump power, as expected in the presence of amplitude and phase noise (see Sec. 4.1.5). We also observe a feature at frequency $\omega_s/2\pi = 1.448$ MHz emerging from the noise floor around the upper sideband at high power. We attribute this feature to a relatively massive vibrational mode of the silicon chip substrate, and

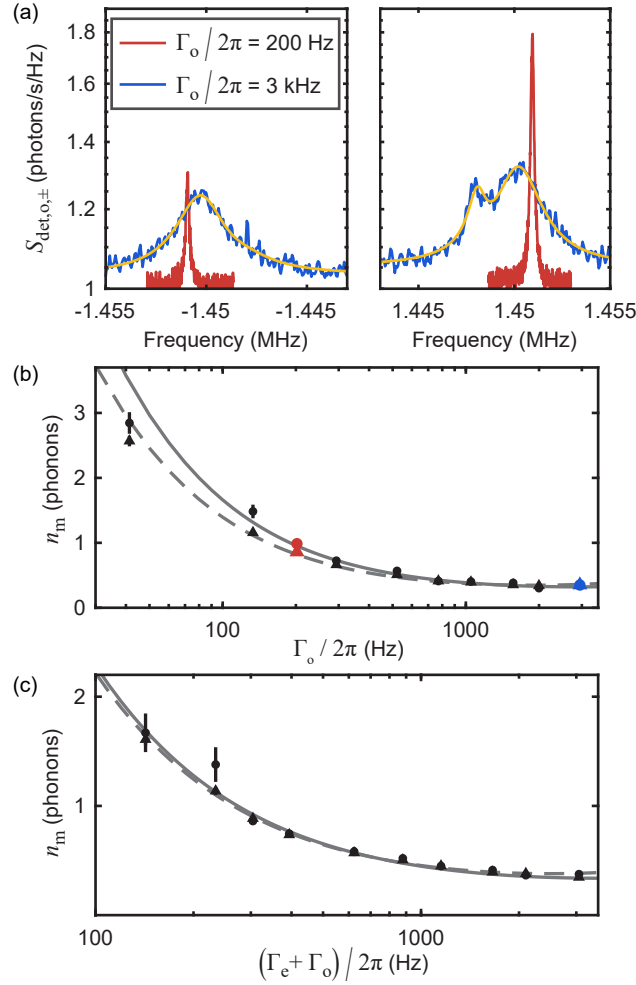


Figure 5.11: **Analysis of optomechanical spectra.** (a) Two examples of optomechanical spectra normalized to shot noise at the optical heterodyne detector input. The left (right) panel displays the lower (upper) sideband spectra. At high damping, we fit the spectra to the squared sum of two complex Lorentzian amplitudes (yellow line) to account for the 1.448 MHz mode of the silicon substrate. (b) Comparing two methods of inferring membrane mode occupancy n_m from purely optomechanical damping sweep, with corresponding fits. Circles and solid line (also shown in Fig. 5.2(a)) are obtained using sideband asymmetry thermometry, while triangles and dashed line are obtained from the upper-sideband spectra $S_{\text{det},o,+}$ and optical measurement chain efficiency ξ_o . Red and blue data points correspond to the spectra shown in panel (a). (c) Comparing the two methods of inferring n_m for optomechanical damping sweep at fixed electromechanical damping (sideband asymmetry data and fit also shown in Fig. 5.2(b)). All error bars represent one standard deviation.

account for it by fitting spectra to the squared sum of two complex Lorentzians with the amplitude of the second mode and the relative phase as additional fit parameters. This fit was motivated by observations of destructive interference between a membrane mode and a substrate mode in

unpublished data from a previous device. For the present device, this coherent double-Lorentzian fit differs by less than 10% from a fit to the incoherent sum of two squared Lorentzian amplitudes. The fits are not appreciably altered by including antisymmetric terms as in Eqs. (4.17) and (4.18).

The upper and lower sideband amplitudes $N_{\text{det,o},\pm}$ obtained from these fits must then be corrected to undo the effects of amplitude and phase noise (Sec. 4.1.5) before we can infer the membrane mode occupancy. Correcting the peak amplitudes in this way in turn requires knowledge of the relative magnitudes of the shot noise-normalized technical noise spectral densities C_{xx} , C_{yy} , and C_{xy} , which are related to the spectral densities of fractional amplitude and phase fluctuations by $C_{xx} = 4\dot{N}_o (S_{\delta A, \delta A}(\omega_m)/2)$, $C_{yy} = 4\dot{N}_o (S_{\delta \phi, \delta \phi}(\omega_m)/2)$, and $C_{xy} = 4\dot{N}_o (S_{\delta A, \delta \phi}(\omega_m)/2)$, where $\dot{N}_o = P_o/\hbar\omega_{p,o}$ is the photon flux of the pump incident on the optical cavity. Here $S_{\delta A, \delta A}(\omega)$, $S_{\delta \phi, \delta \phi}(\omega)$, and $S_{\delta A, \delta \phi}(\omega)$ are defined as one-sided double-sideband spectral densities in units of rad^2/Hz . The factor of 1/2 in each expression accounts for the transition from the two-sided spectral densities assumed by the theoretical framework in Sec. 4.1.5 to the one-sided convention more commonly used by experimentalists.

We made independent measurements of $S_{\delta A, \delta A}(\omega)$ and $S_{\delta \phi, \delta \phi}(\omega)$ shortly before the cooldown in which we acquired the data presented in this work, using a very similar optomechanical cavity in the same setup as an AM/ Φ M transducer. We measured amplitude noise by comparing the sum and difference photocurrents in balanced direct detection of the pump beam promptly reflected off the cavity far from resonance [Yu 2015]. We then repeated this measurement with the laser locked to the cavity and the pump beam near resonance to transduce amplitude noise to phase noise, and calibrated the cavity's AM/ Φ M transduction coefficient using a phase modulation sideband of known modulation depth [Safavi-Naeini et al. 2013]. From these measurements we obtained $10 \log_{10} (S_{\delta A, \delta A}(\omega)) \approx -155 \text{ dBc/Hz}$ and $10 \log_{10} (S_{\delta \phi, \delta \phi}(\omega)) \approx -136 \text{ dBc/Hz}$ at frequencies near ω_m . It is worth noting that we made these measurements at room temperature, and highly-occupied mirror substrate vibrational modes (see Sec. 6.2.5) added significant noise to the measurement. The level quoted above was the noise floor we observed between the thermomechanical peaks, though the vibrational noise could have limited that floor as well. However, it seems unlikely that our noise

level was much better at the time of this measurement, because after making the improvements to the filter cavity PDH lock loop described in Sec. 6.2.2, we measured -139 dBc/Hz using another method, also described in Sec. 6.2.2. Using these numbers in Eq. (4.16) at the highest-power data point in Fig. 5.2, we reproduce the measured values of the off-resonance excess noise $\tilde{S}_{o,\pm}$ to within a factor of 1.5 if we assume maximal positive correlations $C_{xy} = \sqrt{C_{xx}C_{yy}}$, lending credence to the conclusion that $C_{yy} \gg C_{xx}$, C_{xy} . The results obtained from the subsequent analysis assuming phase noise only differ negligibly from the results obtained assuming the measured ratio between $S_{\delta A, \delta A}(\omega_m)$ and $S_{\delta \phi, \delta \phi}(\omega_m)$ and maximal positive correlations, so we discuss the former analysis here for simplicity.

Here, I'd like to note that the above analysis to extract the phase noise is based on an assumption that is invalid for our measurement. The measurement was performed by detuning the pump by $\Delta = -\kappa/2$ from the cavity resonance. At this detuning, slow frequency fluctuations of the beam will be transduced to amplitude modulation, and, because the cavity response is steepest at this detuning, the cavity's AM/ Φ M transduction coefficient is maximized. At the same time, slow amplitude fluctuations would be unaffected, so subtracting the amplitude noise measured in direct detection with the cavity far detuned would give the desired phase noise measurement. However, this model is applies to changes in amplitude and frequency that are adiabatic, such that the light field in the cavity reaches equilibrium. The noise spectrum we are interested in at frequencies around ω_m , on the other hand, is of comparable frequency to the cavity decay rate. Instead, the cavity would cause the amplitude and phase quadratures of the light to rotate into one another, which in practice could change the result of the analysis.

For $C_{yy} \gg C_{xx}$, C_{xy} , the anti-symmetric contribution to both the Stokes and the anti-Stokes peak lineshapes is suppressed and the Lorentzian contribution is negative for both peaks, resulting in squashing. We undo this phase noise squashing by taking

$$N_{\text{det},o,\pm} \rightarrow N_{\text{det},o,\pm} \pm \frac{\partial \left(\epsilon_{\text{CL}} \epsilon_{\text{PC}} \frac{\kappa_{o,\text{ext}}}{\kappa_o} \kappa_o^2 \mathcal{A}_o \frac{\Gamma_o}{\Gamma_T} \text{Re}[\tilde{B}_{\pm}] \right)}{\partial C_{yy}} \left(\frac{\partial \tilde{S}_{o,\pm}}{\partial C_{yy}} \right)^{-1} \tilde{S}_{\text{det},o,\pm}, \quad (5.10)$$

where the two partial derivatives can be evaluated using independently measured transducer param-

eters, and $\tilde{S}_{\text{det},\text{o},\pm}$ is the measured excess over the white noise background level of 1 photon/Hz/s. This method exploits the fact that the phase noise is responsible for both the squashing of the peak amplitudes and the increase in the white noise background to correct for squashing without assuming a specific value of C_{yy} . We then use the corrected peak amplitudes to infer the membrane mode occupancy n_{m} via Eq. (4.12), and fit the data as shown in Figs. 5.2(a) and (b), with an additional $\gamma_{\text{lock}}n_{\text{min,lock}}$ term in Eq. (5.1) to account for the lock backaction effects discussed in Sec. 5.6.4.

As a final check on the consistency of the sideband asymmetry analysis, we can use the value of the coefficient a_{o} obtained from the fit shown in Fig. 5.2(a) to get an independent estimate of the phase noise spectral density. Setting the effective optical mode occupancy $n_{\text{eff,o}} = a_{\text{o}}\Gamma_{\text{o}}$ in Eq. (4.15) and dropping the C_{xx} and C_{xy} terms, we obtain $10 \log_{10} (S_{\delta\phi,\delta\phi}(\omega_{\text{m}})) = -135^{+2}_{-3}$ dBc/Hz, consistent with the independent measurement.

As discussed in Sec. 5.7.2, we can also obtain the membrane mode occupancy n_{m} from the normalized upper-sideband spectrum $S_{\text{det},\text{o},+}$, given the optical measurement chain efficiency ξ_{o} and independent measurements of other transducer parameters. The n_{m} values obtained this way and the corresponding fit are overlaid on the values obtained from the sideband asymmetry analysis in Fig. 5.11(b) and (c) for the purely optomechanical and electro-optomechanical damping sweeps respectively. Fitting to the purely optomechanical data to Eq. (5.1) yields $n_{\text{th}} = 750 \pm 50$ and $a_{\text{o}} = 5 \pm 1$, and fixing these values in the fit to the electro-optomechanical data yields $n_{\text{eff,e}} = 1.02 \pm 0.07$ for the effective microwave mode thermal occupancy. The origin of the discrepancy between these values and those obtained from the sideband asymmetry analysis is unknown, but the analysis using only the upper sideband data is more susceptible to miscalibration of parameters. The equilibrium occupancy n_{th} inferred from this fits is very sensitive to the first data point, and the discrepancy between the values of $n_{\text{eff,e}}$ obtained from the two electro-optomechanical fits is strongly correlated with the discrepancy between the values of n_{th} .

5.7.4 Added noise measurements

The values of $N_{\text{add,up}}$ plotted in Fig. 5.2(d) are obtained from the inferred phonon occupancy data in Fig. 5.2(b) via Eq. (4.13), with an additional term accounting for the off-resonance level $\tilde{S}_{o,\pm}$ of each spectrum. This off-resonance level (normalized at the transducer output) can be obtained from the directly measured spectra without need for independent measurement of the optical measurement chain efficiency ξ_o , because a measurement of n_m calibrates the optical output spectra.

The analysis of the data shown in Fig. 5.3(b) is simpler than the ground-state cooling analysis discussed above, as both optical pump phase noise and interference from the substrate mode shown in Fig. 5.11(a) were negligible with the optomechanical damping fixed at $\Gamma_o/2\pi = 85$ Hz. However, the lower sideband amplitude $N_{\text{det,o,-}}$ was a very small fractional excess over the white noise background at high microwave pump power, where $\Gamma_o \ll \Gamma_e$. We thus calibrate the added noise using the optical measurement chain efficiency ξ_o obtained from the temperature sweep rather than sideband asymmetry. We fit the upper sideband amplitude $N_{\text{out,o,+}} = N_{\text{det,o,+}}/\xi_o$ as a function of damping to the expected behavior obtained by substituting Eq. (5.1) into Eq. (4.10) and evaluating at frequency ω_m :

$$N_{\text{out,o,+}} = 4\epsilon_{\text{CL}}\mathcal{A}_o \frac{\kappa_{o,\text{ext}}}{\kappa_o} \Gamma_o \frac{\gamma_m n_{\text{th}} + \Gamma_e n_e + \Gamma_o n_o}{\Gamma_T^2}. \quad (5.11)$$

As in the analysis of the data shown in Fig. 5.2, we also include a contribution from lock beam backaction in the fit. To obtain the data and the theory curve shown in Fig. 5.3(b), we divide by the product of efficiency η_t and transducer gain $\mathcal{A}_e\mathcal{A}_o$. Calibrating the data with the sideband asymmetry analysis instead yields slightly lower added noise, so our choice to present the data calibrated using ξ_o is also conservative.

At the optimal electromechanical damping rate $\Gamma_e/2\pi = 135$ Hz where added noise is minimized, the total added noise arises from the weighted sum of couplings to the various baths that determine n_m : 1.4 photons/s/Hz from $n_{\text{eff,e}}$, 1.0 photons/s/Hz from n_{th} , 0.4 photons/s/Hz from $n_{\text{min,lock}}$, 0.2 photons/s/Hz from $n_{\text{min,o}}$, and 0.1 photons/s/Hz from $n_{\text{min,e}}$. The final 0.1 pho-

tons/s/Hz comes from excess white noise referred to the transducer's microwave input.




Figure 5.12: **Microwave pump effects on superconducting circuit.** (a) Internal loss of the microwave circuit $\kappa_{e,int}$ vs. $\bar{n}_{p,e}$, the resonator coherent state mean photon number. (b) Effective microwave mode occupancy $n_{eff,e}$ vs. electromechanical damping Γ_e , with linear fit. All error bars represent one standard deviation.

5.8 Microwave pump effects on superconducting circuit

We observed two adverse effects of the microwave pump on the superconducting microwave circuit, whose origins are likely related. The power-dependence of the circuit's internal loss $\kappa_{e,int}$ is plotted in Fig. 5.12(a) as a function of $\bar{n}_{p,e}$, the mean photon number of the coherent state induced in the microwave mode by the microwave pump, over the range of power used in Fig. 5.3. As noted in Sec. 5.6, this power-dependence results in power-dependent electromechanical transducer gain \mathcal{A}_e , modifies the expected behavior of transducer efficiency η_t vs. electromechanical damping, and leads to sublinear scaling of damping vs. microwave pump power.

A more significant constraint on the transducer performance is the power-dependent mi-

microwave mode effective occupancy $n_{\text{eff},e}$. To calibrate $n_{\text{eff},e}$, we first refer the measured microwave noise spectral density $S_{\text{det},e}$ to the transducer output by dividing by the measurement chain efficiency ξ_e , then divide by $4\kappa_{e,\text{ext}}/\kappa_e$ to translate the spectral density into a mode occupancy. Operationally, we are interested in the behavior of $n_{\text{eff},e}$ as a function of electromechanical damping Γ_e , plotted in Fig. 5.12(b). The noise is seen to scale as $n_{\text{eff},e} = a_e\Gamma_e + b_e$, with $a_e = 1.1 \times 10^{-3} \text{ Hz}^{-1}$ and $b_e = 0.077$, with 1% and 5% fractional uncertainty in the slope and offset, respectively. Evaluated at $\Gamma_e = 2\pi \times 100 \text{ Hz}$, as in the data shown in Fig. 5.2(b), the fit yields $n_{\text{eff},e} = 0.77 \pm 0.01$, consistent with the value obtained from the Fig. 5.2(b) fit. Fitting the data in Fig. 5.3(b), in which Γ_e is swept, yields $a_e = (1.17 \pm 0.05) \times 10^{-3} \text{ Hz}^{-1}$ and $b_e = 0.1 \pm 0.1$, consistent with the values inferred from direct measurements of the noise described above.

One possible contribution to $n_{\text{eff},e}$ is the phase noise of the microwave pump. We used a Rohde & Schwartz SMA100B signal generator to source the pump specifically for the ultra-low phase noise afforded by option SMAB-B711, and measured its DSB noise *in situ* to be $10 \log_{10}(S_{\delta\phi,\delta\phi,e}(\omega)) = -149 \text{ dBc/Hz}$ at frequencies near ω_m . Following the procedure used to infer the phase noise spectral density of the optical pump from the coefficient a_o , we find that microwave pump phase noise can account for about 25% of the measured value of a_e . The remainder of the power-dependent noise must be attributed to the behavior the superconducting circuit itself.

We note that microwave power-dependent noise was also observed in the device described in Ref. [Higginbotham et al. 2018], and motivated our decision to fabricate the superconducting circuit from NbTiN rather than Nb as in past work. In test circuits without released membranes that were otherwise identical to the flip-chip circuits used in our transducer, we observed reduction of the power-dependent noise by a factor of 4 in circuits fabricated from NbTiN. This improvement was largely preserved in our integrated electro-optomechanical device, and the microwave power-dependent noise remained a chief limitation primarily because of the large capacitor pad separation d of the device used in this work. As noted in Sec. 5.6.3, reproducing the values of d obtained in previous devices would reduce the power required to obtain a given electromechanical damping by at least a factor of 25, which would greatly reduce the impact of power-dependent noise on

transducer performance even without further reduction of the noise itself.

We now believe that the power-dependent noise and internal loss of the microwave circuit have a common origin: the presence of hydrogen impurities in the silicon nitride membrane [Mittal et al. 2023]. In addition to contributing the commonly recognized saturable two-level-system (TLS) loss and noise at drive strengths much lower than those shown in Fig. 5.12, the hydrogen would also contribute dynamics described by a relaxation-TLS model at higher microwave drive powers. These dynamics include temperature-dependent dielectric loss, and non-saturable power-dependent noise. The power dependent loss can then be described by a self-heating model, in which the microwave pump increases the temperature of the silicon nitride in the membrane. These hydrogen impurities are removed by annealing the silicon nitride membrane before depositing the superconducting circuit. We expect this will greatly reduce the added noise and power-dependent loss of our next transducer devices, and expect they should demonstrate quantum transduction, with $N_{\text{add}} < 1$.

5.9 Dispersive readout of a transmon qubit using optical heterodyne detection via the transducer

In this section I'll relate the results of an experiment in which we optically read out the state of a superconducting transmon qubit by upconverting a microwave dispersive readout pulse, and observe that the transducer imparts very low backaction on the qubit [Delaney et al. 2022]. This work is covered in detail in the thesis of Rob Delaney, who designed the experiment and built the qubit-cavity system [Delaney 2022], so here I'll briefly summarize the results and how what we learned from this experiment informs design choices we make as we orient towards incorporating single-photon counting detection into our experiments as described in Ch. 6.

To test the transducer's compatibility with a superconducting qubit despite added noise preventing quantum transduction, we perform dispersive readout via the transducer [Blais et al. 2004]. A microwave readout pulse of approximately 400 photons incident on the qubit cavity acquires a phase shift dependent on the state of the qubit. We direct the pulse towards the transducer, and measure the phase of the output optical field using heterodyne detection (Fig. 5.13).




Figure 5.13: Setup to perform dispersive readout via electro-optic transduction and optical heterodyne detection. (a) A circuit QED system is prepared and dispersively read out with microwave pulses at frequencies ω_{prep} and ω_{read} respectively. A sapphire rod is used to tune the readout cavity into resonance with the electro-optic transducer. (b) Microwave isolation and pump cancellation protect the qubit from the transducer’s microwave pump, and provides a microwave readout channel, while routing the readout pulse towards the transducer. (c) The electro-optic transducer comprises a microwave-frequency LC circuit (blue) and optical cavity mode (red) coupled to the same vibrational mode of a mechanical oscillator (purple). The photo inset shows the electromechanical subsystem of the transducer before it is incorporated into an optical cavity. The silicon chip supporting the metalized silicon nitride membrane and phononic shielding is mounted above the chip supporting the remainder of the LC circuit. (d) The transduced readout pulse is measured with optical heterodyne detection. [Delaney et al. 2022]

To calibrate this readout scheme, we compile a histogram of optical heterodyne readout voltages for repeated measurements initializing the qubit in either its excited and ground states (Fig. 5.14.a). With this characterization, we set a voltage threshold of 0 V for single-shot readout, which we use to perform a Rabi experiment, observing Rabi oscillations with optical detection (Fig. 5.14.b).

We can assess the backaction [Hatridge et al. 2013] of the transducer on the qubit by performing Ramsey experiments to measure the dephasing time T_2 of the qubit, using the auxiliary microwave readout port shown in Fig. 5.13.b. We perform this experiment three times: with the transducer’s pumps off, with only the optical pump on, and with the transducer fully running with both microwave and optical pumps (Fig. 5.14.c). Comparing the extracted values of T_2 , notably, we observe no discernable degradation in the qubit’s dephasing rate with the application of the optical pump. We do observe a slight reduction in T_2 with applying the microwave pump, equivalent to an additional $\Delta n = (3 \pm 1) \times 10^{-3}$ photons occupying the microwave readout cavity.




Figure 5.14: **Dispersive readout measurement results.** (a) Histogram of measured optical heterodyne voltage with qubit prepared in ground (blue) and excited (red) states. Orange dashed line indicates the voltage threshold chosen for single-shot readout. (b) Rabi experiment using measurement threshold indicated in (a). (c) Ramsey experiments using microwave readout with the transducer off (red circles), with the transducer's optical pump on (cyan diamonds), and with the transducer's microwave and optical pumps on (blue squares). Extracted T_2 values show little dependence on the transducer pumps. The diminished visibility associated with the microwave pump is due to increasing microwave resonator loss in the transducer, causing decreased signal-to-noise in microwave readout.[Delaney et al. 2022]




Figure 5.15: **Using qubit readout to calibrate the total quantum efficiency.** The SNR of dispersive readout gives the quantum efficiency of the entire readout chain η_q (left). This model combines the effects of both loss and noise on the SNR, treating it as the effect of a single beamsplitter with efficiency η_q (upper center) This is equivalent to a model that treats the total path efficiency η separately from the additional noise detected at the detector N_{det} , which includes contributions from vacuum fluctuations and ideal heterodyne detection, such that $\eta_q = \eta / N_{\text{det}}$ (lower center). The independent effects of efficiency due to loss η and additional noise added by the chain $N_{\text{det}} - 1$ on the SNR are depicted on the right.

We can also use the qubit state readout as a method to calibrate the total quantum efficiency $\eta_q = \eta/N_{\text{det}}$, where η is the total efficiency from the qubit cavity to the optical detection and N_{det} is the total noise measured at the optical heterodyne detector, including all sources of transducer and technical noise, as well as vacuum noise and the added noise of an ideal heterodyne detector, as illustrated in Fig. 5.15. We measure a maximum $\eta_q = 8 \times 10^{-4}$, and of which the chain efficiency contributes $\eta = 1.9 \times 10^{-3}$, using independent measurement of N_{det} . We further break down the contributions to η_q , summarized in Table 5.2.

description of inefficiency	symbol/expression for efficiency	efficiency value
bandwidth mismatch	η_{bw}	0.15
microwave transmission loss	η_{MW}	0.34
optical detection efficiency	ξ_{o}	0.28
transducer LC circuit loss	$\kappa_{\text{e,ext}}/\kappa_{\text{e}}$	0.53
optical cavity loss	$\kappa_{\text{o,ext}}/\kappa_{\text{o}}$	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 added noise	$1/N_{\text{det}}$	0.42

Table 5.2: Contributions to η_q in qubit readout experiment.

This demonstration of operating a microwave-optical transducer without significantly impacting the performance of the superconducting is promising in that we can operate the transducer continuously with an interfaced qubit, and that the qubit will retain information as demanded by the networking protocols described in Ch. 3.

Chapter 6

Subsequent improvements towards a quantum-enabled transducer

In this chapter I describe the developments made after the demonstrations of ground state cooling and superconducting qubit readout detailed in Ch. 5. First, I describe operating our optomechanical measurement setup in a new, cryogen-free dilution refrigerator, a modification to the optical cavity in order to clamp the front mirror in place, and our success in replicating the ability to optomechanically ground-state cool the membrane mode. I then orient towards the tighter constraints that an experiment using single-photon detection places on laser noise, and consider different sources of noise. I detail improvements to the optical setup in order to reach these noise thresholds, and measurements of noise imprinted on the light due to thermal occupation of vibrational modes of the optical cavity mirrors themselves.

6.1 Optomechanics in a cryogen-free dilution refrigerator

All of our measurements up through Refs. [Brubaker et al. 2022] and [Delaney et al. 2022] were done in the low-vibration environment of a “wet fridge”—a dilution refrigerator that uses liquid helium to precool the dilution unit, shown in the photograph in Fig. 6.1(a). However, there are drawbacks associated with such systems: they are expensive to run without a highly efficient recovery system, helium supply is subject to shortage, initial cool-down consumes more helium and must be done manually, they require routine transfers to refill the helium tank, and there is a possibility to plug the helium input line when transferring, making very long experimental runs with the same device less likely. By switching to a “dry” cryogen-free dilution refrigerator that

uses a pulse-tube (PT) cooler in lieu of a liquid helium tank, we gain turn-key functionality that allows for rapid and frequent cycling of devices, and the ability to retain a given device at cryogenic temperatures for longer periods when needed, especially alluring for an experiment with multiple modules of increasing complexity. Additionally, it is more practical to engineer a large sample space in a dry system [Uhlig 2012].

However, the PT causes vibrations to which our high finesse optical cavities are very sensitive. To offer the convenience of a dry system, but the ability to temporarily shut off the PT for vibration-sensitive measurements, commercial dilution refrigerator companies are installing “helium batteries” in their systems [BlueFors 2023]. This is a small tank attached to the 4 K plate into which gaseous helium can be condensed while the PT is running during normal operation of the fridge, and then can keep the 4 K plate cold as the helium boils off when the PT is shut off. The BlueFors LD400 cryogen-free dilution refrigerator we purchased (installation planned by Luca Talamo) has a helium battery that allows for approximately two hours of measurement with the PT off, after which the battery must be refilled for 8 hours or more. Not only is the battery more efficient with helium consumption because it only is boiled off during the measurement time, but because it takes helium gas, there is not the trouble of sourcing and handling liquid helium, which boils off when not in use. The next sections will describe the more salient considerations in transferring our optomechanical measurement setup to this new cryostat.

6.1.1 Coupling optics

As shown in Fig. 6.1(a), our previous setup, the Cryoconcept Horizontal 200 wet system, is mounted directly to our optical table, which ensures relative stability between the coupling optics and the cavity mode. The fridge is oriented such that thermal contraction occurs along the optical axis, ensuring that upon cooling down, the cavity mode’s translation orthogonal to the beam axis is minimal, minimizing coarse readjustments to the coupling optics upon cooling down.

The current setup, the dry system shown in Fig. 6.1(b), is free-standing and oriented in the traditional vertical orientation of dilution refrigerators. We choose to mount our optics to




Figure 6.1: **Dilution refrigerator measurement setup photographs** (a) The CryoConcept Horizontal 200 dilution refrigerator wet system in which the measurements described in previous chapters was performed. Measurement is performed in reflection along the beam line indicated in red, along the direction in which the fridge contracts upon cooldown. The vacuum chamber enclosure and radiation shields are not in place in the photograph for initial alignment. The breadboard with the input coupling optics at the bottom of the image slides out of the way for installation of the radiation shields and vacuum chamber enclosure. (b) The BlueFors LD400 cryostat with measuring optics mounted to the frame of the dilution refrigerator. The measurements are performed in reflection through the side windows using the beam path indicated by the red line.

a breadboard mounted to the frame supporting the cryostat. We have found that any relative vibration between the device and the coupling optics from this somewhat less direct connection does not limit the coupling to our optomechanical devices. Though our fridge was purchased with an active vibration dampening system, we found it better to bypass this system and rigidly affix the cryostat to the frame to prohibit long-timescale drifts in our optical coupling to the cavity mounted in the cryostat, due to e.g. load changes on the fridge or thermal fluctuations in the room affecting the support of the still line. The vertical orientation does mean we have to steer the optics to follow the cavity as it translates upward during cooldown from room temperature. We have found this to be relatively reproducible, and dominantly a pure change in the upward direction, so in practice it is sufficient to steer the beam upward with a single mirror knob to re-find cavity resonance, when starting from the initial coupling configuration with the cryostat at room temperature.

6.1.2 Vibrations from the pulse-tube

Though the helium battery allows for low-vibration measurement, because of its limited hold time, we would like the capability to perform whatever measurements are possible while the PT is running. Coarse alignment tasks can be done with the PT running, though we observe large frequency fluctuations of the cavity resonance with every pulse from the cooler. Measurements that require the laser to be locked to the cavity but would be insensitive to additional phononic occupation due to vibration from the PT, such as the swept-sideband measurements described in Sec. 5.6.3, cannot be carried out because the lock is unable to correct for these large frequency excursions. These fluctuations may be due to the PT exciting vibrational modes of either mirror, or of a differential mode of the spacer supporting the mirrors. The dominant spectral contribution of the cavity motion we observe is approximately 20 kHz. The Schliesser group has engineered a promising solution to isolate the optical cavity from vibrations along the cavity axis with a compliant mounting bracket [Planz et al. 2023].

6.1.3 Challenges working with the helium battery

Over the course of running the helium battery during measurement, the temperature of the 50 K plate increases to approximately 100 K in two hours, which can lead to additional complications. First, the change in temperatures causes the 50 K stage to expand, leading to a noticeable change in modematching to the cavity. A rough calculation explains this. From Table 3.3 in Ref. [Ventura and Risegari 2008], stainless will expand by about 0.03% between 40 K and 100 K. So for an approximate length scale of 30 cm for the 50 K stage, we would expect the optical cavity to shift downwards by approximately 100 μm . Compared to the cavity mode radius of 40 μm , this would lead to a noticeable change in modematching.

Additionally, we have observed that the dilution unit will sometimes form a plug after running the battery when the fridge has been cold for several weeks. Our qualitative observations are that plugging seems correlated with restarting the PT, or when the temperature of the 4 K plate is increased if the battery is accidentally run dry, for example. It also seems to be more dependent on the length of the cool down, rather than the number of times the battery has been cycled. We have also fixed a partial plug by collecting the mix and elevating the mixing chamber temperature to approximately 20 K temporarily. Upon recondensing after this attempt, the fridge impedance had returned to its normal value.

There is an internal cold trap that is located at the 50 K plate in our system. It is possible that every time we run the battery and the 50 K stage reaches 77 K, that any air trapped in the internal trap migrates to the colder parts of the fridge and condenses. Another possibility is that hydrogen, e.g. as a byproduct of the circulation scroll pump operating, is accumulating in the plumbing of the 4 K stage, and when the temperature of that stage increases, that moves to the dilution unit where it results in a plug. The accumulation of hydrogen being the cause would be consistent with the correlation between plugging and the 4 K plate's increase in temperature, rather than a general correlation with battery use, as well as the observation that increasing the mixing chamber temperature to 20 K resolved the partial plug.

6.1.4 Spring-mounted curved cavity mirror

In his thesis, Mark Stone observed fluctuations of a similar frequency scale of 3 kHz due to vibrational modes of mirrors held with cryogenic epoxy being excited by a PT [Stone 2021]. Initially inspired by the improvement in performance he reported in clamping his cavity mirrors in place with Belleville spring washers, and shifting up the mechanical resonance frequencies of the cavity mirrors, we decided to make a similar design change in how our front mirror is mounted. Fig. 6.2 shows a rendering and photographs of this design, which uses two 3/32-inch Belleville springs (McMaster part number 9713k52) with 15 lbs working load in the parallel configuration to increase the clamping force to 30 lbs. Sarah Dickson built a cavity using this design, and measured the mechanical resonances in the spectrum of the PDH lock error signal, and found that this did not resolve the lowest-lying 20 kHz resonance. Nonetheless, clamping our curved mirrors to make our optical cavities is an improvement in that it is a simpler and faster assembly process that doesn't rely on epoxy, which we have observed to degrade from the stress of repeated thermal cycling, ultimately leading to the failure of the glue joint.

In using this new mirror mounting style, care must be taken to avoid increased cavity birefringence. Previously, optomechanical cavities with glued mirrors often had birefringences of about 1 MHz, and once we introduced the bonded etalon design described in Sec. 4.2.2.1 this increased to 2-3 MHz, likely from additional strain on the bonded mirror coating. On the second cooldown measuring a device with a clamped mirror, we observed a birefringence of 8 MHz. Birefringence of a few MHz can be useful for isolating any lock leakage from contaminating the transduced signal, but if it becomes too large can be inconvenient to separate the lock and damp beams with our current beam preparation setup (Fig. 5.5). Furthermore, a fractional change in a larger static birefringence would have magnified effects, and could lead to drift of the pump frequency relative to the cavity mode, if the birefringence changes e.g. due to strain changing with optical circulating power. By loosening and re-tightening the screws clamping the mirror assembly in place before every cooldown, it seems we have alleviated this problem, as we have observed birefringences of

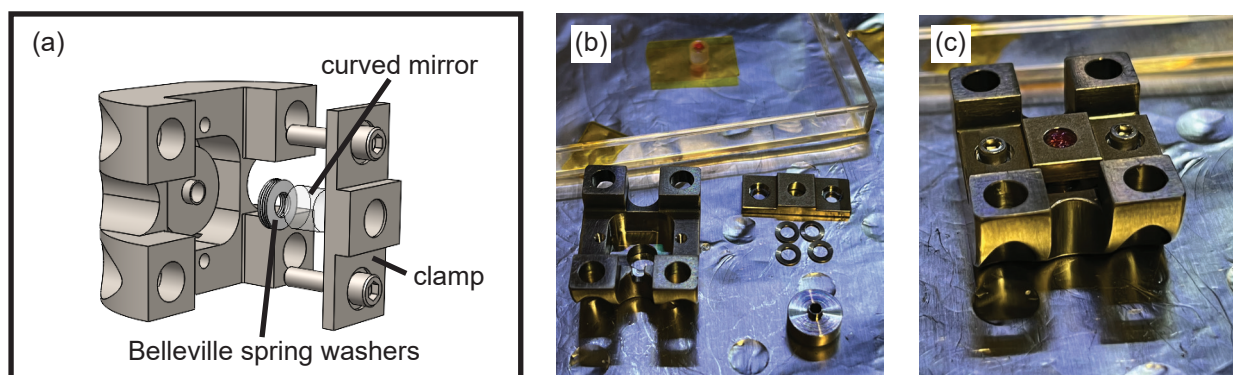


Figure 6.2: **Spring-mounted mirror design** (a) Computer Aided Design graphic illustrating the spring-mounted mirror design. The mirror substrate (transparent cylinder) is mounted to the curved mirror holder (dark gray) by being compressed against the spacer (light gray cylinder) and Belleville spring washers (annular disks) by the clamp bar (light gray) and fastened with screws as indicated. Springs are oriented such that the surface at the inner diameter of the spring exerts force against the mirror substrate. Assembly is aided by holding the spring washers in place gently with First Contact lens cleaning polymer, which can be easily and cleanly removed. This is the curved front mirror, and optical access is provided through the center holes in the spacer, springs and clamp. All metal parts are made from Invar other than the stainless steel screws and 302 stainless steel Belleville spring washers. (b) and (c) Photographs of parts before and after assembly, respectively. Device shown was assembled by Sarah Dickson.




Figure 6.3: **Clamping the cold lenses.** (a) Cold input-coupling lenses epoxied to Invar lens flange. A crack is visible in the lower portion of the right lens. (b) Cold lens clamped to Invar lens flange with beryllium copper shim spring. Caution: beryllium copper dust is very hazardous if inhaled. It should therefore not be machined or sanded—use snips and punches.

approximately 2 MHz in the three devices we’ve cooled down with screws reset in this manner.

6.1.5 Clamping the cold lenses

We have also observed the epoxy joints strain and crack the “cold lenses” used to couple into the optical cavity, as shown in Fig. 6.3. These are relatively fast lenses (18 mm focal length, Edmund Optics part number 87-941), that we use to decrease the beam diameter required to efficiently couple our beams to the cavity mode. To avoid the strain of an epoxy joint we use a beryllium copper shim spring to hold the lens in place, shown in Fig. 6.3(b).

6.1.6 Optomechanical ground-state cooling in a cryogen-free dilution refrigerator

It is important that we are able to optomechanically cool the membrane’s transduction mode below an occupation of one phonon in our new fridge setup. Figure 6.4(a) shows the results of an optomechanical damping sweep, plotting the mechanical occupation n_m vs. the optomechanical damping rate Γ_o . The occupation is inferred by both sideband asymmetry thermometry as described in Sec. 5.7.3 and using the thermally driven peak height and independently measured readout efficiency as described in Sec. 5.7.2. However, rather than sweeping the base plate temperature to calibrate the measurement efficiency which is more time consuming, we instead independently measured individual contributions: κ_o , modematching coefficients, optical path losses, the detector

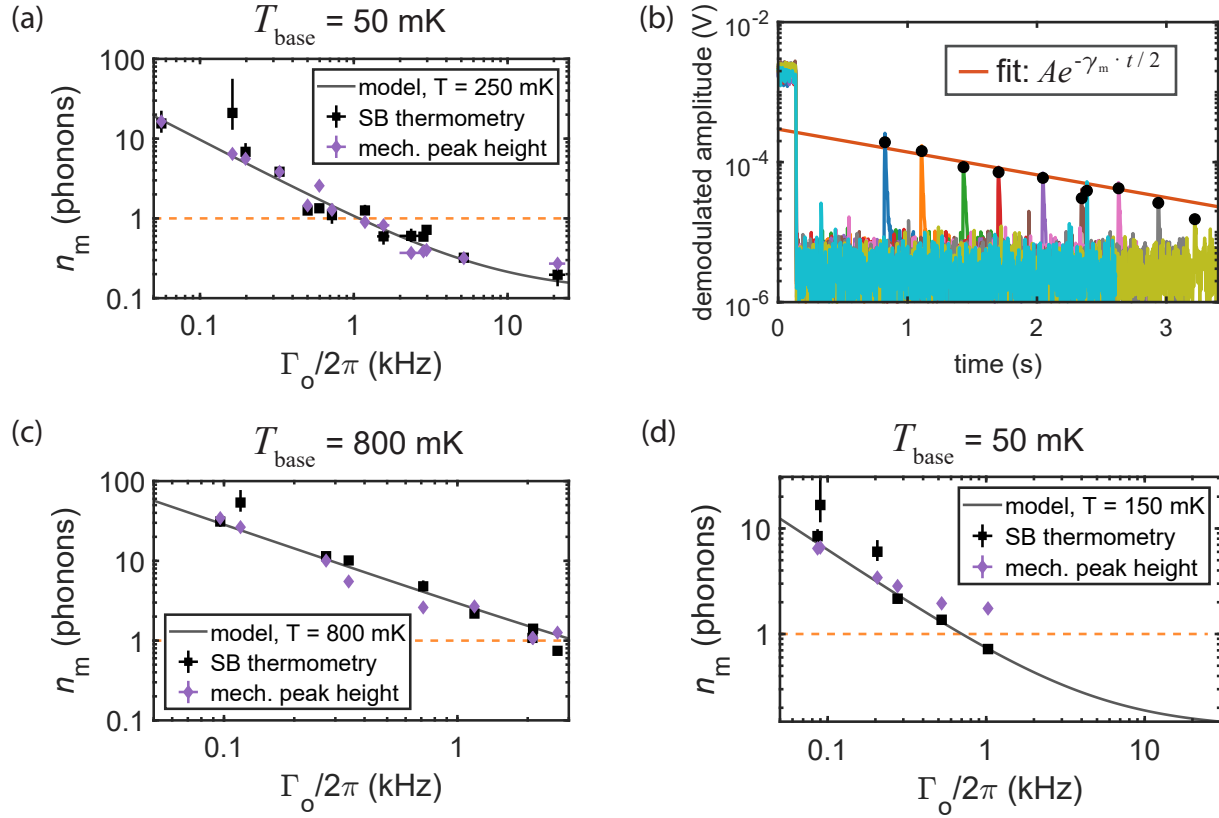


Figure 6.4: **Ground state cooling in a cryogen-free dilution refrigerator** (a) Mechanical occupation n_m vs. Γ_o as inferred from sideband asymmetry thermometry (black points) and mechanical peak height and measurement efficiency (black points). Increasing Γ_o cools the mode to below an occupation of one phonon (orange dashed line), close to its ground state. Gray curve indicates theoretical model defined by Eq. 6.1 using parameters as listed in 6.1, and an environmental temperature $T = 250$ mK, indicating thermalization to a higher temperature compared with previous devices. (b) Overlaid traces from stroboscopic ringdown measurements of variable delay time. Line of best fit through the initial values of each decay (black points) indicates a value of $\gamma_m/2\pi = 240 \pm 40$ mHz. (c) n_m vs. Γ_o as in (a) with an elevated base plate temperature of $T_{\text{base}} = 800$ mK. Theoretical model assumes an environmental temperature of $T = 800$ mK. (d) n_m vs. Γ_o in a subsequent cooldown after making some improvements to thermalization. Theoretical model assumes an environmental temperature of $T = 150$ mK.

efficiency, and the expected value of $\kappa_{\text{o,ext}}$ from cavity simulations similar to those in Sec. 4.2.2.3. As we increase Γ_{o} we again cool below a mechanical occupation of one phonon and approach the backaction limit for this device, $n_{\text{min}} = 0.11$.

6.1.7 Mechanical ringdown with stroboscopic optical readout

An independent measurement of γ_{m} was used to get information about the environmental occupation from this sweep. Because this transducer device was not connected for microwave measurement, we could not use the technique described in Sec. 5.6.3 to measure γ_{m} . Instead, we performed an optomechanical ringdown measurement, using stroboscopic readout. To avoid the damping from the lock beam γ_{lock} dominating the measurement, we performed the measurement without the lock beam, and relied on the passive stability of the laser with respect to the cavity. Similar to the microwave measurement, we turned on a temporary SSB probe tone detuned $+\omega_{\text{m}}$ from the pump to ring up the mechanical mode. We then blocked the copropagating pump and probe beams, and allowed the mechanical mode to decay at γ_{m} before unblocking the pump beam after a delay time to read out the mechanical motion. Repeating this procedure with a variable delay time, and interleaving OMIT measurements to ensure the pump had not drifted relative to the cavity, allowed us to coarsely measure γ_{m} , as shown in Fig. 6.4(b). Fitting an exponential decay through the initial value of each more rapid decay, gives a value of $\gamma_{\text{m}}/2\pi = 240$ mHz. The more rapid decays occur at rate Γ_{o} due to the readout beam. The initial value is defined as the average of the values of the maximum voltage point and the two subsequent data points, and is indicated by the black points.

6.1.8 Inferring the mechanical environmental temperature

With the measured value of $\gamma_{\text{m}} = 240 \pm 40$ mHz, the environmental occupation n_{th} appeared to be significantly higher than expected, corresponding to a temperature of approximately 250 mK.

The gray curve in Fig. 6.4(a) is model with no free parameters of the occupation

$$n_m = \frac{n_{\text{th}}\gamma_m + n_{\text{min,o}}\Gamma_o + n_{\text{min,lock}}\gamma_{\text{lock}}}{\Gamma_o}, \quad (6.1)$$

using measured κ_o and assuming similar lock backaction to that determined in Sec. 5.6.4, and an environmental bath temperature of 250 mK, though our mixing chamber temperature was approximately $T_{\text{base}} = 50$ mK. The parameters used for the model are in Table 6.1. As an additional check on our thermal calibration, we elevated the temperature of the base plate to $T_{\text{base}} = 800$ mK and repeated a measurement of n_m vs Γ_o , shown in Fig. 6.4(c). Comparing the measured occupations with a model assuming a thermal bath of 800 mK gave consistent results, so something was preventing the device from thermalizing to the temperature base plate, whereas our previously measured devices thermalized to temperatures of approximately 100 mK.

subfigure	$\gamma_m/2\pi$	$\kappa_o/2\pi$	$n_{\text{min,o}}$	$n_{\text{min,lock}}\gamma_{\text{lock}}$
(a) and (c)	240 mHz	2.0 MHz	0.11	$2\pi \cdot 100$ Hz
(d)	240 mHz	2.1 MHz	0.13	$2\pi \cdot 100$ Hz

Table 6.1: Parameters used for ground-state cooling model in Fig. 6.4

To improve the thermalization of the device, we warmed up the fridge, tightened the screws of the mounting bracket and shielded the transducer with an OFHC radiation shield. Upon cooling the device down, measurement indicated that the device was now reaching a temperature of 150 mK, as shown in Fig. 6.4(d). After this cooldown, we removed the bracket to plate it with gold for a better thermal connection at the interfaces, and saw that the interfaces of the bracket had been contaminated with thermal grease which is ineffective at providing good thermal contact at millikelvin temperatures. In the time since, we haven't remeasured an optomechanical device to verify the environmental temperature, but we believe that it will be lower than 150 mK after cleaning the bracket, plating it, and replacing it.

In summary, we have successfully cooled a mechanical transduction mode to its ground state in our new cryogen-free dilution refrigerator setup, and have characterized many differences compared with our previous setup to allow for measurement of a transducer device.

6.2 Noise considerations for future experiments

Although we already have strict requirements on our laser noise in order to achieve optomechanical ground state cooling with a MHz-frequency mechanical mode, the noise requirements for an experiment involving photon counting are even more stringent. Moreover, because the single-photon detectors are sensitive to radiation across the electromagnetic spectrum, there are more potential noise sources to contaminate the experimental signal. As with our previous characterization of the transducer, we must contend with the noise from the transducer as given by Eq. 4.13, which we will continue to mitigate as much as possible. In the following sections, I consider additional sources of noise that have the potential to limit our capability to measure a nonclassical $g^{(2)}$ or perform a heralding experiment. As discussed in Ch. 3, we are sensitive to small amounts of technical noise on our pump that, when integrated over the bandwidth of the three cascaded filter cavities indicated in Fig. 3.4, can contribute significant count rates, as well as any pump leakage through the filter cavities. Any mechanical motion of our transducer cavity mirrors at frequencies close to that of our membrane mode can also add noise counts, as we read out this undesired optomechanical signal in addition to that of our membrane.

6.2.1 Laser technical noise

To determine the amount of technical noise we can tolerate, I'll calculate an expression for the noise we would require to achieve an SNR of one. Expressing the signal rate as an equivalent white noise spectral density in units of dBc/Hz can be useful as this is how the noise performance of our lasers, and microwave and RF signal generators is specified, and is a value we have quantified for our optical measurement setup in the past. I'll consider the signal rate for an experiment in which we pitch excitations from a superconducting qubit. At the transducer output, the signal power is given by

$$P_{\text{sig}} = \frac{\eta_{\text{MW}}\eta_{\text{t}}}{T_{\text{rep}}} \hbar\omega_{\text{o}}, \quad (6.2)$$

where η_{MW} is the microwave transmission efficiency from the qubit to the transducer input, and I'll set the repetition time to $T_{\text{rep}} = 2\pi/\Gamma_{\text{t}}$. The optomechanical damping rate is given by

$$\Gamma_{\text{o}} = \frac{4 \frac{\kappa_{\text{o,ext}}}{\kappa_{\text{o}}} \epsilon_{\text{PC}} P_{\text{pump}}}{\hbar \omega_{\text{o}}} \frac{g_{\text{o}}^2}{\omega_{\text{m}}^2} \quad (6.3)$$

in the resolved-sideband limit, where P_{pump} is the pump power incident on the cavity [Schliesser et al. 2008]. The ratio of signal power to pump power is then found to be

$$P_{\text{sig}}/P_{\text{pump}} = \frac{4\eta_{\text{MW}}\eta_{\text{t}} \frac{\kappa_{\text{o,ext}}}{\kappa_{\text{o}}} \epsilon_{\text{PC}}}{2\pi} \frac{\Gamma_{\text{t}}}{\Gamma_{\text{o}}} \frac{g_{\text{o}}^2}{\omega_{\text{m}}^2}. \quad (6.4)$$

Assuming that $\Gamma_{\text{t}} = \Gamma_{\text{e}} + \Gamma_{\text{o}} = 2\Gamma_{\text{o}}$ for matched transduction, and parameters given in Table 6.2, this is a ratio of -96 dBc. Distributing this power relative to the pump over the 30-kHz effective

parameter	η_{MW}	η_{t}	$\frac{\kappa_{\text{o,ext}}}{\kappa_{\text{o}}}$	ϵ_{PC}	$g_{\text{o}}/2\pi$	$\omega_{\text{m}}/2\pi$
value	0.3	0.5	0.9	0.9	60 Hz	1.5 MHz

Table 6.2: Parameters for estimation of noise required for $\text{SNR} = 1$.

bandwidth of the filter cavities gives a value of -141 dBc/Hz. So to achieve an SNR of one, we would require a white noise density of -141 dBc/Hz and an otherwise noiseless transducer.

It is also worth noting, from the dependence on ω_{m} in Eq. 6.3, that increasing the membrane resonance frequency does not reduce the noise filtered through the filter cavities for the same value of Γ_{o} , up to any change in g_{o} . Though the technical noise follows a Lorentzian roll-off, the power would need to increase by the same amount.

6.2.2 Increasing filter cavity lock bandwidth

One source of laser technical noise can be from the active feedback in a PDH lock. At sufficiently high frequency, the phase delay acquired by the signals propagating through the feedback loop will cause the feedback loop to amplify perturbations rather than suppress them. By decreasing the phase delay by shortening cable and fiber lengths, and ensuring that the bandwidth of electrical filters do not add delay, it is possible to then increase the gain and bandwidth of the feedback loop.

We were able to do this on the PDH servo loop locking our the laser to the filter cavity, by shortening the electrical cables associated with the fast loop and the optical fibers, illustrated in Fig. 6.5(a). We increased the proportional gain of the FALC110 feedback loop controller to see the fast loop ring at 300 kHz, and then adjusted the fast limited integrator and fast limited differentiator stages to decrease gain at this frequency, and compensate for the phase delay at higher frequencies. We also decreased the gain of a preamplifier at the input of the FALC110, as we found it was adding excess noise. With these steps, we were able to decrease the error spectrum noise at 1.5 MHz by over 10 dB, though perhaps the lock parameters had drifted since performing the previous optimization, and so the improvement in laser noise was mitigated. We then measured the phase noise of our laser downstream, using an unbalanced Mach-Zehnder interferometer with a delay line of length 67 m tuned to measure phase noise at 1.5 MHz (details about this setup will be described in Luca Talamo’s thesis) [Parniak et al. 2021]. We found that the phase noise had improved moderately, to -139 dBc/Hz.

6.2.3 Additional filter cavity

To further suppress the technical noise at 1.5 MHz, we introduced a 130 kHz filter cavity (designed and built by Sarah Dickson) before the fiber amplifier, as shown in Fig 6.5(b). We lock the cavities sequentially, such that the laser is first locked to the first filter cavity with the fast output of the PDH lock feeding back to the laser current, and the slow output feeding back to the laser piezo, as described in Sec. 5.5 for the previous setup. A seed for the fiber amplifier is then provided with the light transmitted through the first cavity, allowing the amplifier to be powered on. The first cavity is then locked to the second filter cavity, with the PDH lock only feeding back to the piezo of the second cavity. Finally, we lock the transducer cavity as described in Sec. 5.5, with the slow channel of the PDH lock feeding back to the second filter cavity, and the fast channel feeding back to the double-pass AOM (not shown in Fig. 6.5).

Placing the new filter cavity before the amplifier allowed us to keep our beam paths unmodified, and allowed the circulating power in the new filter cavity to remain low. Measuring the phase

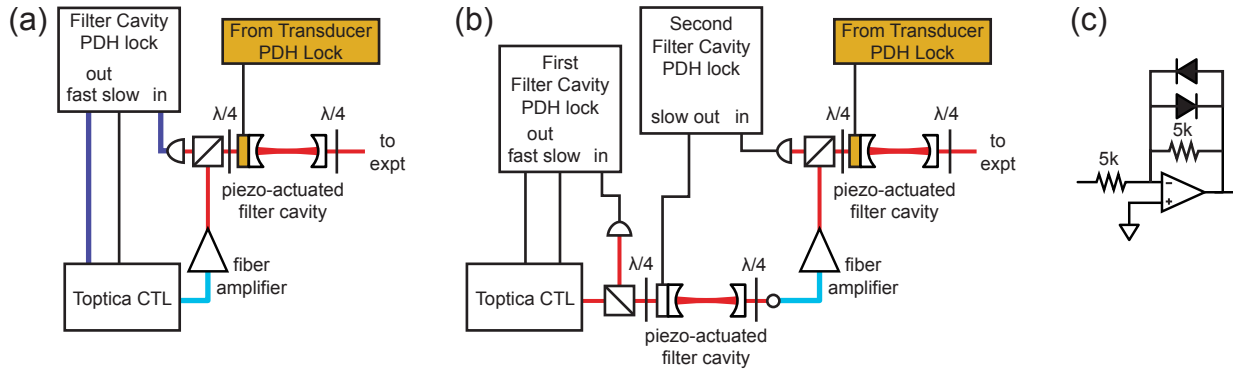


Figure 6.5: Decreasing laser technical noise (a) Diagram of feedback loop locking the laser to single filter cavity, as described in Sec. 5.5. By shortening the electrical cables (purple lines) and optical fibers (blue line) to reduce the phase delay, we could increase the loop bandwidth and gain and reduce the error signal spectrum at 1.5 MHz. (b) Diagram of setup after introducing the first filter cavity before the fiber amplifier. The laser is first locked to the first filter cavity with feedback to the laser current and piezo as in the previous setup, described in Sec. 5.5. The amplifier can then be powered on, and the first filter cavity locked to the second filter cavity by feeding back to the first filter cavity's piezo. The transducer can then be locked as before, with the fast feedback channel modulating the frequency of a double-pass AOM (not shown here, see Fig. 5.5), and the slow feedback channel actuating the second filter cavities piezo. The first filter cavity and laser then follow the second filter cavity. (c) Voltage clamping circuit in second filter cavity and transducer PDH lock loops, to limit the loop output voltage, preventing integrator wind up from unlocking the laser from the first filter cavity and causing the fiber amplifier from losing seed power.

noise at 1.5 MHz before the first filter cavity and after the amplifier indicated that the cavity was suppressing the noise by 26 dB, within 1 dB of the 27 dB expected for a 130 kHz-linewidth cavity. The phase noise directly out of the second filter cavity was then measured to be -158 dBc/Hz (again using the convention of one-sided double-sideband spectral densities), 4 dB greater than the expected suppression for a cavity with a linewidth of 80 kHz. With additional improvements to the RF signals driving the pump AOMs, we were able to reduce the noise at the input of the transducer by an additional 12 dB, from -142 dBc/Hz to -154 dBc/Hz. The first improvement was to improve the impedance matching circuit of a surface acoustic wave (SAW) bandpass filter that prepares the RF tone supplied to the AOM. By mounting the SAW filter to a PCB designed by Terry Brown, the rejection of the filter was improved by 20 dB to 60 dB. The second improvement was bypassing the SSB modulation module, which was limiting noise to the -145 dBc/Hz level. We can bypass the SSB modulation module for the sensitive spectrum measurements, as as we only use SSB modulation for OMIT measurements.

There is a risk associated with introducing a filter cavity before the fiber amplifier: if the laser comes unlocked from the first cavity, then the amplifier will lose its seed, which risks damaging the amplifier. Because we have two other cavities locked downstream, this is likely to occur if one of the other locks is perturbed and loses lock, causing the integrator to “wind up,” or integrate a small voltage offset over time. Integrator wind-up will then cause the cavity to shift too far in frequency for the laser to follow. As a safeguard, we introduced simple voltage clamping circuits into the lock servo loops of the second filter cavity and the transducer cavity in order to limit the maximum voltage supplied to the piezos of both the first and second cavity. The voltage clamp is shown in Fig. 6.5(c) and is a follower circuit shunted with two 1N4148 switching diodes. This limits the output of the follower circuit to the forward bias voltage of the diodes, which is 600 mV for the 1N4148 diode. If a different voltage is needed, additional diodes can be placed in series, and Schottky diodes have a forward bias voltage approximately half that of the switching diodes we used. We thank Terry Brown for this clever and simple solution.

With these improvements, the noise level of -154 dBc/Hz would contribute 0.05 counts for

every signal count from a superconducting qubit, assuming the experimental parameters in Table 6.2.

6.2.4 Leakage of the coherent pump

Any leakage of the pump through the filter cavities will be an additional source of extraneous counts. The filter cavities have been measured to suppress a probe tone by 107 dB at a detuning of 1.5 MHz relative to the transmission on resonance. We can again compare this to an equivalent white noise power spectral density by dividing by the bandwidth of the cascaded filter cavities, and get -152 dBc/Hz. This contribution is below our signal power by 11 dB, so it should not limit us for most of the protocols considered in Ch. 3. Again comparing this source of noise counts to the experimental signal described in Sec. 6.2.1, pump leakage would add 0.08 noise counts per signal count from a superconducting qubit.

If we require greater suppression of this tone, we could interferometrically cancel the tone by splitting off the pump upstream of the transducer, and recombining it with the appropriate phase shift. Because of the frequency-dependent phase response of the cavity, this would however lead to constructive interference of the technical noise on cavity resonance. A suitable delay line could be used to manipulate the phase of the canceling tone in order to cancel both the coherent pump and the noise on cavity resonance simultaneously.

6.2.5 Thermal motion of cavity mirrors

Though our optical cavity is operating at cryogenic temperatures, thermomechanical motion of the cavity mirror substrates will imprint noise onto the light because of its relatively short, mm-scale length. This noise source is particularly insidious, as it is written onto the light at the same time as our transduced signal, so, unlike other sources of technical noise on our light, it cannot be canceled or filtered out upstream. To get a rough idea of the scale of the noise, we can use an earlier measurement: the phase noise measurements described in Sec. 5.7.3 used an optical cavity of the dimensions of our transducer cavity, but without a membrane, operating at room temperature.

The spectrum included peaks from the thermomechanical motion of the mirror substrates. At frequencies near 1.5 MHz, the noise peaks rose to a level of approximately -115 dBc/Hz. Scaling the temperature from 300 K to 100 mK, we would expect peaks to contribute noise at a level of roughly -150 dBc/Hz, approximately the level we are sensitive to. We don't know the temperature the mirror substrates thermalize to when operated in the cryostat, and one can expect that in general the loss of the mirror substrate modes can change upon cooling down. The Ye group has measured the loss tangent of the $\text{SiO}_2/\text{Ta}_2\text{O}_5$ dielectric mirror stacks used for mirror coatings using their ultrastable crystalline silicon Fabry-Pérot cavity at cryogenic temperatures of 124 K, 16 K, and 4 K [Robinson et al. 2021]. They observed weak temperature-dependence, with an *increase* of roughly a factor of two in loss when cooling from 124 K to 4 K. However, measurements of the mirror coating loss have not been carried out at mK-temperatures. Furthermore, our cavity geometry is substantially different from the crystalline silicon cavities, which could lead to different temperature dependence. To determine if this noise is large enough to affect a photon counting experiment, we must measure it.

To measure the mirror substrate motion, we perform spectrum measurements similar to those used to characterize the thermal occupation of our mechanical mode, as described in Sec. 5.7.2, instead measuring a cavity without a membrane. This is essentially the same measurement described in Sec. 5.7.3 that we performed to measure the phase noise of the laser using an empty cavity, with the caveat that we were not optimally detuned in the earlier measurement, and our calibration method was less direct. Here, we first obtain shot-noise normalized spectra $S_{\text{det},\text{o},+}$ by normalizing to the corrected shot noise value $\text{SN}_{\text{corr}} = \text{SN}_{\text{meas}} - \text{DN}$, by subtracting the detector dark noise contribution DN from the measured value of the shot noise SN_{meas} obtained by blocking the pump beam. This extra care was necessary due to a larger dark noise fraction, as we needed to reduce the LO power to avoid saturating the detector at the higher optical powers required for this measurement. We cast the normalized spectrum into position spectral density units [Teufel et al.

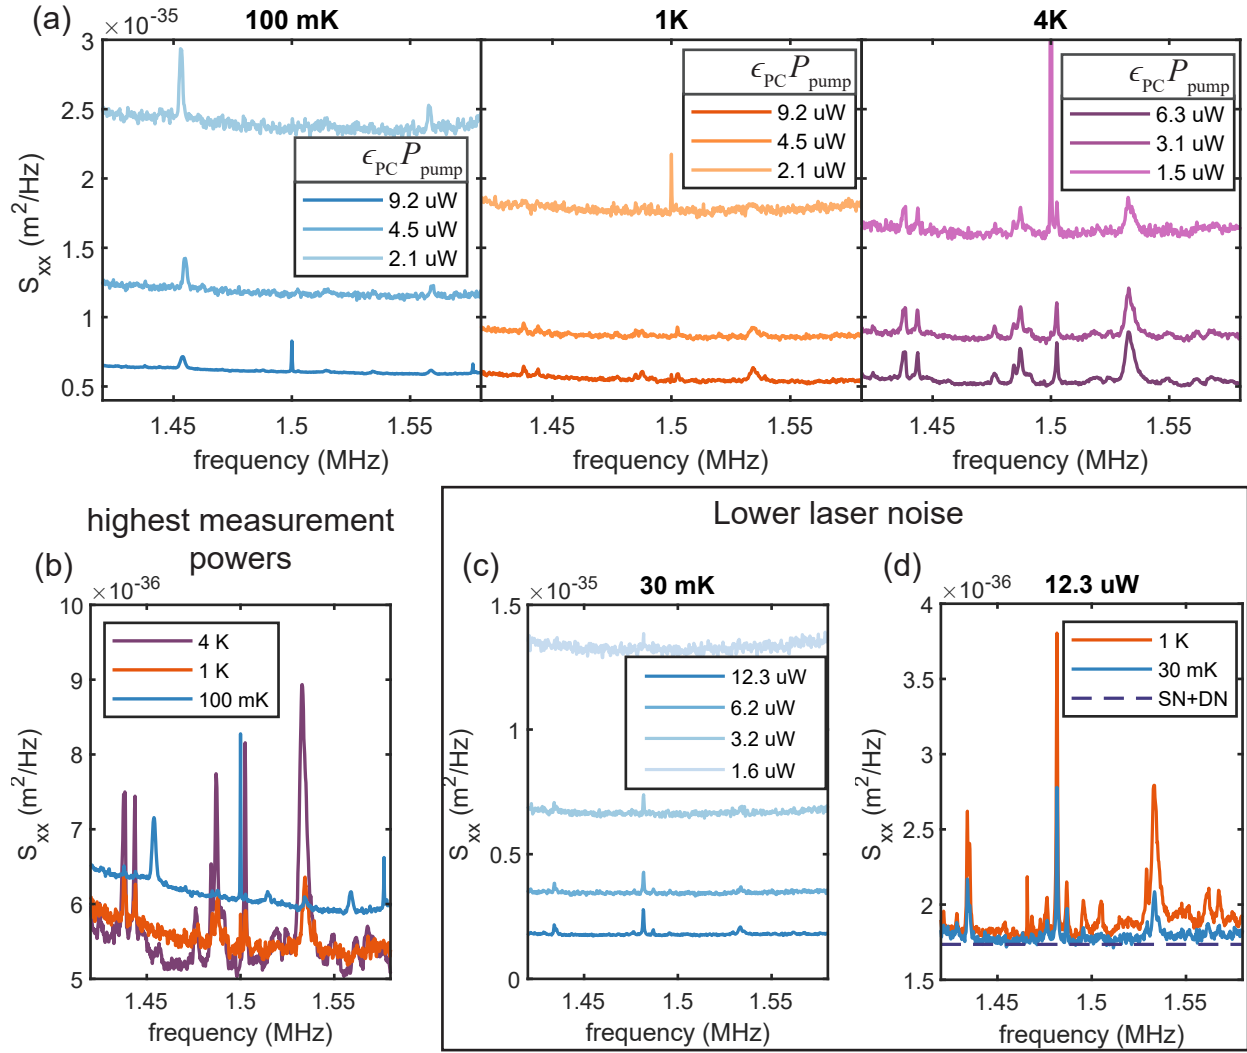


Figure 6.6: Optical spectra showing thermomechanical mirror motion. (a) Optical spectra measured at base plate temperatures of $T_{\text{base}} = 100 \text{ mK}$ (blue), $T_{\text{base}} = 1 \text{ K}$ (orange), and $T_{\text{base}} = 3.5 \text{ K}$ (purple) calibrated in position units given by Eq. 6.5. At each temperature, increasing mode-matched optical power $\epsilon_{\text{PC}} P_{\text{pump}}$ reduces imprecision noise, while technical noise and thermomechanical motion are independent of power in these units. The $T_{\text{base}} = 100 \text{ mK}$ measurements also show electrical noise peaks that decrease with laser power. (b) Comparison of the highest power measurements at each temperature. The $T_{\text{base}} = 100 \text{ mK}$ spectrum taken at high power was averaged 200,000 times, and is possible to resolve the largest thermal peak superimposed on the higher levels of technical noise. All other traces were taken with 10,000 averages. (c) Optical spectra measured at a base plate temperature $T_{\text{base}} = 30 \text{ mK}$ at indicated values of $\epsilon_{\text{PC}} P_{\text{pump}}$, in a device similar to the device used in (a), after improving the laser technical noise to -154 dBc/Hz . The thermomechanical motion is readily identifiable at the lower temperature with the improved laser noise. (d) Optical spectra measured with improved laser noise at base plate temperatures $T_{\text{base}} = 30 \text{ mK}$ and $T_{\text{base}} = 1 \text{ K}$. The dashed line indicates the measurement noise due to shot noise imprecision and detector dark noise.

2011], as

$$S_{xx} = \left(\xi_o \epsilon_{\text{CL}} \frac{2G_{\text{cav}}^2 n_{\text{cav}}}{\kappa_o} \frac{\kappa_{o,\text{ext}}}{\kappa_o} \right)^{-1} S_{\text{det,o,+}}, \quad (6.5)$$

where $G_{\text{cav}} = \omega_o/L$ is the cavity frequency shift per mirror displacement, and n_{cav} is the number of photons circulating in the optical cavity.

Fig. 6.6(a) shows a first attempt at measuring mirror motion at cryogenic temperatures. The spectrum measurements were repeated at three pump powers, at base plate temperatures of $T_{\text{base}} = 100 \text{ mK}$, 1 K , and 4 K , with the indicated optical measurement powers. The spectra taken at $T_{\text{base}} = 1 \text{ K}$ and 4 K show clear mechanical features that are common to the two temperatures and increase with T_{base} . The mechanical peak heights remain relatively constant with increasing optical power because the modes are too massive to be damped and cooled by the radiation pressure of the light. An increase in peak height with optical power could indicate increased thermal occupation with higher optical circulating powers, although it is likely superimposed with the artifact of changing mode matching over the course of using the helium battery discussed in Sec. 6.1.3. These two effects would result in a trend in the same direction, as the highest power points were measured first. Reversing the order in which these measurements were taken in one measurement run indicated that the changing modematching was a more dominant effect than an increase in thermal occupation from higher optical powers.

The $T_{\text{base}} = 100 \text{ mK}$ measurements show structured noise that decreases with increasing readout power, as would be expected for an independent noise source in these position units—in this case it was electrical pickup from a damaged cable connected to the HF2LI. Unfortunately, the dilution unit plugged at the beginning of that helium battery run, so we did not have time to diagnose and fix that issue before the fridge warmed to 1 K . To compensate, we took a final measurement at the highest power with 200,000 averages, hoping to be able to resolve a thermomechanical peak superimposed on the technical noise.

Fig. 6.6(b) shows the highest measurement powers at each value of T_{base} , and the largest thermomechanical peak at 1.53 MHz can be seen on the 100 mK trace. Crudely comparing the

height of that peak to the heights of the same mechanical feature at other temperatures (as a proxy for the peak area assuming the quality factor is somewhat constant) suggests that the mirror temperature is approximately 0.2–0.3 K. However, this value is likely to be optimistically cold, as we now suspect the modematching would have degraded at the end of the battery run.

Fig. 6.6(c) and (d) show measurements of a similar device after the laser noise was improved to a level of -154 dBc/Hz at 1.5 MHz. The optical cavity was constructed from the same mirrors, but with the curved mirror reseated in the spring clamp as described in Sec. 6.1.2, which could perturb the vibrational modes we measure. From the mm-scale wavelength of acoustic vibrations at 1.5 MHz in our mirror substrates, I expect that the vibrational modes of the flat mirror chip that we optically read out would remain the same, with similar couplings, though thermal cycling of the glue joints holding it in place could also perturb those vibrational modes as well. Thermal peaks are easily identifiable with the reduced laser noise, even at the low value of $T_{\text{base}} = 30$ mK. We see sparse, approximately kHz-wide features with spectral heights of approximately $1 \times 10^{-36} \text{ m}^2/\text{Hz}$. As a preliminary analysis, when comparing the heights of the thermal peaks with others measured at higher values of T_{base} , it appears that mirror modes are at a temperature of approximately 400 mK. However, it is possible that different mechanical resonances cool to different temperatures, depending on which mirror they are associated with. For example, in Fig. 6.6(d), the feature at 1.53 MHz appears to decrease in height more significantly than the feature at 1.48 MHz when the base plate temperature is decreased from 1 K to 30 mK.

To compare position spectral densities to an optical cavity of a given length, we can invert Eq. 6.5 to get

$$\frac{S_{\text{det,o,+}} \hbar \omega_o / \xi_o \epsilon_{CL}}{P_{\text{pump}}} = 2 \frac{\omega_o^2}{L^2} \frac{\kappa_{\text{o,ext}}^2}{\kappa_o^2} \frac{1}{\omega_m^2} S_{xx} \quad (6.6)$$

in the resolved sideband limit. For an overcoupled $L = 1$ mm cavity operating at 1064 nm, the prefactor to S_{xx} is $7 \times 10^{22} \text{ m}^{-2}$. So a value of $1 \times 10^{-36} \text{ m}^2/\text{Hz}$ would imprint -132 dBc/Hz of noise onto the light, and a single 1-kHz Lorentzian peak of that height, such as the feature at 1.48 MHz, would contribute the equivalent of -149 dBc/Hz of white noise distributed over the

30 kHz effective bandwidth of the filter cavities, or 0.16 noise photons per signal photon in the experimental context described in Sec. 6.2.1. So thermal occupation of a single mirror mode could appreciably eat into the noise budget in a qubit readout experiment if it falls within the bandwidth of the filter cavities.

Of course, not only the sparse peaks will contribute to noise counts, however it is difficult to resolve the noise floor of the thermal motion at these power levels and at mK temperatures without being sensitive to any systematic offsets in these measurements, as the contributions to the measurement imprecision background from shot noise and the dark noise of the detector dominate the low levels of thermomechanical background noise between peaks, as shown in Fig. 6.6(d). We can however, roughly infer the noise floor by comparing the spectrum at a higher temperature where we are more confident in the measurement of the mechanical background, such as 1 K, and scaling it down to the apparent mirror substrate temperature of approximately 400 mK, making the assumption that the quality factor of the mechanical resonances does not change appreciably between these temperatures. An excess mechanical noise level of $1 \times 10^{-37} \text{ m}^2/\text{Hz}$ from the 1 K measurement, for example, when scaled down in temperature by a factor of 0.4, would cause -146 dBc/Hz of noise to be imprinted on the light, and such a level of white noise integrated over the filter cavity bandwidth would result in 0.32 noise photons per signal photon in the qubit readout experiment described in Sec. 6.2.1, in addition to the noise contribution from any individual peak.

6.2.6 Phononically engineering mirror substrates

Though we cannot filter or cancel the noise from the mirrors' thermally driven motion, we could phononically engineer the mirrors to avoid resonances at frequencies near our transducer signal, as depicted in Fig. 6.7(a). Ideally, we would use lithographic techniques to make high-quality phononic crystals like those in Ref. [Yu et al. 2014]. The Rakich group has recently made exciting advances in microfabrication techniques to manufacture high-finesse mirrors with the cm-scale radii of curvature we require for our optical cavities [Jin et al. 2022]. Such mirrors made at wafer-scale using photoresist reflow and reactive ion etch techniques would in principle be compatible with




Figure 6.7: **Phononically engineered mirror substrates.** (a) Concept of cavity with mirrors integrated into a phononic crystal. The blue squares indicate mirrors integrated into structures with a phononic band gap at the transducer signal frequency. The mirror chips would then be affixed in such a way as to avoid differential modes between the chips, either by intimate contact or further phononic engineering. (b) Photograph of mini-mirror affixed to phononic crystal. The mirror substrate was lapped to a thickness of approximately $500\text{ }\mu\text{m}$ and diced into a $600\text{ }\mu\text{m} \times 600\text{ }\mu\text{m}$ square. Stycast 2850FT was then placed on the silicon phononic crystal structure and the mirror glued in place, with the high-reflectivity side facing down.

patterning phononic crystals in the silicon chip, to create mirrors that do not support acoustic modes near our transducer signal frequency, however this would be a somewhat involved undertaking. As a more rapid initial test, we have pursued lapping and then cutting superpolished mirrors to $600\text{ }\mu\text{m} \times 600\text{ }\mu\text{m}$ squares of few-hundred- μm thickness, to shift the lowest resonance frequency supported by the mirror substrate to MHz-frequencies. These “mini-mirrors” can then be affixed in a phonic crystal with epoxy, as shown in Fig. 6.7(b). However this is a slow and difficult process. In Ref. [Huang et al. 2024] the Kippenberg group engineers a phononic band gap directly in their fused silica mirror substrates in an innovative way using a dicing saw. This seems like a promising direction for our transducer, though ultimately the cleanest spectra and most reliable transducers will likely result from mirrors made using lithographic techniques.

References

- R. Acharya, I. Aleiner, R. Allen, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. Atalaya, R. Babbush, D. Bacon, J. C. Bardin, J. Basso, A. Bengtsson, S. Boixo, G. Bortoli, A. Bourassa, J. Bovaird, L. Brill, M. Broughton, B. B. Buckley, D. A. Buell, T. Burger, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, B. Chiaro, J. Cogan, R. Collins, P. Conner, W. Courtney, A. L. Crook, B. Curtin, D. M. Debroy, A. Del Toro Barba, S. Demura, A. Dunsworth, D. Eppens, C. Erickson, L. Faoro, E. Farhi, R. Fatemi, L. Flores Burgos, E. Forati, A. G. Fowler, B. Foxen, W. Giang, C. Gidney, D. Gilboa, M. Giustina, A. Grajales Dau, J. A. Gross, S. Habegger, M. C. Hamilton, M. P. Harrigan, S. D. Harrington, O. Higgott, J. Hilton, M. Hoffmann, S. Hong, T. Huang, A. Huff, W. J. Huggins, L. B. Ioffe, S. V. Isakov, J. Iveland, E. Jeffrey, Z. Jiang, C. Jones, P. Juhas, D. Kafri, K. Kechedzhi, J. Kelly, T. Khattar, M. Khezri, M. Kieferová, S. Kim, A. Kitaev, P. V. Klimov, A. R. Klots, A. N. Korotkov, F. Kostritsa, J. M. Kreikebaum, D. Landhuis, P. Laptev, K.-M. Lau, L. Laws, J. Lee, K. Lee, B. J. Lester, A. Lill, W. Liu, A. Locharla, E. Lucero, F. D. Malone, J. Marshall, O. Martin, J. R. McClean, T. McCourt, M. McEwen, A. Megrant, B. Meurer Costa, X. Mi, K. C. Miao, M. Mohseni, S. Montazeri, A. Morvan, E. Mount, W. Mruczkiewicz, O. Naaman, M. Neeley, C. Neill, A. Nersisyan, H. Neven, M. Newman, J. H. Ng, A. Nguyen, M. Nguyen, M. Y. Niu, T. E. O'Brien, A. Opremcak, J. Platt, A. Petukhov, R. Potter, L. P. Pryadko, C. Quintana, P. Roushan, N. C. Rubin, N. Saei, D. Sank, K. Sankaragomathi, K. J. Satzinger, H. F. Schurkus, C. Schuster, M. J. Shearn, A. Shorter, V. Shvarts, J. Skrzynny, V. Smelyanskiy, W. C. Smith, G. Sterling, D. Strain, M. Szalay, A. Torres, G. Vidal, B. Villalonga, C. Vollgraft Heidweiller, T. White, C. Xing, Z. J. Yao, P.

- Yeh, J. Yoo, G. Young, A. Zalcman, Y. Zhang, N. Zhu, and Google Quantum AI, “Suppressing quantum errors by scaling a surface code logical qubit”, *Nature* **614**, 676–681 (2023).
- T. P. M. Alegre, A. Safavi-Naeini, M. Winger, and O. Painter, “Quasi-two-dimensional optomechanical crystals with a complete phononic bandgap”, *Opt. Express* **19**, 5658–5669 (2011).
- R. W. Andrews, “Quantum signal processing with mechanical oscillators” (University of Colorado, Boulder, 2015).
- 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**, 321–326 (2014).
- M. Aspelmeyer, T. J. Kippenberg, and F. Marquardt, “Cavity optomechanics”, *Reviews of Modern Physics* **86**, 1391–1452 (2014).
- C. J. Axline, L. D. Burkhardt, W. Pfaff, M. Zhang, K. Chou, P. Campagne-Ibarcq, P. Reinhold, L. Frunzio, S. M. Girvin, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, “On-demand quantum state transfer and entanglement between remote microwave cavity memories”, *Nature Physics* **14**, 705–710 (2018).
- K. C. Balram, M. I. Davanço, J. D. Song, and K. Srinivasan, “Coherent coupling between radiofrequency, optical and acoustic waves in piezo-optomechanical circuits”, *Nature Photonics* **10**, 346–352 (2016).
- S. D. Barrett and P. Kok, “Efficient high-fidelity quantum computation using matter qubits and linear optics”, *Physical Review A* **71**, 060310 (2005).
- 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).
- J. Bechhoefer, “Feedback for physicists: a tutorial essay on control”, *Reviews of Modern Physics* **77**, Publisher: American Physical Society, 783–836 (2005).

- C. H. Bennett and G. Brassard, “Quantum cryptography: public key distribution and coin tossing”, in Proceedings of IEEE international conference on computers, systems, and signal processing, event-place: Bangalore (1984), p. 175.
- M. Bild, M. Fadel, Y. Yang, U. von Lüpke, P. Martin, A. Bruno, and Y. Chu, “Schrödinger cat states of a 16-microgram mechanical oscillator”, *Science* **380**, 274–278 (2023).
- E. D. Black, “An introduction to pound–drever–hall laser frequency stabilization”, *American Journal of Physics* **69**, 79–87 (2000).
- A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, “Cavity quantum electrodynamics for superconducting electrical circuits: an architecture for quantum computation”, *Physical Review A* **69**, 062320 (2004).
- BlueFors, “Helium Battery”, <https://bluefors.com/products/cryogenic-measurement-system-options/helium-battery/>
- J. Bochmann, A. Vainsencher, D. D. Awschalom, and A. N. Cleland, “Nanomechanical coupling between microwave and optical photons”, *Nature Physics* **9**, 712–716 (2013).
- K. Børkje, F. Massel, and J. G. E. Harris, “Nonclassical photon statistics in two-tone continuously driven optomechanics”, *Physical Review A* **104**, 063507 (2021).
- B. M. Brubaker, J. M. Kindem, M. D. Urmev, 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).
- P. Burns, “Reducing added noise in a microwave-mechanical-optical converter”, Publication Title: Department of Physics Volume: Ph.D., PhD thesis (University of Colorado Boulder, Boulder, Dec. 2019).
- C. Cabrillo, J. I. Cirac, P. García-Fernández, and P. Zoller, “Creation of entangled states of distant atoms by interference”, *Physical Review A* **59**, 1025–1033 (1999).

- E. T. Campbell and S. C. Benjamin, “Measurement-based entanglement under conditions of extreme photon loss”, *Physical Review Letters* **101**, 130502 (2008).
- L. Catalini, Y. Tsaturyan, and A. Schliesser, “Soft-clamped phononic dimers for mechanical sensing and transduction”, *Physical Review Applied* **14**, 014041 (2020).
- C. M. Caves, “Quantum limits on noise in linear amplifiers”, *Physical Review D* **26**, 1817–1839 (1982).
- J. Chan, T. P. M. Alegre, A. H. Safavi-Naeini, J. T. Hill, A. Krause, S. Gröblacher, M. Aspelmeyer, and O. Painter, “Laser cooling of a nanomechanical oscillator into its quantum ground state”, *Nature* **478**, 89–92 (2011).
- S. Clemmen, A. Farsi, S. Ramelow, and A. L. Gaeta, “Ramsey interference with single photons”, *Physical Review Letters* **117**, 223601 (2016).
- J. P. Covey, A. Sipahigil, and M. Saffman, “Microwave-to-optical conversion via four-wave mixing in a cold ytterbium ensemble”, *Physical Review A* **100**, 012307 (2019).
- R. D. Delaney, M. D. Urmev, 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**, 489–493 (2022).
- R. Delaney, “Superconducting qubit readout via low-backaction electro-optomechanical transduction”, PhD thesis (University of Colorado Boulder, Boulder, CO, Mar. 2022).
- H. M. Doleman, T. Schatteburg, R. Benevides, S. Vollenweider, D. Macri, and Y. Chu, “Brillouin optomechanics in the quantum ground state”, *Physical Review Research* **5**, 043140 (2023).
- L.-M. Duan and H. J. Kimble, “Efficient engineering of multiatom entanglement through single-photon detections”, *Physical Review Letters* **90**, 253601 (2003).
- L.-M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, “Long-distance quantum communication with atomic ensembles and linear optics”, *Nature* **414**, 413–418 (2001).

- M. Eichenfield, J. Chan, R. M. Camacho, K. J. Vahala, and O. Painter, “Optomechanical crystals”, *Nature* **462**, 78–82 (2009).
- L. Fan, C.-L. Zou, R. Cheng, X. Guo, X. Han, Z. Gong, S. Wang, and H. X. Tang, “Superconducting cavity electro-optics: a platform for coherent photon conversion between superconducting and photonic circuits”, *Science Advances* **4**, eaar4994 (2018).
- M. Forsch, R. Stockill, A. Wallucks, I. Marinković, C. Gärtner, R. A. Norte, F. van Otten, A. Fiore, K. Srinivasan, and S. Gröblacher, “Microwave-to-optics conversion using a mechanical oscillator in its quantum ground state”, *Nature Physics* **16**, 69–74 (2020).
- W. Fu, M. Xu, X. Liu, C.-L. Zou, C. Zhong, X. Han, M. Shen, Y. Xu, R. Cheng, S. Wang, L. Jiang, and H. X. Tang, “Cavity electro-optic circuit for microwave-to-optical conversion in the quantum ground state”, *Physical Review A* **103**, 053504 (2021).
- I. Galinskiy, G.ENZIAN, M. Parniak, and E. Polzik, “Non-classical correlations between phonons and photons in a MHz-frequency mechanical oscillator coupled to an optical cavity”, arXiv:2312.05641
- I. Galinskiy, Y. Tsaturyan, M. Parniak, and E. S. Polzik, “Phonon counting thermometry of an ultracoherent membrane resonator near its motional ground state”, *Optica* **7**, 718 (2020).
- 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 times in on-chip superconducting quantum memories by optimizing materials, processes, and circuit design”, arXiv:2308.15539
- J. Gao, J. Zmuidzinas, B. A. Mazin, H. G. LeDuc, and P. K. Day, “Noise properties of superconducting coplanar waveguide microwave resonators”, *Applied Physics Letters* **90**, 102507 (2007).
- C. Gerry and P. Knight, **Introductory quantum optics** (Cambridge University Press, Cambridge, 2004).

- A. H. Ghadimi, S. A. Fedorov, N. J. Engelsen, M. J. Bereyhi, R. Schilling, D. J. Wilson, and T. J. Kippenberg, “Elastic strain engineering for ultralow mechanical dissipation”, *Science* **360**, 764–768 (2018).
- D. Gottesman, T. Jennewein, and S. Croke, “Longer-baseline telescopes using quantum repeaters”, *Physical Review Letters* **109**, 070503 (2012).
- J. Grebel, H. Yan, M.-H. Chou, G. Andersson, C. R. Conner, Y. J. Joshi, J. M. Miller, R. G. Povey, H. Qiao, X. Wu, and A. N. Cleland, “Bidirectional multiphoton communication between remote superconducting nodes”, *Physical Review Letters* **132**, Publisher: American Physical Society, 047001 (2024).
- M. Hafezi, Z. Kim, S. L. Rolston, L. A. Orozco, B. L. Lev, and J. M. Taylor, “Atomic interface between microwave and optical photons”, *Physical Review A* **85**, 020302 (2012).
- X. Han, W. Fu, C. Zhong, C.-L. Zou, Y. Xu, A. A. Sayem, M. Xu, S. Wang, R. Cheng, L. Jiang, and H. X. Tang, “Cavity piezo-mechanics for superconducting-nanophotonic quantum interface”, *Nature Communications* **11**, 3237 (2020).
- 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**, 178–181 (2013).
- W. Hease, A. Rueda, R. Sahu, M. Wulf, G. Arnold, H. G. Schwefel, and J. M. Fink, “Bidirectional electro-optic wavelength conversion in the quantum ground state”, *PRX Quantum* **1**, 020315 (2020).
- 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**, 1038–1042 (2018).

- 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).
- S. Hong, R. Riedinger, I. Marinković, A. Wallucks, S. G. Hofer, R. A. Norte, M. Aspelmeyer, and S. Gröblacher, “Hanbury brown and twiss interferometry of single phonons from an optomechanical resonator”, *Science* **358**, 203–206 (2017).
- G. Huang, A. Beccari, N. J. Engelsen, and T. J. Kippenberg, “Room-temperature quantum optomechanics using an ultralow noise cavity”, *Nature* **626**, Publisher: Nature Publishing Group, 512–516 (2024).
- W.-Y. Hwang, “Quantum key distribution with high loss: toward global secure communication”, *Physical Review Letters* **91**, 057901 (2003).
- V. S. Ilchenko, A. A. Savchenkov, A. B. Matsko, and L. Maleki, “Whispering-gallery-mode electro-optic modulator and photonic microwave receiver”, *JOSA B* **20**, 333–342 (2003).
- A. Imamoglu, “Cavity QED based on collective magnetic dipole coupling: spin ensembles as hybrid two-level systems”, *Physical Review Letters* **102**, 083602 (2009).
- 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 Si_3N_4 membrane and classical laser noise”, *New Journal of Physics* **14**, 115018 (2012).
- W. Jiang, C. J. Sarabalis, Y. D. Dahmani, R. N. Patel, F. M. Mayor, T. P. McKenna, R. Van Laer, and A. H. Safavi-Naeini, “Efficient bidirectional piezo-optomechanical transduction between microwave and optical frequency”, *Nature Communications* **11**, 1166 (2020).
- 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**, 965–970 (2022).

- N. Kalb, A. A. Reiserer, P. C. Humphreys, J. J. W. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham, and R. Hanson, “Entanglement distillation between solid-state quantum network nodes”, *Science* **356**, 928–932 (2017).
- N. S. Kampel, R. W. Peterson, R. Fischer, P.-L. Yu, K. Cicak, R. W. Simmonds, K. W. Lehnert, and C. A. Regal, “Improving broadband displacement detection with quantum correlations”, *Physical Review X* **7**, 021008 (2017).
- J. M. Kindem, “Quantum nanophotonics with ytterbium in yttrium orthovanadate”, PhD thesis (California Institute of Technology, 2019).
- Y. Kubo, F. R. Ong, P. Bertet, D. Vion, V. Jacques, D. Zheng, A. Dréau, J.-F. Roch, A. Auffeves, F. Jelezko, J. Wrachtrup, M. F. Barthe, P. Bergonzo, and D. Esteve, “Strong coupling of a spin ensemble to a superconducting resonator”, *Physical Review Letters* **105**, 140502 (2010).
- A. G. Kuhn, J. Teissier, L. Neuhaus, S. Zerkani, E. van Brackel, S. Deléglise, T. Briant, P.-F. Cohadon, A. Heidmann, C. Michel, L. Pinard, V. Dolique, R. Flaminio, R. Taïbi, C. Chartier, and O. Le Traon, “Free-space cavity optomechanics in a cryogenic environment”, *Applied Physics Letters* **104**, 044102 (2014).
- 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**, 614–619 (2023).
- K. C. Lee, M. R. Sprague, B. J. Sussman, J. Nunn, N. K. Langford, X.-M. Jin, T. Champion, P. Michelberger, K. F. Reim, D. England, D. Jaksch, and I. A. Walmsley, “Entangling macroscopic diamonds at room temperature”, *Science* **334**, 1253–1256 (2011).
- M. M. Leivo and J. P. Pekola, “Thermal characteristics of silicon nitride membranes at sub-kelvin temperatures”, *Applied Physics Letters* **72**, 1305–1307 (1998).
- S.-K. Liao, W.-Q. Cai, J. Handsteiner, B. Liu, J. Yin, L. Zhang, D. Rauch, M. Fink, J.-G. Ren, W.-Y. Liu, Y. Li, Q. Shen, Y. Cao, F.-Z. Li, J.-F. Wang, Y.-M. Huang, L. Deng, T. Xi, L.

- Ma, T. Hu, L. Li, N.-L. Liu, F. Koidl, P. Wang, Y.-A. Chen, X.-B. Wang, M. Steindorfer, G. Kirchner, C.-Y. Lu, R. Shu, R. Ursin, T. Scheidl, C.-Z. Peng, J.-Y. Wang, A. Zeilinger, and J.-W. Pan, “Satellite-relayed intercontinental quantum network”, *Physical Review Letters* **120**, 030501 (2018).
- R. Loudon, **The quantum theory of light** (Clarendon Press, 1973).
- G. G. Macfarlane, T. P. McLean, J. E. Quarrington, and V. Roberts, “Fine structure in the absorption-edge spectrum of si”, *Physical Review* **111**, 1245–1254 (1958).
- 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).
- F. Marsili, V. B. Verma, J. A. Stern, S. Harrington, A. E. Lita, T. Gerrits, I. Vayshenker, B. Baek, M. D. Shaw, R. P. Mirin, and S. W. Nam, “Detecting single infrared photons with 93% system efficiency”, *Nature Photonics* **7**, 210–214 (2013).
- S. Meesala, D. Lake, S. Wood, P. Chiappina, C. Zhong, A. D. Beyer, M. D. Shaw, L. Jiang, and O. Painter, “Quantum entanglement between optical and microwave photonic qubits”, *arXiv:2312.13559*
- 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*, Publisher: Nature Publishing Group, 1–7 (2024).
- M. Mirhosseini, A. Sipahigil, M. Kalaei, and O. Painter, “Superconducting qubit to optical photon transduction”, *Nature* **588**, 599–603 (2020).
- S. Mittal, K. Adachi, N. E. Frattini, M. D. Urmey, S.-X. Lin, A. E. Emser, C. Metzger, L. Talamo, S. Dickson, D. Carlson, S. B. Papp, C. A. Regal, and K. W. Lehnert, “Annealing reduces Si₃N₄ microwave-frequency dielectric loss in superconducting resonators”, *arXiv:2312.13504*

- D. L. Moehring, M. J. Madsen, K. C. Younge, J. R. N. Kohn, P. Maunz, L.-M. Duan, C. Monroe, and B. B. Blinov, “Quantum networking with photons and trapped atoms (invited)”, *JOSA B* **24**, 300–315 (2007).
- A. Narla, S. Shankar, M. Hatridge, Z. Leghtas, K. M. Sliwa, E. Zalys-Geller, S. O. Mundhada, W. Pfaff, L. Frunzio, R. J. Schoelkopf, and M. H. Devoret, “Robust concurrent remote entanglement between two superconducting qubits”, *Physical Review X* **6**, 031036 (2016).
- M. A. Page, M. Goryachev, H. Miao, Y. Chen, Y. Ma, D. Mason, M. Rossi, C. D. Blair, L. Ju, D. G. Blair, A. Schliesser, M. E. Tobar, and C. Zhao, “Gravitational wave detectors with broadband high frequency sensitivity”, *Communications Physics* **4**, 1–8 (2021).
- M. Parniak, I. Galinskiy, T. Zewtler, and E. S. Polzik, “High-frequency broadband laser phase noise cancellation using a delay line”, *Optics Express* **29**, 6935–6946 (2021).
- R. W. Peterson, T. P. Purdy, N. S. Kampel, R. W. Andrews, P.-L. Yu, K. W. Lehnert, and C. A. Regal, “Laser cooling of a micromechanical membrane to the quantum backaction limit”, *Physical Review Letters* **116**, 063601 (2016).
- W. Pfaff, C. J. Axline, L. D. Burkhardt, U. Vool, P. Reinhold, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, “Controlled release of multiphoton quantum states from a microwave cavity memory”, *Nature Physics* **13**, 882–887 (2017).
- A. P. M. Place, L. V. H. Rodgers, P. Mundada, B. M. Smitham, M. Fitzpatrick, Z. Leng, A. Premkumar, J. Bryon, A. Vrajitoarea, S. Sussman, G. Cheng, T. Madhavan, H. K. Babla, X. H. Le, Y. Gang, B. Jäck, A. Gyenis, N. Yao, R. J. Cava, N. P. de Leon, and A. A. Houck, “New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds”, *Nature Communications* **12**, 1779 (2021).
- 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 (June 13, 2022), QM4B.7.

- E. Planz, X. Xi, T. Capelle, E. C. Langman, and A. Schliesser, “Membrane-in-the-middle optomechanics with a soft-clamped membrane at milliKelvin temperatures”, *Optics Express* **31**, 41773 (2023).
- M. Pompili, S. L. N. Hermans, S. Baier, H. K. C. Beukers, P. C. Humphreys, R. N. Schouten, R. F. L. Vermeulen, M. J. Tiggelman, L. dos Santos Martins, B. Dirkse, S. Wehner, and R. Hanson, “Realization of a multinode quantum network of remote solid-state qubits”, *Science* **372**, 259–264 (2021).
- I. M. Pop, K. Geerlings, G. Catelani, R. J. Schoelkopf, L. I. Glazman, and M. H. Devoret, “Coherent suppression of electromagnetic dissipation due to superconducting quasiparticles”, *Nature* **508**, 369–372 (2014).
- T. P. Purdy, R. W. Peterson, P.-L. Yu, and C. A. Regal, “Cavity optomechanics with Si_3N_4 membranes at cryogenic temperatures”, *New Journal of Physics* **14**, 115021 (2012).
- C. L. Rau, A. Kyle, A. Kwiatkowski, E. Shojaei, J. D. Teufel, K. W. Lehnert, and T. Dennis, “Entanglement thresholds of doubly parametric quantum transducers”, *Physical Review Applied* **17**, 044057 (2022).
- 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).
- C. A. Regal and K. W. Lehnert, “From cavity electromechanics to cavity optomechanics”, *Journal of Physics: Conference Series* **264**, 012025 (2011).
- R. Riedinger, A. Wallucks, I. Marinković, C. Löschner, M. Aspelmeyer, S. Hong, and S. Gröblacher, “Remote quantum entanglement between two micromechanical oscillators”, *Nature* **556**, 473–477 (2018).

- J. M. Robinson, E. Oelker, W. R. Milner, D. Kedar, W. Zhang, T. Legero, D. G. Matei, S. Häfner, F. Riehle, U. Sterr, and J. Ye, “Thermal noise and mechanical loss of $\text{SiO}_2/\text{Ta}_2\text{O}_5$ optical coatings at cryogenic temperatures”, *Optics Letters* **46**, 592–595 (2021).
- A. Rueda, F. Sedlmeir, M. C. Collodo, U. Vogl, B. Stiller, G. Schunk, D. V. Strekalov, C. Marquardt, J. M. Fink, O. Painter, G. Leuchs, and H. G. L. Schwefel, “Efficient microwave to optical photon conversion: an electro-optical realization”, *Optica* **3**, 597–604 (2016).
- 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).
- R. Sahu, L. Qiu, W. Hease, G. Arnold, Y. Minoguchi, P. Rabl, and J. M. Fink, “Entangling microwaves with light”, *Science* **380**, 718–721 (2023).
- 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).
- A. Schliesser, R. Rivière, G. Anetsberger, O. Arcizet, and T. J. Kippenberg, “Resolved-sideband cooling of a micromechanical oscillator”, *Nature Physics* **4**, 415–419 (2008).
- 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).
- R. Shaniv, C. Reetz, and C. A. Regal, “Direct measurement of a spatially varying thermal bath using brownian motion”, *Physical Review Research* **5**, Publisher: American Physical Society, 043121 (2023).
- A. E. Siegman, **Lasers** (University Science Books, Mill Valley, Calif, 1986).
- R. Singh and T. P. Purdy, “Detecting acoustic blackbody radiation with an optomechanical antenna”, *Physical Review Letters* **125**, 120603 (2020).
- 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).

- R. Stockill, M. Forsch, F. Hijazi, G. Beaudoin, K. Pantzas, I. Sagnes, R. Braive, and S. Gröblacher, “Ultra-low-noise microwave to optics conversion in gallium phosphide”, *Nature Communications* **13**, 6583 (2022).
- M. Stone, “Hybrid cavity QED with rydberg atoms”, PhD thesis (University of Chicago, 2021).
- 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**, 265–270 (2023).
- D. V. Strekalov, H. G. L. Schwefel, A. A. Savchenkov, A. B. Matsko, L. J. Wang, and N. Yu, “Microwave whispering-gallery resonator for efficient optical up-conversion”, *Physical Review A* **80**, 033810 (2009).
- J. D. Teoh, P. Winkel, H. K. Babla, B. J. Chapman, J. Claes, S. J. de Graaf, J. W. O. Garmon, W. D. Kalfus, Y. Lu, A. Maiti, K. Sahay, N. Thakur, T. Tsunoda, S. H. Xue, L. Frunzio, S. M. Girvin, S. Puri, and R. J. Schoelkopf, “Dual-rail encoding with superconducting cavities”, *Proceedings of the National Academy of Sciences* **120**, e2221736120 (2023).
- 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**, 359–363 (2011).
- J. D. Thompson, B. M. Zwickl, A. M. Jayich, F. Marquardt, S. M. Girvin, and J. G. E. Harris, “Strong dispersive coupling of a high-finesse cavity to a micromechanical membrane”, *Nature* **452**, 72–75 (2008).
- Y. Tsaturyan, A. Barg, E. S. Polzik, and A. Schliesser, “Ultracoherent nanomechanical resonators via soft clamping and dissipation dilution”, *Nature Nanotechnology* **12**, 776–783 (2017).

- Y. Tsaturyan, A. Barg, A. Simonsen, L. G. Villanueva, S. Schmid, A. Schliesser, and E. S. Polzik, “Demonstration of suppressed phonon tunneling losses in phononic bandgap shielded membrane resonators for high-q optomechanics”, *Optics Express* **22**, 6810–6821 (2014).
- K. Uhlig, “Cryogen-free dilution refrigerators”, *Journal of Physics: Conference Series* **400**, 052039 (2012).
- K. Uhlig, “Dry dilution refrigerator with pulse tube shutoff option”, *Cryogenics* **130**, 103649 (2023).
- M. Underwood, D. Mason, D. Lee, H. Xu, L. Jiang, A. B. Shkarin, K. Børkje, S. M. Girvin, and J. G. E. Harris, “Measurement of the motional sidebands of a nanogram-scale oscillator in the quantum regime”, *Physical Review A* **92**, 061801 (2015).
- G. Ventura and L. Risegari, **The art of cryogenics** (Elsevier, 2008).
- 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).
- 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).
- C.-H. Wang, F. Li, and L. Jiang, “Quantum capacities of transducers”, *Nature Communications* **13**, 6698 (2022).
- L. A. Williamson, Y.-H. Chen, and J. J. Longdell, “Magneto-optic modulator with unit quantum efficiency”, *Physical Review Letters* **113**, 203601 (2014).
- 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).
- I. Wilson-Rae, N. Nooshi, W. Zwerger, and T. J. Kippenberg, “Theory of ground state cooling of a mechanical oscillator using dynamical backaction”, *Physical Review Letters* **99**, 093901 (2007).
- O. Wipfli, “Membrane mirror etalons for cavity optomechanics”, PhD thesis (ETH Zürich, 2015).

- C. Xiong, W. H. P. Pernice, and H. X. Tang, “Low-loss, silicon integrated, aluminum nitride photonic circuits and their use for electro-optic signal processing”, *Nano Letters* **12**, 3562–3568 (2012).
- Y. Xu, A. A. Sayem, L. Fan, C.-L. Zou, S. Wang, R. Cheng, W. Fu, L. Yang, M. Xu, and H. X. Tang, “Bidirectional interconversion of microwave and light with thin-film lithium niobate”, *Nature Communications* **12**, 4453 (2021).
- T. Yoon, D. Mason, V. Jain, Y. Chu, P. Kharel, W. H. Renninger, L. Collins, L. Frunzio, R. J. Schoelkopf, and P. T. Rakich, “Simultaneous brillouin and piezoelectric coupling to a high-frequency bulk acoustic resonator”, *Optica* **10**, 110–117 (2023).
- 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).
- P.-L. Yu, “Quantum optomechanics with engineered membrane resonators”, PhD thesis (University of Colorado, Boulder, 2015).
- E. Zeuthen, A. Schliesser, A. S. Sørensen, and J. M. Taylor, “Figures of merit for quantum transducers”, *Quantum Science and Technology* **5**, 034009 (2020).

Appendix A

Details of optical cavity assembly

A.1 Mounting the curved mirror

Prior to assembling the transducer cavity, the curved mirror must be mounted to the mirror holder. It can either be held by cryogenic-compatible epoxy in the geometry shown in Fig. 5.8, or clamped in place as shown in Fig. 6.2 and described in Sec. 6.1.4.

To accommodate our short optical cavity lengths of 1-2 mm, the mirrors are first “cored down” to a diameter of 4 mm using brass tubing and a diamond slurry, while protecting the mirror coating surface with Crystalbond mounting wax, a process pioneered by Hans Green.

Epoxying the mirror to the curved mirror holder is a process that involves two epoxy joints. During assembly, the mirror surface is protected with a layer of FirstContact optical cleaning polymer. The mirror substrate is first affixed to an Invar spacer that determines the length of the optical cavity. The mirror substrate is centered over the spacer by eye, and then clamped in place with a cotton swab, using approximately 5 lbs of force, as shown in Fig. A.1(a). The mirror is then epoxied in place with a thin and even fillet of Stycast 2850ft mixed with Catalyst 9. The mirror and spacer are then inverted and clamped against the curved mirror holder, as shown in Fig. A.1(c), using the spring clamp assembly shown in Fig. A.1(b). A fillet of Stycast 2850ft with Catalyst 9 is then applied to affix the Invar spacer to the curved mirror holder, as shown in Fig. A.1(d). Symmetry of the glue joints is important for maintaining alignment of the optical cavity mode, and reducing strain on the mirror that can lead to birefringence of the cavity mode or even crack the mirror substrate over repeated thermal cycling. The design shown in Fig. A.1 was a later




Figure A.1: **Epoxying the curved mirror.** (a) Curved mirror clamped to Invar spacer for gluing. A fillet of epoxy can be seen in black joining the mirror substrate to the Invar spacer. (b) Spring clamp assembly used to the clamp Invar spacer and mirror assembly to the curved mirror holder. The epoxied mirror assembly can be seen at the top of the image. (c) Spring clamp assembly holding the mirror-spacer assembly to the curved mirror holder. (d) Epoxy joint between the Invar spacer and the curved mirror holder can be seen in the hole in the curved mirror holder.

iteration that was more robust to thermal cycling, and glue joints were not observed to fail even after repeated thermal cycling in these later designs.

A.2 Cavity assembly setup

The setup used to assemble the transducer cavities is shown in Fig. A.2. During assembly, the cavity is oriented such that the optical axis is vertical. The chip holder supporting the electromechanical etalon is held in place, while the position of curved mirror holder can be adjusted to position the optical spot on the membrane. The vertical orientation of the cavity axis makes adjusting the cavity mode position easier, as the precise adjustments needed to modify the curved mirror position are not made more difficult by gravity in this orientation.

In assembling the cavity, two light sources are used to simultaneously image the cavity mode and membrane, in order to position the optical cavity mode relative to the membrane. A 940 nm LED is sufficiently out-of-band of the mirror coatings that it can effectively image the membrane. It is injected along the path indicated in orange, before combining with the primary beam path (red) on a polarizing beamsplitter (PBS). Lenses positioned as indicated on the LED path allow even illumination across the membrane.

The primary beam path can be easily steered to couple to the cavity's spatial mode with high efficiency. During assembly, a superluminescent diode (SLD) is used on the primary beam path to locate the optical axis of the cavity by imaging the light transmitted through the cavity. A 10 nm bandpass filter is used to limit the spectrum of the SLD to the operating bandwidth of the cavity's high finesse mirror coatings. A relay lens is positioned its focal length $f = 75$ mm away from the steering mirror, such that adjusting that steering mirror results in a pure translation of the beam at the location of the cavity mode, which is helpful when translating the cavity mode by moving the curved mirror and leaving the flat mirror fixed. The transmitted light is routed through a PBS to a camera for imaging and a photodetector to allow for transmitted power measurements.




Figure A.2: **Cavity assembly setup.** Left: Photograph of the cavity assembly setup, with the primary beam path illustrated in red, and the LED illumination path used for imaging the membrane illustrated in orange. Lens focal lengths are indicated. Right: Chip holder mounted to the assembly setup. The curved mirror holder is placed on top of the chip holder during alignment. In the photograph, the curved mirror holder is resting with the mirror surface facing upward, and is protected with a layer of FirstContact optical cleaning polymer (pink). The assembly setup orients the cavity optical axis vertically for ease of alignment.

A.3 Cavity assembly procedure

To assemble the optical cavity, first the photodetector is aligned to the beam of the SLD, and the location of the optical spot from the SLD on the camera's monitor is noted. The electromechanical etalon, held by the chip holder, is placed at the focus of the SLD, and its tilt is adjusted so that the retroreflection of the SLD overlaps the incident beam, ensuring that the beam is normal to the flat mirror. At the same time, the image of the membrane should be used to ensure the SLD is passing through the membrane. The curved mirror holder can then be put in place above the electromechanical etalon, and its position adjusted until SLD light is seen in transmission. The broadband SLD can excite many transverse modes simultaneously, and the transmitted light is symmetric about the optical axis of the cavity, making locating the cavity mode in the desired position simple using this tool. The position of the optical spot on the membrane can then be adjusted by translating the position of the curved mirror holder, and using the mirror at the focal plane of the relay lens to follow the movement of the cavity's optical axis. The curved mirror can then be fixed to the chip holder using screws.

A.4 Notes on measuring cavity linewidth in the assembly setup

After the optical spot seems to be situated in a suitable position, the cavity linewidth should be measured to ensure the optical cavity has sufficiently low loss. This is most easily done by sweeping the laser across resonance, while phase modulating the light with sidebands of a given frequency. In practice 8 to 15 MHz modulation is convenient for measuring cavities with linewidths between of 1-2 MHz. Fitting the three transmitted Lorentzians can then extract the cavity linewidth, with the sideband spacing calibrating the width of the Lorentzians.

Sweeping the laser across cavity resonance must be done sufficiently slowly, such that the transmitted light can reach a steady state to trace out the cavity's frequency response. If the laser is swept too quickly, the cavity will ring down during the measurement, broadening the transmission peak—this can also lead to interference effects explored in Oliver Wipfli's master's

thesis [Wipfli 2015]. On the other hand, the sweep must be done sufficiently quickly such that the laser's frequency can be considered stable, i.e. the sweep must be finished within the coherence time of the laser: $t_{\text{sweep}} < 2\pi/\kappa_l$, where $\kappa_l/2\pi$ is the linewidth of the laser. This constrains the sweep rate to approximately 10^{11} Hz/s for cavities with linewidths of approximately 1 MHz, and using sideband modulations frequencies of 5-10 MHz and a laser with a linewidth of 1 kHz. Air currents can cause frequency fluctuations of the cavity. A vacuum enclosure can be placed over the cavity in-situ to help get more precise measurements of the cavity linewidth. Evacuating the enclosure can help mitigate thermally-induced air movement during sweeping, though often in this setup blocking the ambient air currents with the enclosure is sufficient, and evacuation is often unnecessary.

Appendix B

Details of operating the optical detection setup

B.1 Detail on the optical measurement setup

Fig. B.1 shows the optical measurement setup used to measure transducers in the cryogen-free Bluefors LD400 dilution refrigerator. The measurement setup occupies three aluminum optical breadboards of 1/2" thickness mounted to the frame of the dilution refrigerator, with two 12" \times 12" breadboards fastened to the larger 36" \times 12" breadboard. The frequencies of the lock, pump, and LO beams are prepared on a separate optical table (not shown in Fig. B.1, see Fig. 5.5), and their frequencies are as defined in Sec. 5.5, where ω_m , Δ_B , and Δ_{LO} may vary depending on the device being measured. The lock, pump and local oscillator (LO) beams are coupled into PM980 polarization maintaining optical fiber, in order to transfer them to the measurement setup. Upon exiting the fiber, the three beams are each collimated before passing through a polarizing beam splitter (PBS) acting as a "clean up" cube, to ensure the propagating mode has linear polarization, as thermal drift of the fiber can lead to polarization drifts of the output beam. A half-wave plate ($\lambda/2$), second PBS and Thorlabs PDA10CF fixed-gain InGaAs photodetector (IS det) are used to intensity-stabilize each of the beams.

The lock beam uses an neutral density filter with optical density 1.0 (ND 1.0) to lower the power of the lock beam to reduce its optomechanical backaction on the transducer, while maintaining sufficient power on the intensity stabilization (IS) detector. A Faraday rotator (Thorlabs IO-5-1064-HP) allows the lock beam to propagate towards the transducer, to combine with the pump beam on a PBS, and, upon reflection, to be routed towards a Thorlabs APD 110C avalanche


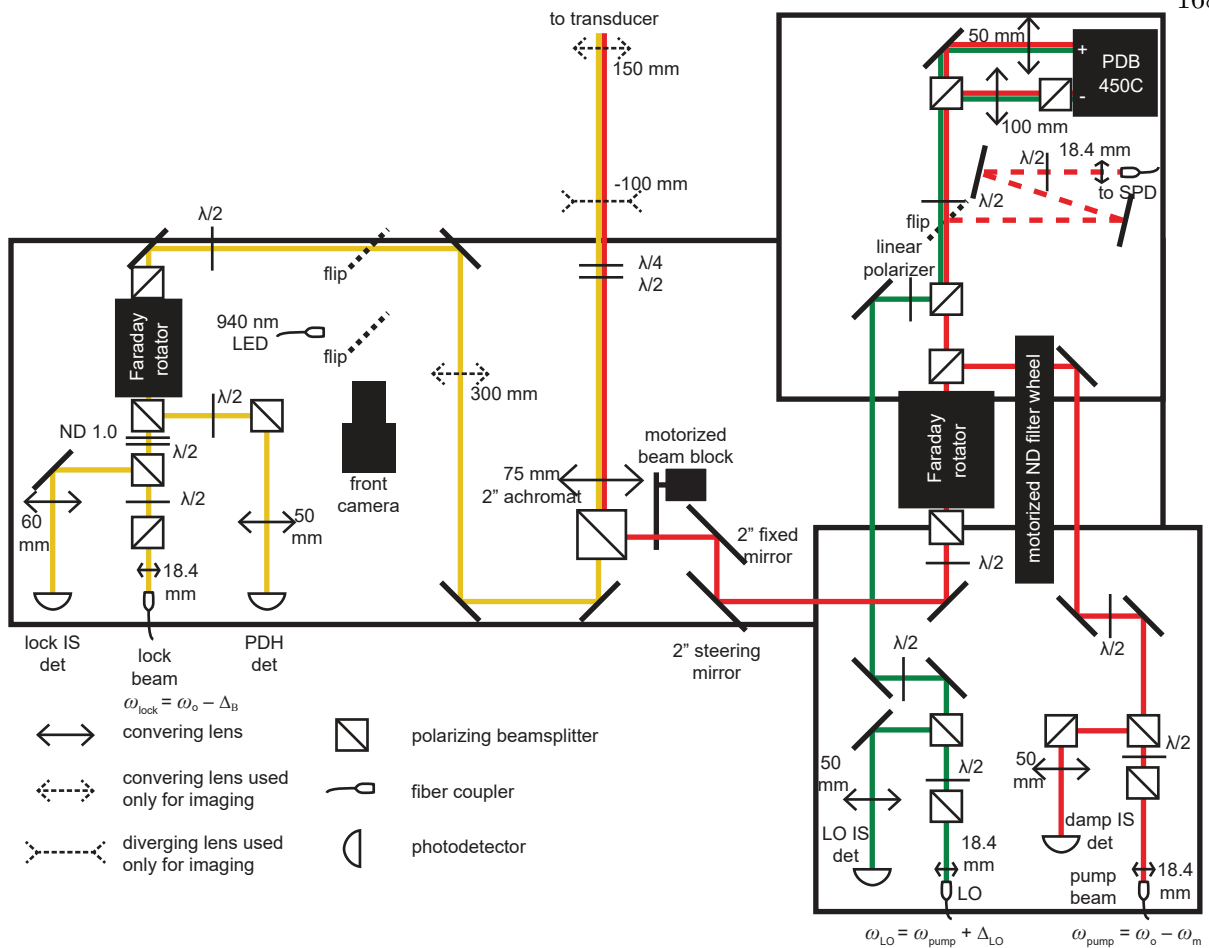


Figure B.1: **Detail on optical measurement setup.** Schematic (top) and photograph (bottom) of the optical measurement setup used with the Bluefors LD400 dilution refrigerator.

photodetector for frequency-locking the beams to the transducer cavity using the PDH technique. The sensitive avalanche photodetector is used here to allow the lock beam to be operated with low optical powers of tens of nW. A half-wave plate ($\lambda/2$) and quarter-wave plate ($\lambda/4$) allow polarization control of the pump and lock beams propagating towards the transducer (see Sec. B.2).

The path of the lock beam has additional optics used for imaging when the lock beam is not being used to stabilize the frequency of the laser to that of the optical cavity. These additional optics are inserted in the positions indicated by the dashed lens symbols in Fig. B.1. The “front camera” can be used with the flip mirror (flip) and an $f = 300$ mm lens inserted in the position indicated to image a “back beam” propagating in transmission from the back side of the transducer cavity (not shown in diagram). The back beam is used to measure the modematching between the cavity mode and the pump beam and LO beam, as described in Sec. 5.6.2. In order to image the membrane in transmission, a camera on the back-side of the transducer cavity (not shown) is used with 940 nm LED illumination from the front side of the transducer. To illuminate the membrane, the 940 nm LED is used with the two flip mirrors to direct the LED illumination along the lock path towards the transducer. The $f = 300$ mm, $f = -100$ mm, and $f = 150$ mm lenses are inserted in the positions indicated to focus the LED illumination in front of the cavity approximately at the focal plane of the “cold lens” (see Sec. 6.1.5) at the transducer input, in order to evenly illuminate the membrane with coarsely collimated light, over a maximum transverse extent.

The pump beam power can be controlled using a series of neutral density (ND) filters mounted on an Edmund Optics 84-899 USB Motorized Filter Wheel. Care must be taken to account for the effect of the ND filters on the modematching of the pump beam to the cavity and LO beam. An Electro Optics Technology (EOT) Faraday rotator with 12 mm aperture allows the pump beam to propagate towards the transducer, and upon reflection, be routed towards the heterodyne detection setup. A New Focus 8891 Motorized Flipper Mount is used to block the pump beam path for measurements of shot noise using the heterodyne detection setup. Importantly, the beam block in this position also blocks any leakage lock light that can lead to an artificially elevated shot-noise measurement. A manual flip mirror allows the reflected pump beam and optomechanical signal to

be routed to the single-photon detection setup (SPD).

Upon cryogenic cooling, the transducer cavity rises vertically by approximately 3 mm from thermal contraction of the dilution refrigerator, and optics that were aligned to the cavity when it was at room temperature must be realigned. To better accommodate this realignment, two larger 2"-diameter mirrors are used on the pump beam path close to the PBS on which the lock beam and pump beam are combined. A 2"-diameter achromatic $f = 750$ mm doublet is used as the primary focusing lens for the pump and lock beams, as this choice of lens results in less aberration when the beam is not centered, compared to a singlet lens. In order to efficiently couple into the cavity mode of waist $30\text{ }\mu\text{m}$ from a distance outside the dilution refrigerator, the beams propagate with a $1/e^2$ radius of 1.9 mm. These somewhat large beams are susceptible to clipping on smaller-diameter optics, so we use Faraday rotators designed for high-power applications because they have large apertures; the Thorlabs IO-5-1064-HP has a 4.7 mm aperture, and the EOT Faraday rotator has a 12 mm aperture.

The LO beam operates with a power of approximately 2 mW, compared with the tens-of- μW power levels of the pump beam. A linear polarizer reduces the leakage s-polarization of the LO beam from being reflected towards the heterodyne detection setup. A $\lambda/2$ and PBS route the LO and pump beams to the Thorlabs PDB450C InGaAs balanced detector. An additional PBS on the (-) input to the balance detector rejects unwanted s-polarization from the LO, which leads to a DC offset, and a measurement systematic in measuring the modematching factors as described in Sec. 5.6.2. An additional pitfall in these modematching measurements is that some of the Minicircuits low-pass filters have different insertion losses for DC and MHz-frequency signals, which can dramatically change the visibility measurements in a modematching measurement. I have found the 5-MHz BLP-5+ and 50-MHz BLP-50+ filters to faithfully recover visibilities of ~ 8 MHz-frequency signals, but all others I've tested do not.

B.2 Isolating the polarizations between the damp and lock

The lock and pump beams use orthogonal polarizations in order to prevent crosstalk between the beams, and to limit additional power in the signal bandwidth of the transducer due to the lock beam. Upon cooling the transducer to cryogenic temperatures, the cavity becomes birefringent from induced strain on the mirror coatings. The substrate of the input mirror also becomes birefringent. Even if the birefringent axis of the substrate is not aligned to the eigenpolarizations of the birefringent cavity, complete polarization control given by a half-wave plate (HWP) and a quarter-wave plate (QWP) is sufficient to satisfy two constraints simultaneously to ensure no crosstalk between the beams occurs upon reflection from the cavity, provided that the eigenpolarizations of the cavity are linear polarization states. The two constraints that must be satisfied are (1) that the input beams can be mapped to the cavity's orthogonal eigenpolarizations, and (2) that the prompt reflection from the mirror coating surface is returned to the input polarization state. The proof that these constraints can be simultaneously satisfied requires the eigenpolarizations of the cavity be to be linear, which I believe to be a good assumption for a Fabry-Pérot cavity whose birefringence is induced by strain on one or both of the cavity mirror coatings. If the eigenpolarizations are not linear, however, then control of the polarization state at the input of the dilution refrigerator is not sufficient to guarantee arbitrarily low crosstalk between the lock and pump beams.

Let p be a vector in the Jones polarization matrix formalism describing one eigenpolarization of the optical cavity, and let h be the polarization vector of the beam to address this pol, assumed to be linearly polarized in the horizontal direction, as indicated in Fig. B.2. Together, the HWP and QWP can realize an arbitrary polarization transformation described by the matrix W , and the strained mirror substrate results in the transformation described by the matrix S . The first constraint is then described by

$$p = SWh, \tag{B.1}$$

and the second constraint is described by

$$h = W^{-1}S^{-1}RSWh, \tag{B.2}$$




Figure B.2: **Jones matrices describing the desired polarization configuration of the experiment.** A input beam with polarization h is mapped to the cavity eigenpolarization p by transformations W due to the input half-wave plate (HWP) and quarter-wave plate (QWP) and S due to the mirror substrate.

where R is the matrix describing prompt reflection from the front mirror surface. Applying SW to both sides of Eq. B.2 and substituting Eq. B.1 into it twice, we see that the constraints can be simultaneously satisfied if

$$p = Rp, \quad (\text{B.3})$$

which is true if and only if p describes a linear polarization state.

The following procedure achieves the desired polarization separation between the lock and pump beams in practice, i.e. finding the solutions to Eqs. B.1 and B.2. The procedure assumes that the two beams are sufficiently aligned to the cavity such that their reflections will interfere on a photodetector, giving a signal at their difference frequency to minimize by manipulating the polarizations. In practice I use the science cavity lock detector for this optimization. (If the birefringence is such that beatnote occurs within the lock loops bandwidth of a few hundred kHz, this can lead to difficulties. The polarizations of the lock and damp can then be interchanged, and the beam detunings changed, as indicated in Fig. 5.6).

First the laser frequency is changed such that both beams are tuned away from cavity resonance by approximately 10 MHz. The polarizations of promptly reflected beams are then only affected by their double-pass through the mirror substrate. The QWP at the input of the fridge is removed from the setup. The HWP is rotated in order to minimize the beatnote between the beams as detected by the lock detector. The input QWP is replaced, and rotated to minimize the

beatnote on the detector. Eq. B.2 is now satisfied for each beam, with W resulting in two linear polarization states, such that the polarizations of the beams are now aligned to the fast and slow axes of the mirror substrate.

The laser is now locked to the cavity using the PDH technique. Generally, the eigenpolarizations of the cavity are not aligned to the birefringent axes of the mirror substrate, so the beatnote on the lock detector will increase. The HWP and QWP should now be walked together in order to find the solution to Eq. B.1, while ensuring that Eq. B.2 continue to be satisfied. During this optimization, it is possible for the PDH lock to disengage due to the changing amplitude of the lock beams returning to the photodetector, so care must be taken during this step.

B.3 Procedure for locating the cavity mode with the pump and lock beams

In the Bluefors LD400 system, cryogenic cooling translates the optical cavity up by approximately 3 mm. After cooling down, we have found that we can usually find cavity resonance using a Thorlabs SLD1050S-A60 1050-nm superluminescent diode (SLD) by purely tilting one of the steering mirrors purely upward. However, it is possible that does not lead to an observation of power transmitted through the optical cavity, and the prospect of being lost in a vast parameter space can induce some amount of fear. The following tools can be used to relocate the cavity's TEM_{00} mode, in order from more coarse and robust, to more refined though they require sufficient pre-alignment to be effective.

Spectrophotometer measurements we conducted of our mirror coatings deposited on silicon chips revealed a transmission maximum at approximately 955 nm, so one powerful tool is a Qphotonics QFBGLD-950-5BTF fiber-coupled laser diode with a center frequency of 955 nm. At this operating wavelength, the cavity is a low-finesse etalon, and minimally perturbs the transmitted light. The 955-nm laser can be used to maximize the power transmitted to coarsely align, and then two steering mirrors can be walked to improve the symmetry of the image of the beam transmitted through the low-finesse etalon formed by the cavity to refine the alignment. A 940-nm LED can also be used as a source of incoherent light to image the transducer's membrane, and to provide

an expectation of where the cavity mode will form, given the location of the optical spot during cavity assembly.

The 1050-nm SLD provides a broadband source of light in the wavelength band at which our cavity acts as a high-finesse optical cavity. Light is then injected into many cavity modes simultaneously, all of which are then imaged in transmission. We can optimize on the symmetry of these transmitted cavity modes to increase our coupling to the TEM₀₀ transverse modes of the optical cavity. Monitoring the retroreflected beam on the input side of the fridge, and steering the beam to ensure the input and reflected beams overlap spatially can be a more sensitive, though slower procedure for increasing the coupling to the TEM₀₀ cavity modes. Finally, the monochromatic light from the CTL can be used to optimize the modematching to a single TEM₀₀ mode by monitoring the transmitted power, and most sensitively by measuring the modematching between the cavity mode and the input beam, according to the procedure described the following section.

B.4 Procedure for optimizing modematching

When operating a new device, the heterodyne detection setup needs to be realigned. The following outlines the procedure for doing so.

To measure mode overlap between the LO and pump beams and the optical cavity, an auxiliary beam is passed through the cavity in transmission from the back back port of the optical cavity, so that the beams to be measured can be interfered on a photodetector and the modematching factors measured according to the procedure described in Sec. 5.6.2. The transmitted back beam polarization should be adjusted using the half-wave plate and quarter-wave plate located on the back side of the optical cavity, in order to match that of the pump beam, so the transmitted back beam is routed towards the heterodyne detection setup and interferes minimally with the cavity lock. The frequency of the back beam should be adjusted to ensure sufficient power is transmitted through the optical cavity. Assuming the pump beam is optimally detuned from the cavity by the mechanical frequency ω_m , the back beam should then be shifted up in frequency by $\omega_m - \Delta$ relative to the pump beam, where Δ is 100-200 kHz, to ensure the back beam does not contribute

net antidamping and cause the mechanical mode to ring up. It should be noted that because of the cavity birefringence, this is often not close in frequency to the lock beam. The cavity should then be locked, which requires that the lock and damp are reasonably well-aligned from optimizing brightness in transmission. Residual amplitude modulation (RAM) on the lock beam at the PDH sideband frequency can lead to slow fluctuations in the error signal level, and can cause the beam detunings to shift over time. Such detuning fluctuations can lead to the back beam fluctuating above the resonance frequency of the optical cavity, causing the mechanical mode to ring up and potentially lead to the PDH frequency lock failing, so the RAM should be minimized. At this point, the polarization isolation between the pump and lock beams should be optimized according to the procedure in Sec. B.2.

To coarsely align the back beam, it is convenient to use irises to locate the cavity mode in transmission from the front side, using light from the pump beam. A pair of irises can be placed and centered to symmetrically perturb the image of the cavity mode in transmission, giving two points that define the axis of the cavity mode. When aligning the back beam to the cavity mode, the most convenient indication transmitting power from the back beam the observation of it beating against the LO on the heterodyne detector if there is some amount of mode overlap between the cavity mode and the LO. If there is difficulty using this simple method, it is possible to use the camera on the front of the fridge, though this requires the cavity be unlocked to use the camera, and so use of an SLD is preferred as the frequency of the laser will not necessarily stay on resonance with the optical cavity.

Once the back beam is transmitted through the optical cavity, the modematching between the pump beam and cavity mode can be measured by interfering the two beams on the balance detector, and equalizing the power on a single port of the balance detector and blocking the unused port. The measured visibility of this beat note then gives $\sqrt{\epsilon_{\text{PC}}}$. The pump beam steering mirrors can then be walked to maximize this visibility. Because the transmitted cavity mode reflects from these mirrors on the path towards the heterodyne detection setup, ϵ_{PC} must be optimized before ϵ_{CL} . Next, ϵ_{CL} should be optimized by walking the LO steering mirrors. Finally, ϵ_{PL} should

be measured, but not optimized to. The modematching between the pump beam and the LO is important to measure to calibrate our transducer efficiency correctly, but it does not impact the transduction efficiency, unlike ϵ_{CL} and ϵ_{PC} , as described in Secs. 5.6.2 and 5.7.1.

In some cases it can be helpful to use a Shack-Hartmann wavefront sensor to compare the beam profile of the transmitted cavity mode with the LO and pump beams, in order to optimize their modematching factors with more nuanced metrics to feed back on.

B.5 Impact of modematching on reflection measurements made with optical heterodyne detection

Imperfect modematchings between the cavity mode and the pump and LO beams can lead to confounding effects in reflection measurements that present as a decreased overcoupling ratio, or equivalently as a Fano interferometric effect with a path with opposite phase. This effect became appreciable in our reflection measurements over the course of a battery run in the Bluefors LD400 system while running the helium battery, as the beams' modematching factors changed relative to one another over the course of the battery run as described in Sec. 6.1.3, resulting in time dependence of our $\mathcal{S}_{oo}(\omega)$ measurements. Here I derive the effects of differing modematching coefficients on $\mathcal{S}_{oo}(\omega)$ measurements. For simplicity, I'll consider a perfectly overcoupled cavity. I'll define the pump, cavity and LO modes respectively as unit vectors $|P\rangle$, $|C\rangle$, and $|L\rangle$ in a Hilbert space so that their modematching factors are given as the square of their inner product, e.g. $\epsilon_{PL} = |\langle L|P\rangle|^2$.

Far from cavity resonance, the pump beam will be promptly reflected, and the result of a measurement of $\mathcal{S}_{oo} = \langle L|P\rangle$ will simply result in $\sqrt{\epsilon_{PL}}$.

To evaluate the measured cavity response on resonance, the pump mode can be rewritten as

$$|P\rangle = \sqrt{\epsilon_{PC}} |C\rangle + [|P\rangle - \sqrt{\epsilon_{PC}} |C\rangle],$$

where the first term is the portion of the pump mode that is mode matched to the cavity, and the term in square brackets is the unmode-matched portion. Upon reflection from the cavity, the mode-

matched portion will be modified by the cavity response. In the case of a perfectly overcoupled optical on cavity resonance, the mode matched portion of the pump mode will acquire a π phase shift, and the reflected beam is described by:

$$\begin{aligned} |R\rangle &= -\sqrt{\epsilon_{\text{PC}}} |C\rangle + [|P\rangle - \sqrt{\epsilon_{\text{PC}}} |C\rangle] \\ &= -2\sqrt{\epsilon_{\text{PC}}} |C\rangle + |P\rangle \end{aligned}$$

And the result of a measurement of the scattering parameter is then

$$\mathcal{S}_{\text{oo}} = \langle L | R \rangle = -2\sqrt{\epsilon_{\text{PC}}\epsilon_{\text{CL}}} + \sqrt{\epsilon_{\text{CL}}}. \quad (\text{B.4})$$

So in the case of perfect modematchings, $\epsilon_{\text{PC}} = \epsilon_{\text{CL}} = \epsilon_{\text{PL}}$, we recover the expected behavior of a perfectly overcoupled cavity on resonance, the signal reflected with no lost power and a π phase shift. However, if ϵ_{PC} or ϵ_{CL} are reduced relative to ϵ_{CL} , the reflected signal is suppressed, leading to a greater dip in reflected power and an apparent reduction of the overcoupling ratio.

B.6 Transducer cavity lock loop

In this section, I briefly outline the transducer-cavity lock loop and its design considerations. For a helpful review of concepts from control theory used in this discussion, please see Ref. [Bechhoefer 2005].

As discussed in Sec. 5.5, the transducer cavity lock has two actuator channels: a “slow” channel active at frequencies up to a couple kHz, with a large throw at very low frequencies, and a “fast” channel active up to frequencies of around a few hundred kHz. The slow channel realizes a frequency change in the measurement beams by tuning the piezoelectric actuator in the filter cavity to which the laser is locked, and the fast channel shifts the frequency of the light using an acousto-optic modulator (see Fig. 5.5). In some of the earlier experiments described in this thesis, the slow channel used a piezo in the transducer cavity to tune its resonant frequency.

Instability of a control loop will occur if the phase delay of the loop approaches π at the frequency at which the closed-loop gain is one: the unity gain frequency (UGF). In choosing the




Figure B.3: **Simplified transducer cavity lock closed loop response.** The magnitude of the closed loop gain (top) and phase response (bottom) is shown for the slow channel, and the fast channel, and the combined sum of the two channels. The simplified model only includes the frequency-dependence of the control loop response, and does not include effects from piezo resonances, the phase delay of the fast channel response at higher frequencies, or additional filtering beyond the frequencies of loop response. The horizontal dashed line indicates unity gain. From left to right, the vertical dashed lines indicate the frequency at which the slow and fast channel closed loop gains are equal (approximately 70 Hz), the unity-gain frequency of the slow actuator channel (approximately 700 Hz), and the unity gain frequency of that fast actuator channel (approximately 320 kHz).

lock loop parameters, we have therefore ensured the phase response is sufficiently far from π by engineering first-order rolloff at the UGF of both channels.

An additional consideration is the relative phase difference between the two channels at the frequency at which they have equal gain. A π phase difference between them would lead to a cancellation of their responses, and suppress ability of the loop to respond to perturbations near this frequency. In order to prevent the two channels from fighting each other in this way, the difference in the order of the rolloff of each channel should not be greater than one.

There are physical limitations to the bandwidth of each channel. Measurement of the slow-channel response in an older device that used a piezo in the transducer cavity as the slow channel actuator revealed piezo resonances at frequencies of 8 kHz and higher. Resonances result in a π phase shift, so it is important that they fall below the UGF to avoid loop instability. We therefore choose the slow-channel UGF to be 2 kHz. Because we now feed back to the piezo controlling a mirror in the filter cavity at room temperature, it is possible that we can now be more aggressive with the slow-channel gain than we have been in the past. The fast channel is limited by the phase delay acquired in the closed loop at higher frequencies. Measurement of the fast-channel actuator response indicated that the phase delay begins to accumulate at hundreds of kHz. A phase-delay of $\pi/6$ at 300 kHz in addition to a $\pi/2$ -phase delay from first-order rolloff is sufficiently low to avoid instability if the UGF is chosen to be at this frequency or below.

Fig. B.3 shows the designed closed loop magnitude and phase responses that satisfy the constraints laid out above. The responses of both slow and fast channels is shown, as is that of their combined sum. Though the closed-loop gain is shown (and so includes the effect of the actuators) the model is simplified such that it does not contain frequency-dependence stemming from the actuators themselves, i.e. phase delay from the fast-channel actuator, or piezo resonances from the slow-channel actuator. The fast channel has a region of proportional gain between a high-pass filter with a cutoff frequency of 16 Hz, and a pole at 1.6 kHz initiating its first-order rolloff towards its UGF. The level of the proportional gain of 200 is set by the maximum gain permitted by the locking electronics without modifying them to incorporate an additional amplification stage.

The frequency of the pole at 1.6 kHz is then chosen to set the UGF to 300 kHz.

We employ a resistor in series with the capacitance of the filter cavity piezo to make a low-pass filter with a cutoff frequency of 70 Hz. The control loop additionally provides an integral response below a zero of frequency 700 Hz. The result is first-order rolloff in the slow channel at low and high frequencies, and a region of second-order rolloff between 70 Hz and 700 Hz, in order to maximize the gain of the slow channel at lower frequencies. From the UGF of 2 kHz set by the piezo resonances, the slow gain increases towards lower frequencies with first-order rolloff to ensure stability near the UGF, after which it increases to second-order rolloff at 700 Hz. This allows the loop gain to increase from approximately 4 at 700 Hz to approximately 200 at 70 Hz, where first-order rolloff continues down to DC due to the integral response of the control loop. The transition to first-order rolloff at 70 Hz reduces the phase difference between the fast and slow channels, though some suppression of total loop gain does occur due to the phase difference between them, as seen in the dip in the combined response at 70 Hz.

ProQuest Number: 31235296

INFORMATION TO ALL USERS

The quality and completeness of this reproduction is dependent on the quality and completeness of the copy made available to ProQuest.



Distributed by ProQuest LLC (2024).

Copyright of the Dissertation is held by the Author unless otherwise noted.

This work may be used in accordance with the terms of the Creative Commons license or other rights statement, as indicated in the copyright statement or in the metadata associated with this work. Unless otherwise specified in the copyright statement or the metadata, all rights are reserved by the copyright holder.

This work is protected against unauthorized copying under Title 17,
United States Code and other applicable copyright laws.

Microform Edition where available © ProQuest LLC. No reproduction or digitization of the Microform Edition is authorized without permission of ProQuest LLC.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 - 1346 USA