Cavity Optomechanics at Millikelvin Temperatures

Thesis by

Seán Michael Meenehan

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2015

(Defended March 9, 2015)

It never got weird enough for me.

- Hunter S. Thompson

Acknowledgments

To my advisor, Oskar Painter, who provided me with ample support and guidance throughout grad school. It has not been an easy trip. You know, a lotta ins, a lotta outs, a lotta what-have-yous. Still, somehow I have made it through and am stronger for it. Thank you.

To Justin Cohen, my lab partner and friend throughout grad school. I can't think of anyone I would have rather worked with for the last five years. This thesis literally would not have been possible without you.

To Oskar's first generation of students who helped build up the cleanroom and labs. So much of the work done by me and my contemporaries relied upon our ability to walk into a lab which already ran like clockwork. To Raviv in particular, for mentoring me during my initial months in the group. To Amir, for convincing me not to leave grad school after a couple of years. To Jeff and Simon, for supporting Justin and myself as we came up to speed on the silicon devices. To Jasper, for spearheading the initial development of optomechanical crystals and silicon nanobeams to which I have made only a modest contribution in this thesis. To Martin, for sage advice during my early forays into the world of quantum dots. To Tim, for being unreasonably friendly and optimistic when faced with my general crabbiness. To Alex, for science (obligatory). To Richard, for Korean BBQ and general good times. To Tao, for reminding us all that it's okay. To Ale for that sweet dart gun. To Carly and Thiago, for graciously putting up with my incessant drumming in the office. To Johannes for setting up the dilution fridge. To Matt Shaw and Francesco Marsili at JPL, for providing the amazing single photon detectors used in this work along with copious amounts of technical support. To Greg, who quickly came up to speed working under Justin and myself and who will (hopefully) continue this work into the near future.

To Michelle and Christy, without whom I am fairly certain the entire group/department would fall apart.

To Tom Donnelly for teaching me how to do research and convincing me to take a little time off before jumping into grad school.

To Aurora and DGro for providing more support and tolerance than I probably deserved. Your advice and friendship has pulled me through so much over the years. To Andy and Eric for virtuoso guitar and bass work. To Popebear and everyone who came to our concerts. To Josette for Mystikal

rapping. To Tom and Corina for unfailing kindness and generosity. To Matty Fay and Shaun Maguire for constantly reminding me how awesome science really is. To Reuben for serving as an inspiration and/or cautionary tale (I haven't decided yet). To Kenny for building fantastic High Sierra campsites and generally not putting up with any nonsense. To Paul, Tristan, Alice, Helge, and Hesham for beer hour. To Chris for making ice cream and baking amazing pies/pizza *even though you can't eat them.* To Branimir for being classier than everybody else. To Kristjan for movie nights and booking Popebear at Sausagefest. To Mo because SHEEEIIIT.

To my parents for the constant support they have provided throughout the ups and downs of grad school and life in general. To my sister Alex and her husband Dave for getting me out of lab every now and then. To my aunt Sarah and uncle Doug for giving me a second home in Silver Lake whenever necessary. To my grandparents for always being there to listen when needed. To the rest of my family, in Los Angeles and beyond.

To Emma, without whom I might not have made it through these last several years. I don't know what I did to deserve you or why/how you put up with me, but words fail to express how grateful I am to have found you.

Abstract

The field of cavity optomechanics, which concerns the coupling of a mechanical object's motion to the electromagnetic field of a high finesse cavity, allows for exquisitely sensitive measurements of mechanical motion, from large-scale gravitational wave detection to microscale accelerometers. Moreover, it provides a potential means to control and engineer the state of a macroscopic mechanical object at the quantum level, provided one can realize sufficiently strong interaction strengths relative to the ambient thermal noise. Recent experiments utilizing the optomechanical interaction to cool mechanical resonators to their motional quantum ground state allow for a variety of quantum engineering applications, including preparation of non-classical mechanical states and coherent optical to microwave conversion. Optomechanical crystals (OMCs), in which bandgaps for both optical and mechanical waves can be introduced through patterning of a material, provide one particularly attractive means for realizing strong interactions between high-frequency mechanical resonators and near-infrared light. Beyond the usual paradigm of cavity optomechanics involving isolated single mechanical elements, OMCs can also be fashioned into planar circuits for photons and phonons, and arrays of optomechanical elements can be interconnected via optical and acoustic waveguides. Such coupled OMC arrays have been proposed as a way to realize quantum optomechanical memories, nanomechanical circuits for continuous variable quantum information processing and phononic quantum networks, and as a platform for engineering and studying quantum many-body physics of optomechanical meta-materials.

However, while ground state occupancies (that is, average phonon occupancies less than one) have been achieved in OMC cavities utilizing laser cooling techniques, parasitic absorption and the concomitant degradation of the mechanical quality factor fundamentally limit this approach. On the other hand, the high mechanical frequency of these systems allows for the possibility of using a dilution refrigerator to simultaneously achieve low thermal occupancy and long mechanical coherence time by passively cooling the device to the millikelvin regime. This thesis describes efforts to realize the measurement of OMC cavities inside a dilution refrigerator, including the development of fridge-compatible optical coupling schemes and the characterization of the heating dynamics of the mechanical resonator at sub-kelvin temperatures.

We will begin by summarizing the theoretical framework used to describe cavity optomechanical

systems, as well as a handful of the quantum applications envisioned for such devices. Then, we will present background on the design of the nanobeam OMC cavities used for this work, along with details of the design and characterization of tapered fiber couplers for optical coupling inside the fridge. Finally, we will present measurements of the devices at fridge base temperatures of $T_{\mathrm{f}} = 10$ mK, using both heterodyne spectroscopy and time-resolved sideband photon counting, as well as detailed analysis of the prospects for future quantum applications based on the observed optically-induced heating.

Contents

List of Figures

Preface

I joined the Painter group somewhat late in my first year (March 2010, I think; at least that's when I started keeping a lab notebook). While there were a number of projects ongoing in the group at the time, Oskar suggested two in particular where he thought I might be a good fit. The first project, manned by Raviv Perahia and Justin Cohen, involved combining the so-called "zipper" optomechanical cavities with electrostatic actuators in active semiconductor material to make a tunable laser. The other option was to work with Amir Safavi-Naeini and Thiago Alegre, who at the time were attempting to measure an optomechanical analog of electromagnetically-induced transparency as a route to measuring Amir's new "snowflake" crystal design. I chose the former for two reasons. First, Raviv's response to my initial inquiries was something along the lines of "Here's when I'm in lab. Start coming by and we'll eventually find a use for you" whereas Amir's was "I guess I can give you some papers to read". Second, I thought "Hey, cool! Lasers!". I have continued to apply this deep level of consideration and introspection throughout graduate school and life in general.

The tunable laser project would eventually result in a nice little paper and then sort of fizzle out, while the EIT work would lead to a Nature paper and eventually roll into the other silicon optomechanical crystal work headed by Jasper Chan, which resulted in a number of landmark results in optomechanics, including one of the first demonstrations of ground state cooling in, the first observation of motional sideband asymmetry, the first demonstration of wavelength conversion, etc. In retrospect, maybe that was a questionable choice.

Still, the laser project had its upsides. In the several months in which we worked together Raviv proved an excellent mentor, helping me ease into working in the lab and generally functioning well within Oskar's group. This project would also introduce me to Justin, who would remain my constant lab partner throughout the next several years. In addition to starting to assist with actual measurements, Raviv tasked me with figuring out how to properly simulate the expected mechanical tuning of our structures including the intrinsic stress of the material, a task made difficult mostly by the general crumminess of COMSOL's documentation. This proved ideal, as it gave me something concrete on which I could work alone with a reasonable expectation of success and which played to my natural strengths (i.e., sitting in front of a computer and writing code). It also ended up setting

the tone for the working relationship Justin and I would have in the years to come, where I would focus on simulation and device design while Justin would work on device fabrication. This is also how I have come to possess the dubious distinction of having spent the least time in the cleanroom of all of Oskar's students (though not zero time, contrary to popular belief).

By the time the laser work was published, Raviv had moved on to the real world and Oskar suggested that Justin and I work on coupling self-assembled quantum dots in gallium arsenide optomechanical resonators. This built nicely upon the III-V fabrication experience Justin had acquired over the course of the laser work, and had the potential to give us a shortcut of sorts to quantum optomechanics by coupling the mechanical and optical resonators directly to a quantum two-level system. We set about designing, fabricating, and measuring electrostatically tunable cavities in quantum dot material, with the initial goal of enabling simple spectral alignment of an optical cavity with a particular quantum dot. At the same time, I began working on incorporating high-frequency mechanics into the design and attempting to calculate how strong we expected the quantum dot/phonon coupling to be.

Sadly, this project was impeded by a number of factors. For one, all other ongoing projects in the lab at that time involved silicon devices, which meant that by a large margin most of the device processing in our cleanroom was of silicon chips. Due to the finicky nature of reactive ion plasma etching, and the very different etch chemistries of silicon and gallium arsenide, this meant that fabricating good devices was pretty tricky (real talk for anyone not in the Painter group: when I say "good" I mean optical quality factors of *at least* $40 - 50,000$). The best results we obtained came after monopolizing the cleanroom for a solid weekend doing nothing but etching devices, after which the last chip came out pretty well. Another difficulty was testing the devices. As the ongoing effort to cool silicon nanobeam OMCs into the ground state ramped up, it came to encompass all available optical tables across both of the group's labs. This meant that there were long stretches of time where we could not test devices for fear of ruining someone else's experiment.

After a fair amount of struggling on this project, we came to the conclusion that the numbers probably were never going to quite work out: the optical Qs were too low, the estimated coupling between the quantum dots and a high-frequency phonon was too weak, and the lifetimes and dephasing times we expected in our quantum dots (based on earlier measurements by Kartik Srinivasan) were too large. Moreover, we still hadn't really nailed down exactly what we wanted to with such a system. Eventually we made the painful decision to cut our losses, chalk the project up as a dead-end, and move on. Looking back, I think part of the problem was the relative inexperience of Justin and myself. As the earlier experiment with Raviv involved free-space photoluminescence measurements, we started this project with zero experience in many of the basic techniques used in the lab, such as optical fiber taper testing, heterodyne mechanical spectroscopy, pulling a new fiber taper, etc. Over the course of the project we had to learn most of this on our own with relatively

little guidance (typically just enough not to break anything) due to the focus on the ground state experiment, which led to a lot of fumbling about on our part. This was extremely valuable in some sense, as one tends to learn quite a bit in the course of failing, but ended up stretching the project on much longer than was probably advisable.

Around this time the silicon nanobeam experiments had really started to pay off, and the next major goal was to measure the devices inside a dilution refrigerator. At the time Jasper and Simon Gröblacher (and later Amir) had already begun testing free-space couplers as a potential route to optically coupling to OMCs inside a fridge, with mixed results. Oskar suggested that Justin and I work on a different type of coupler involving direct coupling of a cleaved fiber to a tapered on-chip waveguide (the so-called "butt couplers" which are described in this thesis). As we hoped to avoid some of the issues related to two-photon absorption in silicon, this project initially focused on using silicon nitride. While we would eventually move back to silicon due to the much larger achievable optomechanical coupling strengths, the initial coupler work would eventually prove useful for coupling to optomechanical accelerometers developed by Alex Krause and Tim Blasius. In addition, the high stress of the nitride gave us the ability to fabricate extremely long, fully suspended, efficiently coupled waveguides. This would lead to a ongoing collaboration with Jeff Kimble's group focused on trapping atoms in the near-field of the optical photonic crystal waveguide. My involvement in this project was mostly on the front end, helping to create the initial designs for the trapping waveguides and calculating the relevant Casimir-Polder forces on the atoms.

Eventually, Oskar purchased a dilution fridge and our attention turned back to the measurement of devices in the fridge. As most of the original silicon nanobeam team was gone or leaving shortly, Justin and I started to take ownership of the project, assisted heavily in the beginning by Simon, Jeff, and Amir. This led to the work which forms the bulk of this thesis, as well as other measurements detailed in Justin's thesis.

Chapter 1

Fundamentals of Cavity Optomechanics and Optomechanical Crystals

This chapter will present the basic theoretical framework used to describe the optomechanical interaction and the measurement of mechanical motion within such systems, with a particular eye towards the sideband-resolved systems which are the focus of this thesis. Some of the proposed quantum applications for such systems will also be briefly summarized, as they serve to motivate the experiments presented in this work. Finally, the basic theory and design of the specific optomechanical devices studied in the Painter group (the so-called optomechanical crystals, or OMCs) will be presented. As much of this material has already been presented elsewhere, this chapter will aim to present enough detail to support the experimental work while trying to avoid excessive redundancy. In particular, the initial design of the OMC structures presented in section 1.3 was largely performed by previous Painter group students Jasper Chan and Amir Safavi-Naeini, and full details of the design process and associated theory can be found in their theses [1, 2] and related papers.

1.1 Cavity Optomechanics

The canonical model of a cavity optomechanical system is shown in Fig. 1.1. The optical cavity is taken to be a simple Fabry-Perot cavity, with optical resonance frequency ω_{c} and photon annihilation (creation) operator $\hat{a}$ ($\hat{a}^\dagger$), where one end mirror of the cavity is mechanically compliant (shown here as a massive mirror mounted on a spring) and behaves as a simple harmonic oscillator with resonant frequency ω_{m} and phonon annihilation (creation) operator $\hat{b}$ ($\hat{b}^\dagger$). In the absence of any interaction, the optical and mechanical resonators have the bare Hamiltonian

$$\hat{H} = \hbar\omega_{\mathrm{c}}\hat{a}^\dagger\hat{a} + \hbar\omega_{\mathrm{m}}\hat{b}^\dagger\hat{b}, \tag{1.1}$$

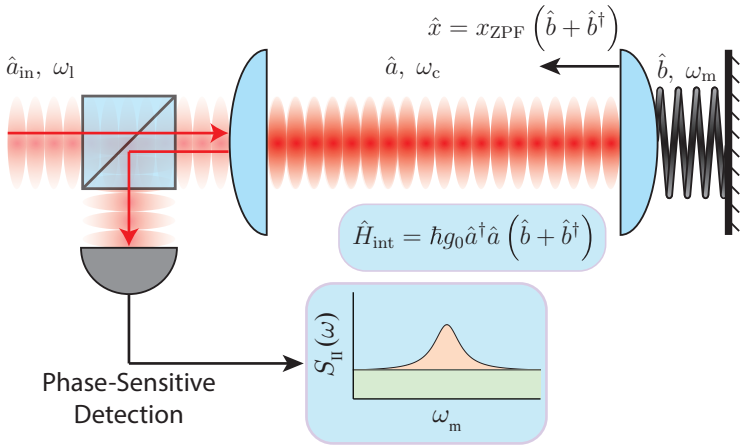

Figure 1.1: **Canonical cavity optomechanical system.** An optomechanical system may be modeled as a Fabry-Perot cavity with a moving end mirror, where the intracavity optical field (photon annihilation operator $\hat{a}$) couples to the mechanical oscillator (phonon annihilation operator $\hat{b}$) via the interaction Hamiltonian $\hat{H}_{\mathrm{int}} = \hbar g_0 \hat{a}^\dagger \hat{a}\left(\hat{b} + \hat{b}^\dagger\right)$. This produces phase modulation on an optical cavity drive (photon annihilation operator $\hat{a}_{\mathrm{in}}$), and phase sensitive detection of the cavity reflection may be used to measure the mechanical oscillation by studying the resulting photocurrent noise power spectral density S_{II}.

where $\hat{a}$ and $\hat{a}^\dagger$ obey the usual bosonic commutation relation $\left[\hat{a}, \hat{a}^\dagger\right] = 1$, and similarly for $\hat{b}$ and $\hat{b}^\dagger$. While the fundamental microscopic interaction between the optical and mechanical degrees of freedom can be rigorously derived from first principles for this model system [3], a more intuitive derivation of the interaction Hamiltonian comes from noting that the optical resonance frequency obeys the relation $\omega_{\mathrm{c}} = nc/2L$, $n \in \mathbb{Z}$, where c is the speed of light in vacuum and L is the total cavity length. As the length in turn depends on the position of the mechanical oscillator, $\hat{x}$, via the relation $L = L_0 + \hat{x}$, the optical frequency will be dependent upon $\hat{x}$. For small amounts of mechanical motion, we can make a Taylor expansion of $\omega_{\mathrm{c}}(\hat{x})$ about $\hat{x} = 0$ to first order and obtain

$$\omega_{\mathrm{c}}(\hat{x}) \approx \omega_{\mathrm{c}} + \frac{\omega_{\mathrm{c}}}{L_0}\hat{x}, \tag{1.2}$$

where ω_{c} is taken henceforth as the optical resonance frequency at the unperturbed length L_0 (i.e., $\omega_{\mathrm{c}}(0)$). Using the definition of canonical position in terms of phonon creation and annihilation operators $\hat{x} = x_{\mathrm{ZPF}}\left(\hat{b} + \hat{b}^\dagger\right)$, where x_{ZPF} is the zero-point amplitude of the mechanical oscillator, we may combine Eqs. 1.1 and 1.2 to obtain the lowest order interacting Hamiltonian of the optomechanical system

$$\hat{H} = \hbar\omega_{\mathrm{c}}\hat{a}^\dagger\hat{a} + \hbar\omega_{\mathrm{m}}\hat{b}^\dagger\hat{b} + \hbar g_0 \hat{a}^\dagger\hat{a}\left(\hat{b} + \hat{b}^\dagger\right), \tag{1.3}$$

where $g_0 = x_{\mathrm{ZPF}}\omega_{\mathrm{c}}/L_0$ is the fundamental optomechanical interaction rate, physically representing the optical frequency shift due to the zero-point motion of the mechanical oscillator. From this

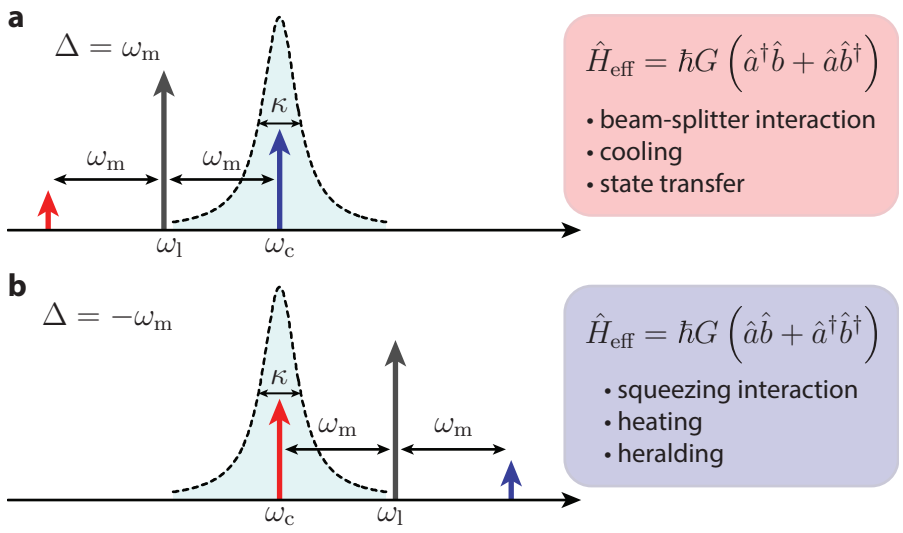

Figure 1.2: **Sideband-resolved optomechanics.** The optomechanical interaction scatters a coherent drive tone (gray arrows) to produce Stokes (red arrows) and anti-Stokes (blue arrows) sidebands corresponding to phonon emission and absorption, respectively. In the sideband resolved regime, the modified density of states provided by the optical cavity may be used to enhance one or the other scattering process. For red-detuned driving this results in an effective beam-splitter-like Hamiltonian which may be used for laser cooling or state transfer. For blue-detuned driving the effective Hamiltonian has the form of two-mode squeezing, useful for amplification of the mechanical motion or the heralded generation of non-classical mechanical states.

Hamiltonian it also follows that the cavity field exerts a radiation pressure force upon the mechanical oscillator $\hat{F}_{\rm rad} = \partial\hat{H}_{\rm int}/\partial\hat{x} \propto \hat{a}^\dagger\hat{a}$, where the per-photon force is equal to $\hbar g_0/x_{\rm ZPF}$. While we have explicitly considered a simple Fabry-Perot cavity, the Hamiltonian of Eq. 1.3 is much more general. Any system where the motion of a mechanical oscillator shifts the resonance frequency of an electromagnetic cavity will have an interaction of this form to lowest order, where $g_0 = x_{\rm ZPF}\partial\omega_c/\partial\hat{x}$.

As the optomechanical interaction dynamically modulates the cavity frequency (i.e., the phase), driving the cavity with an input laser at frequency ω_l will result in phase noise on the output correlated with the mechanical motion, oscillating at frequencies $\omega_l \pm \omega_m$. In general, directing this reflected output to some phase-sensitive detector (e.g., an optical heterodyne receiver) will then produce an output photocurrent $\hat{I}$ with an RF component oscillating at the mechanical frequency. By looking at the noise power spectral density (NPSD) $S_{II}(\omega)$ of this electronic signal, we can measure the properties of the mechanical oscillator via the optical output field. In particular, for the typical case of a mechanical oscillator driven by thermal noise, the NPSD, as illustrated in Fig. 1.1, will consist of a Lorentzian peak (pink shaded region) centered at ω_m, with a total noise power (i.e., area under the curve) related to the average phonon occupancy of the resonator, and a linewidth equal to the mechanical decay rate. This signal sits atop a noise floor (green shaded region), which in the ideal case is due solely to optical vacuum noise on the drive beam.

Beyond enabling precise measurement of the mechanical motion, the optomechanical interaction provides us with the tools necessary to control the state of the mechanics. This is most easily understood by considering the interaction in the frequency domain, as shown in Fig. 1.2. In the weak coupling regime, where the fundamental interaction rate g_0 is less than the cavity decay rate κ, the interaction between the drive laser (gray arrows, at frequency ω_l) and the mechanical oscillator will generate two weak optical sidebands at frequencies $\omega_\mathrm{l} \pm \omega_\mathrm{m}$. As discussed above, this can be understood classically as a weak phase modulation process, but can also be considered microscopically in terms of photon-phonon scattering. The upper sideband at $\omega_\mathrm{l} + \omega_\mathrm{m}$ (blue arrows), often called the anti-Stokes sideband in analogy with Raman scattering, is generated by the interaction term $\hat{a}^\dagger \hat{a} \hat{b}$ and represents the up-conversion of a drive photon accompanied by the absorption of a phonon from the mechanical resonator. Conversely, the lower (Stokes) sideband at $\omega_\mathrm{l} - \omega_\mathrm{m}$ (red arrows) is generated by the $\hat{a}^\dagger \hat{a} \hat{b}^\dagger$ interaction term and corresponds to down-conversion of drive photons accompanied by phonon emission.

In the absence of an optical cavity these two fundamental processes will occur with equal strength, and consequently there is no net effect on the mechanical resonator beyond the overall radiation pressure force discussed above. However, the presence of the cavity provides a strong modification to the density of optical states, represented in Fig. 1.2 by blue shaded Lorentzians. In the so-called "sideband-resolved" regime, when the sideband spacing is much larger than the optical cavity linewidth ($\omega_\mathrm{m} \gg \kappa$), this modified density of states may be used to enhance one or the other of these two basic scattering processes.

In particular, if the drive is placed at a lower frequency (red-detuned) from the cavity such that the anti-Stokes sideband is resonant ($\Delta = \omega_\mathrm{c} - \omega_\mathrm{l} = \omega_\mathrm{m}$), as shown in Fig. 1.2a, one can derive an effective linearized interaction in a rotating wave approximation (RWA) represented by the Hamiltonian $\hat{H}_\mathrm{eff} = \hbar G \left(\hat{a}^\dagger \hat{b} + \hat{a} \hat{b}^\dagger \right)$, where $G = \sqrt{n_\mathrm{c}} g_0$ (n_c is the number of intracavity drive photons) is the parametrically enhanced coupling rate. This beam-splitter-like effective interaction involves the coherent exchange of photons and phonons. The most straightforward application of this interaction is to cool the system by converting phonons to photons which are then extracted from the optical cavity. More generally, however, it may be used for state-transfer between the optical and mechanical systems. This enables, for example, coherent transfer between different optical wavelengths utilizing the mechanical resonator as an intermediary, as discussed further in section 1.2.1.

On the other hand, if the drive is blue-detuned from the cavity ($\Delta = -\omega_\mathrm{m}$) so that the Stokes sidebands is resonantly enhanced, as shown in Fig. 1.2b, one finds an effective RWA interaction $\hat{H}_\mathrm{eff} = \hbar G \left(\hat{a} \hat{b} + \hat{a}^\dagger \hat{b}^\dagger \right)$. This interaction, which has the form of two-mode squeezing, allows one to amplify the mechanical motion, as the process of phonon emission rather than absorption is now dominant. Beyond this, the creation and annihilation of photons and phonons in correlated pairs

allows one to herald the creation of non-classical states of the mechanical resonator, as discussed in section 1.2.2.

1.1.1 Heisenberg-Langevin Equations of Motion

A full theoretical treatment of the cavity optomechanical system starts from the basic equations of motion, working in the Heisenberg picture where the time evolution of an operator $\hat{A}$ is given by $\dot{\hat{A}} = \frac{i}{\hbar}\left[\hat{H}, \hat{A}\right] + \frac{\partial \hat{A}}{\partial t}$. While the optomechanical Hamiltonian given in Eq. 1.3 describes the internal dynamics of the system, it does not incorporate the effects of damping and noise due to the surrounding environment. Such effects can be introduced via the standard input-output formalism for open quantum cavities [4, 5], which is briefly summarized in Appendix A.2, leading to the full Heisenberg-Langevin equations for the photon and phonon annihilation operators:

$$\dot{\hat{a}} = -\left(i\omega_{\mathrm{c}} + \frac{\kappa}{2}\right)\hat{a} - ig_0\hat{a}\left(\hat{b} + \hat{b}^\dagger\right) + \sqrt{\kappa_{\mathrm{e}}}\hat{a}_{\mathrm{in}} + \sqrt{\kappa_{\mathrm{i}}}\hat{a}_{\mathrm{i}}, \tag{1.4}$$

$$\dot{\hat{b}} = -\left(i\omega_{\mathrm{m}} + \frac{\gamma_{\mathrm{i}}}{2}\right)\hat{b} - ig_0\hat{a}^\dagger\hat{a} + \sqrt{\gamma_{\mathrm{i}}}\hat{b}_{\mathrm{in}}, \tag{1.5}$$

where γ_{i} and $\kappa = \kappa_{\mathrm{e}} + \kappa_{\mathrm{i}}$ are the decay rates of the mechanical and optical resonators, respectively. The stochastic noise operator $\hat{b}_{\mathrm{in}}$ arises from the coupling between the mechanical oscillator and its surrounding environment, and obeys the commutation relation [1] $\left[\hat{b}_{\mathrm{in}}(t), \hat{b}^\dagger_{\mathrm{in}}(t')\right] = \delta(t - t')$. The coupling of the optical cavity to its environment, on the other hand, has been explicitly separated into two channels: the extrinsic channel, with coupling rate κ_{e} and noise operator $\hat{a}_{\mathrm{in}}$, and the intrinsic channel, with coupling rate κ_{i} and noise operator $\hat{a}_{\mathrm{i}}$. The extrinsic channel is the specific mode via which the cavity is probed, while the intrinsic channel represents all other environmental loss channels that go unmeasured (e.g., radiation of energy due to scattering, material absorption inside the cavity, etc.). Both optical noise operators obey the same type of commutation relation as $\hat{b}_{\mathrm{in}}$. The output field in the extrinsic channel is given by the input-output boundary condition

$$\hat{a}_{\mathrm{out}} = \hat{a}_{\mathrm{in}} - \sqrt{\kappa_{\mathrm{e}}}\hat{a}. \tag{1.6}$$

While $\hat{b}_{\mathrm{in}}$ and $\hat{a}_{\mathrm{i}}$ are pure noise variables, $\hat{a}_{\mathrm{in}}$ typically consists of both a classical tone (technically, a coherent state) at frequency ω_{l} in addition to stochastic noise. We may explicitly factor out both the harmonic time dependence of the operator as well as the classical amplitude α_{in} by making the substitution $\hat{a}_{\mathrm{in}} \to (\alpha_{\mathrm{in}} + \hat{a}_{\mathrm{in}})\, e^{-i\omega_{\mathrm{l}}t}$. As the optical drive frequency is much faster than the decay rates of the system and the mechanical frequency, it is convenient to move into a frame rotating at the drive frequency, focusing only on the slowly varying dynamics. This is accomplished by making

[1] This is equivalent to the assumption that the environmental coupling is strictly Markovian, so that the noisy environment effectively has no "memory" and the equations of motion only depend on the state of the system at time t. For further details see Appendix A.2 and Ref. [5].

the unitary transformation $\hat{a} \rightarrow \hat{U}^{\dagger}\hat{a}\hat{U}$, where $\hat{U} = e^{i\omega_{\mathrm{l}}\hat{a}^{\dagger}\hat{a}t}$, leading to the final equations of motion:

$$\dot{\hat{a}} = -\left(i\Delta + \frac{\kappa}{2}\right)\hat{a} - ig_0\hat{a}\left(\hat{b} + \hat{b}^{\dagger}\right) + \sqrt{\kappa_{\mathrm{e}}}\left(\alpha_{\mathrm{in}} + \hat{a}_{\mathrm{in}}\right) + \sqrt{\kappa_{\mathrm{i}}}\hat{a}_{\mathrm{i}}, \tag{1.7}$$

$$\dot{\hat{b}} = -\left(i\omega_{\mathrm{m}} + \frac{\gamma_{\mathrm{i}}}{2}\right)\hat{b} - ig_0\hat{a}^{\dagger}\hat{a} + \sqrt{\gamma_{\mathrm{i}}}\hat{b}_{\mathrm{in}}, \tag{1.8}$$

where $\Delta = \omega_{\mathrm{c}} - \omega_{\mathrm{l}}$ is the cavity-laser detuning.

Though we have made no assumptions so far as to the nature of the environmental noise, the mechanical bath is typically taken to be in thermal equilibrium at temperature T_{b}, such that the mechanical noise operator obeys the following correlation relations:

$$\begin{aligned}\left\langle \hat{b}_{\mathrm{in}}^{\dagger}(t)\hat{b}_{\mathrm{in}}(t')\right\rangle &= n_{\mathrm{b}}\delta(t-t'),\\ \left\langle \hat{b}_{\mathrm{in}}(t)\hat{b}_{\mathrm{in}}^{\dagger}(t')\right\rangle &= (n_{\mathrm{b}}+1)\,\delta(t-t'),\end{aligned} \tag{1.9}$$

where $n_{\mathrm{b}} = \left(e^{\hbar\omega_{\mathrm{m}}/k_{\mathrm{B}}T_{\mathrm{b}}} - 1\right)^{-1}$ is the average occupancy of the mechanical resonator in thermal equilibrium at temperature T_{b}, and all other one- and two-time correlations of $\hat{b}_{\mathrm{in}}$ or $\hat{b}_{\mathrm{in}}^{\dagger}$ are equal to zero [2]. The intrinsic optical bath is usually taken to be a thermal bath as well, and in the case of optical systems at room temperature the frequency is large enough ($\hbar\omega_{\mathrm{c}}$, $\hbar\omega_{\mathrm{l}} \gg k_{\mathrm{B}}T_{\mathrm{b}}$) that the bath may be approximated as zero-temperature (vacuum). Thus, the corresponding noise correlation functions are

$$\begin{aligned}\left\langle \hat{a}_{\mathrm{i}}^{\dagger}(t)\hat{a}_{\mathrm{i}}(t')\right\rangle &= 0,\\ \left\langle \hat{a}_{\mathrm{i}}(t)\hat{a}_{\mathrm{i}}^{\dagger}(t')\right\rangle &= \delta(t-t'),\end{aligned} \tag{1.10}$$

where, as in the case of the mechanical noise, all other one- and two-time correlations vanish [3]. The optical noise on the extrinsic channel, $\hat{a}_{\mathrm{in}}$, is typically also assumed to be vacuum noise, and thus has the same type of correlation functions shown in Eq. 1.10 (note that $\hat{a}_{\mathrm{in}}$ and $\hat{a}_{\mathrm{i}}$ belong to independent baths and are thus uncorrelated with each other). It is also possible for $\hat{a}_{\mathrm{in}}$ to have additional noise due to technical noise on the drive laser, which will be discussed further in section 1.1.6. The two-time correlations here can easily be converted into two-frequency correlation functions via Fourier transforms (as defined in Appendix A.1). We find that $\delta(t-t') \rightarrow \delta(\omega+\omega')$. That is, $\left\langle \hat{b}^{\dagger}(\omega)\hat{b}(\omega')\right\rangle = n_{\mathrm{b}}\delta(\omega+\omega')$, and likewise for the other non-zero correlation functions.

[2] In addition to assuming the Markovian nature of the oscillator-bath interaction, the correlation functions given in Eq. 1.9 implicitly assume that the mechanical bandwidth γ_{i} is much smaller than ω_{m} (i.e., large mechanical Q-factor), such that the effective occupancy of the bath can be approximated by the thermal occupation at a single frequency ω_{m} rather than as an integral over all modes comprising the bath.

[3] Only the one- and two-time correlations are required for our purposes. However, as a thermal state is Gaussian (i.e., the associated quasiprobability distributions are Gaussian), any higher order noise correlations can be expressed in terms of the one- and two-time correlation functions [5], and thus Eqs. 1.9 and 1.10 fully characterize the mechanical and optical noise in the system.

1.1.2 Linearization Approximation

While the basic optomechanical interaction is fundamentally non-linear, all current physical realizations of optomechanical systems lie in the so-called *vacuum weak coupling* regime, where $g_0 \ll \kappa$. For such systems vacuum noise has a negligible effect on the system dynamics, and we may linearize Eqs. 1.7 and 1.8 to good approximation. When $g_0 \gtrsim \kappa$, the so-called *vacuum strong coupling regime*, a fully quantum nonlinear treatment is necessary [6–8].

For weakly coupled systems, a coherent drive (the $\alpha_{\rm in}$ term in Eq. 1.7), will result in a coherent steady state for both the optical and mechanical oscillators about which the system fluctuates. Formally we may make a displacement transformation for the system operators ($\hat{a} \rightarrow \alpha + \hat{a}$, $\hat{b} \rightarrow \beta + \hat{b}$) similar to that made for the input noise $\hat{a}_{\rm in}$. In particular we choose displacement amplitudes equal to the classical steady state solution for the driven optomechanical system (i.e., the solution to Eq. 1.7 and 1.8 when $\dot{\alpha} = \dot{\beta} = 0$ and quantum noise operators are neglected). These are found to be

$$\begin{aligned} \alpha &= \frac{\sqrt{\kappa_{\rm e}}\alpha_{\rm in}}{i\Delta' + \kappa/2}, \\ \beta &= \frac{-ig_0 n_{\rm c}}{i\omega_{\rm m} + \gamma_{\rm i}/2}, \end{aligned} \tag{1.11}$$

where the modified cavity-drive detuning is $\Delta' = \Delta + g_0\left(\beta + \beta^*\right)$, and we have used the fact that the steady-state intracavity photon number is given by $n_{\rm c} = |\alpha|^2$. For narrow bandwidth mechanical resonators ($\omega_{\rm m}/\gamma_{\rm i} \gg 1$) the static detuning shift is approximately equal to $-2g_0^2 n_{\rm c}/\omega_{\rm m}$. For the measurements made in this work this detuning shift is well under 1 MHz at all powers. As such a shift is much less than both κ and $\omega_{\rm m}$, it is negligible for our purposes and henceforth we will let $\Delta' \approx \Delta$.

By choosing these displacement amplitudes we remove the classical steady-state solution from our definition of $\hat{a}$ and $\hat{b}$ so that the remaining operators represent pure noise. Inserting this classical solution into Eqs. 1.7 and 1.8, and neglecting products of quantum noise operators (e.g., $\hat{a}^\dagger\hat{a}$) as "small", we arrive at the linearized quantum Heisenberg-Langevin equations:

$$\dot{\hat{a}} = -\left(i\Delta + \frac{\kappa}{2}\right)\hat{a} - iG\left(\hat{b} + \hat{b}^\dagger\right) + \sqrt{\kappa_{\rm e}}\hat{a}_{\rm in} + \sqrt{\kappa_{\rm i}}\hat{a}_{\rm i}, \tag{1.12}$$

$$\dot{\hat{b}} = -\left(i\omega_{\rm m} + \frac{\gamma_{\rm i}}{2}\right)\hat{b} - iG\left(\hat{a} + \hat{a}^\dagger\right) + \sqrt{\gamma_{\rm i}}\hat{b}_{\rm in}, \tag{1.13}$$

where $G = g_0\sqrt{n_{\rm c}}$ and for convenience we have chosen the phase of the input drive such that $\alpha \in \mathbb{R}$. This linearized set of equations now allows for simple analysis of the system output in the frequency domain via Fourier transforms.

Before continuing, we will discuss what "small" means in the context of this linearization approximation. In particular, under what conditions the fluctuations $\hat{a}$ and $\hat{b}$ can be considered small

relative to the classical steady-state. One can quantify the fundamental strength of the optomechanical interaction via a number of dimensionless parameters. The most important is g_0/κ, which specifies the per-phonon optical frequency shift in units of the optical linewidth. This ratio can be thought of as a "quantumness" parameter [6], as it is the only parameter affecting the optomechanical dynamics that vanishes in the classical limit $\hbar \to 0$. Provided this ratio is small, the quantum nature of the non-linearity may be neglected. Indeed, for a sideband resolved system with detuning $\Delta = \pm\omega_{\rm m}$ the ratio of the total number of scattered sideband photons to the steady-state intracavity photon number is given by $(2g_0/\kappa)^2 \langle n \rangle$. Thus, the fluctuations will always be sufficiently small if $g_0/\kappa \ll 1$ even as $n_{\rm c} \to 0$, provided $\langle n \rangle$ does not grow too large. This last point is important, as when $\langle n \rangle$ grows too large the sideband scattering rate will become large enough to substantially deplete the intracavity photons at the drive frequency, at which point the fluctuations can no longer be considered small and the system must be treated using a fully nonlinear, though classical theory [9]. The approximate threshold for when this breakdown occurs can be derived from a classical theory where the optical field is expanded in an infinite series of sidebands, with the value $n_{\rm c} = |\alpha|^2$ representing the energy in the 0-th order sideband. In this case we find the optomechanical interaction can be treated as a classical phase modulation with an effective modulation depth $z \approx 2g_0\sqrt{\langle n \rangle}/\omega_{\rm m}$, where the amplitude of the n-th order sideband is proportional to $J_{\rm n}(z)$ ($J_{\rm n}$ is the n-th order Bessel function of the first kind). Obviously, unless $z \ll 1$ the 0-th order sideband α may become much smaller than the higher order fluctuations, in principle even being extinguished completely, in which case the assumption of "small" fluctuations is no longer valid. In practice, as all current optomechanical systems possess relatively weak coupling, this regime is only reached when $\langle n \rangle$ is amplified well beyond its steady-state thermal value, as in the case of blue-detuned amplification in the sideband resolved regime discussed below in section 1.1.3.

Taking the Fourier transforms of Eqs. 1.12 and 1.13 (as defined in Appendix A.1), we arrive at the expressions for the cavity fluctuations in the frequency domain:

$$\hat{a}(\omega) = \frac{\sqrt{\kappa_{\rm e}}\hat{a}_{\rm in}(\omega) + \sqrt{\kappa_{\rm i}}\hat{a}_{\rm i}(\omega) - iG\left(\hat{b}(\omega) + \hat{b}^\dagger(\omega)\right)}{i\left(\Delta - \omega\right) + \kappa/2}, \tag{1.14}$$

$$\hat{b}(\omega) = \frac{\sqrt{\gamma_{\rm i}}\hat{b}_{\rm in}(\omega) - iG\left(\hat{a}(\omega) + \hat{a}^\dagger(\omega)\right)}{i\left(\omega_{\rm m} - \omega\right) + \kappa/2}. \tag{1.15}$$

We now insert Eq. 1.14 and its conjugate [4] into Eq. 1.15 and, after some algebraic manipulation, arrive at an equation for the mechanical fluctuations solely in terms of the noise operators:

$$\begin{aligned}\hat{b}(\omega) =& \chi_{\rm m}(\omega)\left(\sqrt{\gamma_{\rm i}}\hat{b}_{\rm in}(\omega) - iG\left(\chi_{\rm o}(\omega)\left(\sqrt{\kappa_{\rm e}}\hat{a}_{\rm in}(\omega) + \sqrt{\kappa_{\rm i}}\hat{a}_{\rm i}(\omega)\right)\right.\right.\\ &\left.\left. + \chi_{\rm o}^*(-\omega)\left(\sqrt{\kappa_{\rm e}}\hat{a}_{\rm in}^\dagger(\omega) + \sqrt{\kappa_{\rm i}}\hat{a}_{\rm i}^\dagger(\omega)\right)\right)\right),\end{aligned} \tag{1.16}$$

[4] Keep in mind that we have defined the Fourier transforms such that $\hat{a}^\dagger(\omega) = (\hat{a}(-\omega))^\dagger$, as explained in Appendix A.1.

where we have defined the mechanical and optical response functions $\chi_{\rm m}(\omega) = (i(\omega'_{\rm m} - \omega) + \gamma/2)^{-1}$ and $\chi_{\rm o}(\omega) = (i(\Delta - \omega) + \kappa/2)$, respectively, and where the mechanical frequency and linewidth have been modified by the optomechanical interaction such that $\omega'_{\rm m} = \omega_{\rm m} + \delta\omega_{\rm m}$ and $\gamma = \gamma_{\rm i} + \delta\gamma$, with frequency and linewidth modifications given by

$$\begin{aligned}\delta\omega_{\rm m}(\Delta) &= G^2 {\rm Im}\left[\chi_{\rm o}(\omega_{\rm m}) - \chi_{\rm o}^*(-\omega_{\rm m})\right] \\ &= G^2 {\rm Im}\left[\frac{1}{i(\Delta - \omega_{\rm m}) + \kappa/2} - \frac{1}{-i(\Delta + \omega_{\rm m}) + \kappa/2}\right],\end{aligned} \tag{1.17}$$

and

$$\begin{aligned}\delta\gamma(\Delta) &= 2G^2 {\rm Re}\left[\chi_{\rm o}(\omega_{\rm m}) - \chi_{\rm o}^*(-\omega_{\rm m})\right] \\ &= 2G^2 {\rm Re}\left[\frac{1}{i(\Delta - \omega_{\rm m}) + \kappa/2} - \frac{1}{-i(\Delta + \omega_{\rm m}) + \kappa/2}\right].\end{aligned} \tag{1.18}$$

Note that in obtaining Eq. 1.16, we have dropped terms proportional to $(i(\omega_{\rm m} - \omega) + \gamma_{\rm i}/2)^{-1}\,\hat{b}^\dagger(\omega)$. This is justified by the fact that $\hat{b}^\dagger(\omega)$ will be strongly peaked around $\omega = -\omega_{\rm m}$, and its product with $(i(\omega_{\rm m} - \omega) + \gamma_{\rm i}/2)^{-1}$ (which is peaked at $\omega = \omega_{\rm m}$) will be negligible when $\omega_{\rm m}/\gamma_{\rm i} \gg 1$. Additionally, we take $\omega \approx \omega_{\rm m}$ in the definitions of $\delta\omega_{\rm m}$ and $\delta\gamma$ as the response of $\hat{b}$ is only significant for $|\omega - \omega_{\rm m}| \lesssim \gamma \ll \kappa$. While the damping rate modification is very important, in practice the frequency shift is quite small, and it is often appropriate to let $\omega'_{\rm m} \approx \omega_{\rm m}$, which will be done from here on for notational simplicity.

1.1.3 Sideband-Resolved Systems

We now consider the particular case of a sideband-resolved system, where $\kappa \ll \omega_{\rm m}$, for a laser drive which is detuned to either the red or blue sideband ($\Delta = \pm\omega_{\rm m}$). Under these assumptions, the frequency and damping shifts are given by

$$\delta\omega_{\rm m}(\Delta = \pm\omega_{\rm m}) \approx \pm\frac{G^2}{\omega_{\rm m}}, \tag{1.19}$$

$$\delta\gamma(\Delta = \pm\omega_{\rm m}) \approx \pm\frac{4G^2}{\kappa} = \pm\gamma_{\rm OM}, \tag{1.20}$$

where we have defined the rate $\gamma_{\rm OM} = 4G^2/\kappa$. Physically, this rate represents the scattering rate of the resonant sideband photons due to the optomechanical interaction. Note that here we have defined $\gamma_{\rm OM}$ as a strictly positive quantity such that $\gamma = \gamma_{\rm i} \pm \gamma_{\rm OM}$, depending on detuning. This is in contrast to the typical convention, where the sign is included explicitly in the definition of $\gamma_{\rm OM}$ such that $\gamma = \gamma_{\rm i} + \gamma_{\rm OM}$ always. As many of our measurement results scale with the sideband scattering rate it is notationally convenient to define $\gamma_{\rm OM}$ as positive rather than constantly write $|\gamma_{\rm OM}|$. The

reader should keep this fact in mind when comparing formulae from this thesis to other works. Depending on the detuning, the damping modification will lead to either net cooling (red-detuning, $\Delta = \omega_{\mathrm{m}}$) or amplification (blue-detuning, $\Delta = -\omega_{\mathrm{m}}$) of the mechanical resonator. In particular, as derived in Appendix A.4, the average phonon occupancy $\langle n \rangle = \left\langle \hat{b}^{\dagger}\hat{b} \right\rangle$ is given for red-detuning by

$$\langle n \rangle \,|_{\Delta=\omega_{\mathrm{m}}} = \frac{n_{\mathrm{b}}}{1+C} + \frac{C}{1+C}\left(\frac{\kappa}{4\omega_{\mathrm{m}}}\right)^2, \tag{1.21}$$

and for blue-detuning by

$$\langle n \rangle \,|_{\Delta=-\omega_{\mathrm{m}}} = \frac{n_{\mathrm{b}}}{1-C} + \frac{C}{1-C}, \tag{1.22}$$

where we have defined the cooperativity $C = \gamma_{\mathrm{OM}}/\gamma_{\mathrm{i}}$, which quantifies the relative coupling strength of the mechanical resonator to the effective bath provided by the optical noise compared to the mechanical thermal bath. The second term in Eq. 1.21 represents the limit to backaction cooling due to scattering of the drive into the non-resonant sideband, and is often neglected as it is extremely small for systems with good sideband resolution. Note that for blue-detuning $\langle n \rangle$ does not actually diverge as $C \to 1$. Rather, the mechanical occupation grows until the linear approximation underlying these equations breaks down, as discussed above in section 1.1.2. At this point the stimulated scattering of sideband photons becomes great enough that the "pump" photons (i.e., photons at the drive frequency) will be substantially depleted, effectively depleting the gain of the system until the mechanical oscillator settles into a stable, self-oscillating limit cycle. This behavior can be accurately described using a classical nonlinear treatment of the optomechanical cavity [9, 10].

We can also simplify the equations for the fluctuation operators. Explicitly considering the case of red-detuning ($\Delta = \omega_{\mathrm{m}}$), we note that the contribution of $\hat{a}_{\mathrm{in}}^{\dagger}$ and $\hat{a}_{\mathrm{i}}^{\dagger}$ to $\hat{b}(\omega)$ will be negligible, as $\chi_{\mathrm{o}}^{*}(-\omega)$ and $\chi_{\mathrm{m}}(\omega)$ are sharply peaked at $-\omega_{\mathrm{m}}$ and ω_{m}, respectively, and $\omega_{\mathrm{m}} \gg \kappa, \gamma$. Moreover, as we are only concerned with frequencies $|\omega - \omega_{\mathrm{m}}| \ll \kappa$, we may approximate $\chi_{\mathrm{o}}(\omega) \approx 2/\kappa$. Thus, we obtain

$$\hat{b}(\omega)|_{\Delta=\omega_{\mathrm{m}}} \approx \chi_{\mathrm{m}}(\omega)\left(\sqrt{\gamma_{\mathrm{i}}}\hat{b}_{\mathrm{in}}(\omega) - i\sqrt{\frac{\gamma_{\mathrm{OM}}}{\kappa}}\left(\sqrt{\kappa_{\mathrm{e}}}\hat{a}_{\mathrm{in}}(\omega) + \sqrt{\kappa_{\mathrm{i}}}\hat{a}_{\mathrm{i}}(\omega)\right)\right). \tag{1.23}$$

We may insert this expression into Eq. 1.14, dropping the $\hat{b}^{\dagger}$ term as it only peaks near $-\omega_{\mathrm{m}}$, and utilize the scattering condition of Eq. 1.6 to obtain a solution for the output optical fluctuations in terms of the noise operators

$$\hat{a}_{\mathrm{out}}(\omega)|_{\Delta=\omega_{\mathrm{m}}} \approx r(\omega;+)\hat{a}_{\mathrm{in}}(\omega) + n(\omega;+)\hat{a}_{\mathrm{i}}(\omega) + s(\omega;+)\hat{b}_{\mathrm{in}}(\omega), \tag{1.24}$$

where $r(\omega;+)$, $n(\omega;+)$, and $s(\omega;+)$ are effective scattering matrix elements relating the system output to the relevant noise operators. Similarly, for blue detuning we may discard the $\hat{a}_{\mathrm{in}}$ and $\hat{a}_{\mathrm{i}}$

terms from Eq. 1.16, and the $\hat{b}$ term from Eq. 1.14, to obtain

$$\hat{a}_{\rm out}(\omega)|_{\Delta=-\omega_{\rm m}} \approx r(\omega;-)\hat{a}_{\rm in}(\omega) + n(\omega;-)\hat{a}_{\rm i}(\omega) + s(\omega;-)\hat{b}^{\dagger}_{\rm in}(\omega). \tag{1.25}$$

The effective matrix scattering elements for both detunings are given explicitly in Appendix A.5.

1.1.4 Heterodyne Detection

The total output field is the superposition of the classical cavity reflection amplitude, given by

$$\begin{aligned}\alpha_{\rm out} &= \alpha_{\rm in} - \sqrt{\kappa_{\rm e}}\alpha \\ &= \alpha_{\rm in}\left(1 - \frac{\kappa_{\rm e}}{i\Delta + \kappa/2}\right),\end{aligned} \tag{1.26}$$

and the fluctuation operator $\hat{a}_{\rm out}$ described in section 1.1.3. Placing this total output on a photodetector will result in an photocurrent proportional to $|\alpha_{\rm out}|^2 + \alpha^*_{\rm out}\hat{a}_{\rm out} + \alpha_{\rm out}\hat{a}^{\dagger}_{\rm out} + \hat{a}^{\dagger}_{\rm out}\hat{a}_{\rm out}$. The first term is a DC term, while the last term can be neglected as the fluctuations are assumed much smaller than the classical steady-state amplitude in the linear approximation. Thus, the relevant portion of the detected photocurrent (neglecting conversion factors relating optical power to electrical current) is

$$\hat{I} = \alpha^*_{\rm out}\hat{a}_{\rm out} + \alpha_{\rm out}\hat{a}^{\dagger}_{\rm out}. \tag{1.27}$$

As this signal is linear in the fluctuation field operators, a measurement of this type is sometimes referred to as "linear detection". Recall that, in the lab frame, the sideband noise fluctuates at a frequency $\omega_{\rm l} \pm \omega_{\rm m}$. Thus, beating the noise against the classical coherent amplitude not only amplifies these fluctuations, but also mixes them down in frequency to $\pm\omega_{\rm m}$, where conventional RF measurement equipment may be used.

This detection scheme is a specific case of a measurement technique called "heterodyning". In general, upon photodetection the output noise will beat against some strong coherent tone at frequency $\omega_{\rm l}$, known as the "local oscillator". This may be either the classical cavity reflection, as above, or some separate coherent tone mixed in prior to photodetection. The resulting RF portion of the output photocurrent (assuming the local oscillator is intense enough that small fluctuation terms may be ignored) will be given by

$$\hat{I} = |\alpha_{\rm LO}|\left(\hat{a}_{\rm out}\, e^{i\phi} + \hat{a}^{\dagger}_{\rm out}\, e^{-i\phi}\right), \tag{1.28}$$

where $|\alpha_{\rm LO}|$ is the local oscillator amplitude and ϕ is the relative phase between the local oscillator and the signal [5]. Ignoring for the moment any overall gain factors, the power spectral density (as

[5] We have glossed over a couple finer points here. In particular, the absence of additional amplified noise terms due

defined in Appendix A.1) of this photocurrent is given by

$$S_{\mathrm{II}}(\omega) = \int_{-\infty}^{\infty} d\omega' \left(\left\langle \hat{a}_{\mathrm{out}}(\omega)\hat{a}^{\dagger}_{\mathrm{out}}(\omega') \right\rangle + \left\langle \hat{a}^{\dagger}_{\mathrm{out}}(\omega)\hat{a}_{\mathrm{out}}(\omega') \right\rangle \right. \\ \left. + \left\langle \hat{a}_{\mathrm{out}}(\omega)\hat{a}_{\mathrm{out}}(\omega') \right\rangle \; e^{i2\phi} + \left\langle \hat{a}^{\dagger}_{\mathrm{out}}(\omega)\hat{a}^{\dagger}_{\mathrm{out}}(\omega') \right\rangle \; e^{-i2\phi} \right). \tag{1.29}$$

As shown in section 1.1.3, for a sideband-resolved system driven with detuning $\Delta = \pm\omega_{\mathrm{m}}$, $\hat{a}_{\mathrm{out}}$ will depend only on $\hat{a}_{\mathrm{in}}$, $\hat{a}_{\mathrm{i}}$ and either $\hat{b}_{\mathrm{in}}$ or $\hat{b}^{\dagger}_{\mathrm{in}}$ depending on the detuning. Thus, from the noise correlation relations in Eqs. 1.9 and 1.10 it follows that only the first two terms in the integrand of Eq. 1.29 will be non-vanishing. Consequently, under these conditions the output spectrum will not depend on the relative phase between the local oscillator and the signal.

Using Eqs. 1.24 and 1.25 for the output spectrum along with the noise correlation relations and the relation $|r(\omega;\pm)|^2 + |n(\omega;\pm)|^2 \pm |s(\omega;\pm)|^2 = 1$, it is straightforward to evaluate the spectral density. For red detuning we find

$$\begin{aligned} S_{\mathrm{II}}(\omega)|_{\Delta=\omega_{\mathrm{m}}} &= |r(\omega;+)|^2 + |n(\omega;+)|^2 + |s(\omega;+)|^2(n_{\mathrm{b}}+1) + |s(-\omega;+)|^2 n_{\mathrm{b}} \\ &= 1 + n_{\mathrm{b}}\left(|s(\omega;+)|^2 + |s(-\omega;+)|^2\right) \\ &= 1 + 2\frac{\kappa_{\mathrm{e}}}{\kappa}\gamma_{\mathrm{OM}}\; \bar{S}_{\mathrm{bb}}(\omega;\langle n\rangle\,|_{\Delta=\omega_{\mathrm{m}}}), \end{aligned} \tag{1.30}$$

where $\bar{S}_{\mathrm{bb}}$ is the symmetrized spectral density of $\hat{b}$, as defined in Appendix A.3. Meanwhile, for blue detuning we obtain

$$\begin{aligned} S_{\mathrm{II}}(\omega)|_{\Delta=-\omega_{\mathrm{m}}} &= |r(\omega;-)|^2 + |n(\omega;-)|^2 + |s(\omega;-)|^2 n_{\mathrm{b}} + |s(-\omega;-)|^2(n_{\mathrm{b}}+1) \\ &= 1 + (n_{\mathrm{b}}+1)\left(|s(\omega;+)|^2 + |s(-\omega;+)|^2\right) \\ &= 1 + 2\frac{\kappa_{\mathrm{e}}}{\kappa}\gamma_{\mathrm{OM}}\; \bar{S}_{\mathrm{b}^{\dagger}\,\mathrm{b}^{\dagger}}(\omega;\langle n\rangle\,|_{\Delta=-\omega_{\mathrm{m}}}). \end{aligned} \tag{1.31}$$

In a real measurement setup, there will exist some energy loss between the cavity and detector, including the non-unity quantum efficiency of the detector itself. To treat this loss while preserving the commutator of the output field, we can use the standard practice of introducing a fictitious beam splitter into the optical path prior to detection [5], which mixes the signal with uncorrelated environmental noise (vacuum noise, in the case of an optical signal). This results in a detected photocurrent of

$$\hat{I} = |\alpha_{\mathrm{LO}}|\left(\left(\sqrt{\eta}\,\hat{a}_{\mathrm{out}} + \sqrt{1-\eta}\,\hat{a}_{\mathrm{vac}}\right)\; e^{i\phi} + \left(\sqrt{\eta}\,\hat{a}^{\dagger}_{\mathrm{out}} + \sqrt{1-\eta}\,\hat{a}^{\dagger}_{\mathrm{vac}}\right)\; e^{-i\phi}\right), \tag{1.32}$$

where η is the total power detection efficiency of the measurement setup and $\hat{a}_{\mathrm{vac}}$ represents vacuum

to the local oscillator's vacuum noise assumes a particular setup, known as balanced heterodyning, which is discussed in Refs. [11] and [12]. Moreover, the local oscillator need not have frequency ω_{l}, in which case the optical sideband fluctuations will merely be shifted to frequencies other than $\pm\omega_{\mathrm{m}}$. These points are discussed further in section 3.1.

noise similar to (but uncorrelated with) $\hat{a}_{\rm i}$ and $\hat{a}_{\rm in}$. Due to the admission of this additional vacuum noise, we find that the noise floor is unchanged, while the Lorentzian terms centered at $\omega = \pm\omega_{\rm m}$ are reduced by a factor of η. Note that the formulae given in Eqs. 1.30 and 1.31 are normalized to the vacuum noise level. For red detuning, the actual physical detected spectral density, including electronic noise on the detector and the conversion between optical and electronic power, is

$$S_{\rm II}(\omega)|_{\Delta=\omega_{\rm m}} = S_{\rm dark} + \frac{G_{\rm e}^2}{R_{\rm L}} S_{\rm vac}^2 \left(1 + 2\eta\frac{\kappa_{\rm e}}{\kappa}\gamma_{\rm OM}\, \bar{S}_{\rm bb}(\omega;\langle n\rangle)\right), \tag{1.33}$$

where $S_{\rm dark}$ is the spectral density of the electronic noise, $G_{\rm e}$ and $R_{\rm L}$ are the optical conversion gain of the detector and input impedance of the spectrum analyzer, respectively, and $S_{\rm vac}$ is the optical vacuum noise power spectral density. Obviously, for blue detuning the same relation holds, with $\bar{S}_{\rm bb} \to \bar{S}_{\rm b^\dagger\, b^\dagger}$.

As the integrals of $\bar{S}_{\rm bb}$ and $\bar{S}_{\rm b^\dagger\, b^\dagger}$ are equal to $\langle n\rangle$ and $\langle n\rangle + 1$ (see Appendix A.3), we can perform calibrated thermometry of the mechanical resonator using the above relations. This is particularly simple when the local oscillator is separated from the cavity output. In this case, blocking the signal will result in a measurement of the amplified vacuum, with measured spectral density $S_{\rm noise} = S_{\rm dark} + \frac{G_{\rm e}^2}{R_{\rm L}} S_{\rm vac}^2$, while blocking both the signal and local oscillator will measure the dark noise of the detector $S_{\rm dark}$. For a red-detuned drive, the occupancy can be determined from the relation

$$\begin{aligned}\int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \frac{S_{\rm II}(\omega) - S_{\rm noise}}{S_{\rm noise} - S_{\rm dark}} &= \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \bar{S}_{\rm bb}(\omega;\langle n\rangle) \\ &= 2\eta\frac{\kappa_{\rm e}}{\kappa}\gamma_{\rm OM}\langle n\rangle,\end{aligned} \tag{1.34}$$

provided we know $\kappa_{\rm e}/\kappa$, $\gamma_{\rm OM}$ and η. For a blue-detuned measurement the same formula holds true, with $\langle n\rangle \to \langle n\rangle + 1$. This fact allows for an alternative method of thermometry for low $\langle n\rangle$. In particular, if we define $P_\pm$ as the total integrated sideband power for $\Delta = \pm\omega_{\rm m}$, then for small cooperativity (i.e., $\langle n\rangle|_{\Delta=\omega_{\rm m}} \approx \langle n\rangle|_{\Delta=-\omega_{\rm m}}$) we can define the sideband asymmetry parameter:

$$\xi = \frac{P_-}{P_+} - 1 = \frac{\langle n\rangle|_{\Delta=-\omega_{\rm m}} + 1}{\langle n\rangle|_{\Delta=\omega_{\rm m}}} - 1 \approx \frac{1}{\langle n\rangle}, \tag{1.35}$$

which allows for determination of $\langle n\rangle$ without calibration of the detection efficiency or any related cavity parameters. Provided that one can obtain sufficiently small $\langle n\rangle$ and maintain good signal quality for $C \ll 1$, this proves an accurate form of thermometry [13–16].

1.1.5 Phonon Counting

Linear detection via optical heterodyning allows for extremely sensitive measurements of the mechanical noise (in principle down to $\langle n\rangle \approx 0$), and has the advantage of being compatible with off-the-shelf photodetectors and detection electronics. However, a number of proposed applications

of optomechanical systems, such as the heralded preparation of non-classical mechanical states discussed in section 1.2.2, require the ability to measure single phonon emission or absorption events. This can be accomplished by counting scattered photons in the Stokes (anti-Stokes) sideband, as each photon is correlated with the emission (absorption) of a single phonon, effectively performing a phonon counting measurement [6].

As the full reflected signal from the cavity includes the reflected drive amplitude, which is many orders of magnitude larger than the sideband signal, we must first filter the cavity output to reject the drive. Provided the filter is sufficiently high-finesse, it can be modeled in the frequency domain by a single Lorentzian function:

$$F(\omega;\omega_{\rm f}) = \frac{\kappa_{\rm f}/2}{i(\omega_{\rm f}-\omega)+\kappa_{\rm f}/2}, \tag{1.36}$$

where $\kappa_{\rm f}$ and $\omega_{\rm f}$ are the resonant frequency of the filter, respectively. Now, explicitly considering the case of red-detuned driving in the sideband-resolved regime, the filtered cavity output will be the product of $F(\omega;\omega_{\rm f})$ with the frequency domain output of the cavity. As the resonantly enhanced anti-Stokes sideband photons will be detuned by $\omega_{\rm m}$ from the drive, we choose $\omega_{\rm f} = \omega_{\rm m}$. Using Eq. 1.24, we obtain

$$\hat{a}_{\rm filt}(\omega) = F(\omega;\omega_{\rm m})\left(\alpha_{\rm out}\delta(\omega) + r(\omega;+)\hat{a}_{\rm in}(\omega) + n(\omega;+)\hat{a}_{\rm i}(\omega) + s(\omega;+)\hat{b}_{\rm in}(\omega)\right). \tag{1.37}$$

Performing photon counting on the filtered output then results in an average count rate of

$$\begin{aligned}
\Gamma(t) &= \left\langle \hat{a}^\dagger_{\rm filt}(t)\hat{a}_{\rm filt}(t)\right\rangle \\
&= \frac{1}{2\pi}\int_{-\infty}^{\infty} d\omega \int_{-\infty}^{\infty} d\omega'\; e^{i(\omega+\omega')t}\left\langle \hat{a}^\dagger_{\rm filt}(\omega)\hat{a}_{\rm filt}(\omega')\right\rangle \\
&= \frac{1}{2\pi}\left(|F(0;\omega_{\rm m})|^2|\alpha_{\rm out}|^2 + \frac{\kappa_{\rm e}}{\kappa}\gamma_{\rm OM}\int_{-\infty}^{\infty} d\omega |F(\omega;\omega_{\rm m})|^2\; S_{\rm bb}(\omega;\langle n\rangle)\right) \\
&\approx A|\alpha^2_{\rm out}| + \frac{\kappa_{\rm e}}{\kappa}\gamma_{\rm OM}\langle n\rangle,
\end{aligned} \tag{1.38}$$

where $A = \frac{1}{2\pi}|F(0;\omega_{\rm m})|^2$ is the drive attenuation factor, and where we have assumed a filter bandwidth $\kappa_{\rm f} \gg \gamma$, allowing us to approximate $|F(\omega;\omega_{\rm m})|^2 \approx |F(\omega_{\rm m};\omega_{\rm m})|^2 = 1$ inside the integral over $S_{\rm bb}$. A similar analysis for blue detuning (where $\omega_{\rm f} = -\omega_{\rm m}$ to filter the Stokes sideband) yields a comparable result, with $\langle n\rangle \to \langle n\rangle + 1$. Note that this result also assumes that $\hat{a}_{\rm in}$ and $\hat{a}_{\rm i}$ are vacuum noise, and thus do not contribute to counting of real photons.

The total count rate, including noise of the photon counter and reduction of the drive and signal

[6] We use the term "phonon counting" here in the same sense as common usage of the term "photon counting". That is, counting individual phonons as they are created or destroyed. One should not confuse this with either a quantum nondemolition measurement of the phonon number or projective measurement of a phonon Fock state.

due to measurement inefficiency, is given for red-detuning by

$$\Gamma_{\rm tot} = \Gamma_{\rm dark} + \Gamma_{\rm drive} + \eta\frac{\kappa_{\rm e}}{\kappa}\gamma_{\rm OM}\left\langle n\right\rangle, \tag{1.39}$$

where $\Gamma_{\rm dark}$ is the dark count rate of the photon detector, $\Gamma_{\rm drive} = \eta A|\alpha_{\rm out}|^2$ and η is the total measurement efficiency. These expressions can be used to perform thermometry in a similar fashion to linear detection, either by calibrating the cavity parameters and total measurement efficiency or by measuring the asymmetry between the red- and blue-detuned count rates.

To assess the sensitivity of this counting scheme, it is convenient to express the measurement noise floor $\Gamma_{\rm dark} + \Gamma_{\rm drive}$ in terms of an equivalent number of mechanical quanta (that is the mechanical occupancy $\langle n\rangle$ that would be needed to yield a signal-to-noise of one). This noise-equivalent phonon number is obtained by dividing the total noise floor by the per-phonon count rate $\Gamma_{\rm SB,0} = \eta(\kappa_{\rm e}/\kappa)\gamma_{\rm OM}$, yielding

$$n_{\rm NEP} = \frac{\Gamma_{\rm dark} + \Gamma_{\rm drive}}{\Gamma_{\rm SB,0}}. \tag{1.40}$$

For a highly sideband-resolved system, the reflected drive in the case of $\Delta = \pm\omega_{\rm m}$ will be approximately given by $\alpha_{\rm out} \approx \alpha_{\rm in}$. This in turn can be expressed in terms of the intracavity photon number by using Eq. 1.11, yielding $|\alpha_{\rm out}|^2 \approx \omega_{\rm m}^2 n_{\rm c}/\kappa_{\rm e}$. Thus, $n_{\rm NEP}$ as a function of $n_{\rm c}$ is given by

$$n_{\rm NEP}(n_{\rm c}) = \frac{\kappa^2\Gamma_{\rm dark}}{4\eta\kappa_{\rm e}g_0^2 n_{\rm c}} + A\left(\frac{\kappa\omega_{\rm m}}{2\kappa_{\rm e}g_0}\right)^2. \tag{1.41}$$

As the above equation shows, the ultimate sensitivity of a phonon counting measurement is dependent primarily on the filter attenuation relative to the figure $g_0/\omega_{\rm m}$. As all current sideband-resolved systems lie in the regime $g_0 \ll \kappa \ll \omega_{\rm m}$, we require an extremely high degree of drive attenuation in order to achieve single phonon sensitivity ($n_{\rm NEP} < 1$). It should also be emphasized that, while in principle one could always increase $n_{\rm c}$ in order to reduce the effects of the dark count term, in practice we are limited in how much optical power we may input to a given cavity. In order to reach the drive-limited regime it is crucial to have efficient detection, low dark-counts, and relatively large g_0/κ.

1.1.6 Technical Laser Noise

Thus far we have assumed that the noise on the coherent drive, $\hat{a}_{\rm in}$, is simply vacuum noise. However, a real laser will have both intensity and phase noise. This additional technical noise will increase the optical noise floor above the ideal limits given in sections 1.1.4 and 1.1.5. More importantly, excess laser noise can also contribute to heating of the mechanical mode [17–20] and systematic errors in thermometry caused by noise squashing [21, 22] and anti-squashing.

For the specific drive laser and optomechanical devices used in this work, intensity noise has been

previously shown to be negligible in the vicinity of the mechanical frequency [23], and this section will deal only with excess phase noise of the laser. A third possibility is intracavity frequency noise, which has been found to be significant in some other optomechanical systems [22]. In dielectric optical microcavities the noise of this type is typically due to thermo-refractive fluctuations, which have been observed in silica microsphere [24] and microtoroid [25] resonators. However, thermo-refractive noise is generally less important at higher frequencies, and scales with the thermo-optic coefficient dn/dT, which decreases sharply as a function of temperature in Si at wavelengths near $\lambda \sim 1550$ nm [26]. As we have previously observed no signs of thermo-refractive noise at base temperatures of $T \approx 16$ K when measuring low frequency ($\omega_{\rm m} \sim 30$ MHz) optomechanical devices similar to those used in this work [27], we do not expect a significant contribution of thermo-refractive noise in our current measurements, which involve mechanical frequencies of $\omega_{\rm m} \sim 4-5$ GHz at base temperatures below 10 K.

Given a laser input amplitude $\alpha_{\rm in}$ (in the rotating frame), the presence of phase noise is captured by the addition of a stochastic phase factor $\phi(t)$, so that

$$\alpha_{\rm in}(t) = |\alpha_{\rm in}|\, e^{i\phi(t)}. \tag{1.42}$$

The phase noise is assumed to be a real, stationary Gaussian process with zero mean, such that it is characterized fully by the two-time correlation function. The phase noise is taken to be delta-correlated in the frequency domain, such that [17, 19, 23]

$$\langle \phi(\omega)\phi(\omega')\rangle = S_{\phi\phi}(\omega)\delta(\omega+\omega'), \tag{1.43}$$

where the expectation value here corresponds to a classical ensemble average. Note that $S_{\phi\phi}(\omega) = \bar{S}_{\phi\phi}(\omega)$ due to the realness of $\phi(t)$. From this, the two-time correlation function is given by

$$\langle \phi(t)\phi(t')\rangle = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi}\, e^{-i\omega(t-t')}\, S_{\phi\phi}(\omega). \tag{1.44}$$

Assuming the phase fluctuations are small, Eq. 1.42 may be expanded so that $\alpha_{\rm in}(t) \approx |\alpha_{\rm in}|\,(1+i\phi(t))$, and the spectral density of the input field is given by

$$S_{\alpha\alpha}(\omega) = |\alpha_{\rm in}|^2\,(2\pi\delta(\omega) + S_{\phi\phi}(\omega))\,. \tag{1.45}$$

As we are only concerned with the phase noise near the mechanical frequency, the $\delta(\omega)$ term may be ignored. The phase noise may be incorporated into our analysis by replacing the input noise operator $\hat{a}_{\rm in}$ by

$$\hat{a}_{\rm in,tot}(t) = \hat{a}_{\rm in}(t) + a_\phi(t), \tag{1.46}$$

where $\hat{a}_{\rm in}(t)$ still represents vacuum noise and $a_\phi(t) = i|\alpha_{\rm in}|\phi(t)$ is an additional (classical) stochastic noise term, with the frequency domain correlation function

$$\left\langle a_\phi^\dagger(\omega) a_\phi(\omega') \right\rangle = S_{\alpha\alpha}(\omega)\delta(\omega+\omega'). \tag{1.47}$$

Note that due to the real, classical nature of $\phi(t)$, we also have the relations

$$\left\langle a_\phi^\dagger(\omega) a_\phi(\omega') \right\rangle = \left\langle a_\phi(\omega) a_\phi^\dagger(\omega') \right\rangle, \tag{1.48}$$

$$\left\langle a_\phi(\omega) a_\phi(\omega') \right\rangle = \left\langle a_\phi^\dagger(\omega) a_\phi^\dagger(\omega') \right\rangle = -S_{\alpha\alpha}(\omega)\delta(\omega+\omega'). \tag{1.49}$$

By inserting this additional noise term into Eq. 1.16, we may repeat the analysis of Appendix A.4 and find that in a sideband-resolved system, for either red or blue detuning $\Delta = \pm\omega_{\rm m}$, the presence of phase noise heats the mechanical resonator. The additional added phonon occupancy is given by [23]

$$\langle n \rangle_\phi \,|_{\Delta=\pm\omega_{\rm m}} = \frac{\kappa_{\rm e}}{\kappa}\frac{\gamma_{\rm OM}}{\gamma} n_\phi, \tag{1.50}$$

where we have defined $n_\phi = S_{\alpha\alpha}(\omega_{\rm m})$, and where we have assumed that $S_{\alpha\alpha}(\omega)$ is sufficiently slow-varying in the vicinity of $\omega = \omega_{\rm m}$ (specifically for $|\omega - \omega_{\rm m}| \lesssim \gamma$) that we may approximate $S_{\alpha\alpha}(\omega) = S_{\alpha\alpha}(\omega_{\rm m})$ when integrating over frequency to obtain $\langle n \rangle$.

1.1.6.1 Phase Noise in Heterodyne Detection

As derived in Ref. [23], the presence of phase noise on the drive laser leads to an additional component of the heterodyne photocurrent power spectral density. Relative to the vacuum noise level, this is equal to

$$S_{\rm I_\phi I_\phi}(\omega)|_{\Delta=\pm\omega_{\rm m}} = n_\phi \left(\frac{2\kappa_{\rm e}}{\kappa}\right)^2 \left(1 \mp \frac{\gamma_{\rm OM}}{2}\frac{\gamma_{\rm i} \pm \gamma_{\rm OM}/2}{(\omega_{\rm m} \pm \omega)^2 + (\gamma/2)^2}\right). \tag{1.51}$$

In the limit of small cooperativity, this leads to squashing of the mechanical noise for red-detuning, and anti-squashing for blue-detuning. That is, the naively inferred phonon occupancy based on integrating the total observed sideband noise (as described in section 1.1.4) is

$$\langle n \rangle_{\rm inf} \,|_{\Delta=\pm\omega_{\rm m}, C\ll 1} \approx \frac{n_{\rm b} \mp (\kappa_{\rm e}/\kappa) n_\phi}{1 \pm C}. \tag{1.52}$$

For large cooperativity the phase noise causes anti-squashing for a red-detuned drive, leading to an inferred phonon occupancy

$$\langle n \rangle_{\rm inf} \,|_{\Delta=\omega_{\rm m}, C\gg 1} \approx \frac{n_{\rm b} + (\kappa_{\rm e}/2\kappa)\, C\, n_\phi}{1 + C}. \tag{1.53}$$

As these equations make clear, in the presence of significant phase noise, squashing/anti-squashing

effects can drastically affect the inferred mode temperature and the observed sideband asymmetry.

1.1.6.2 Phase Noise in Phonon Counting

Including the phase noise term a_ϕ in the derivation presented in section 1.1.5 leads to an additional contribution to the total photon count rate

$$\Gamma_\phi|_{\Delta=\pm\omega_\mathrm{m}} = \eta \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} S_{\alpha\alpha}(\omega)|F(\omega,\pm\omega_\mathrm{m})|^2|r(\omega;\pm)|^2. \tag{1.54}$$

If we make the usual assumptions that $S_{\alpha\alpha}(\omega)$ is slowly-varying in frequency for $|\omega-\omega_\mathrm{m}| \lesssim \gamma$ and that $\kappa_\mathrm{f} \gg \gamma$, this simplifies to

$$\Gamma_\phi|_{\Delta=\pm\omega_\mathrm{m}} = \eta n_\phi \left(\frac{\kappa_\mathrm{f}}{4}\left(1-\frac{2\kappa_\mathrm{e}}{\kappa}\right)^2 + \frac{\kappa_\mathrm{e}\gamma_\mathrm{OM}}{\kappa}\left(\frac{\gamma_\mathrm{OM}\kappa_\mathrm{e}}{\gamma\kappa} \pm \left(1-\frac{2\kappa_\mathrm{e}}{\kappa}\right)\right)\right). \tag{1.55}$$

Using the fact that $n_\phi = |\alpha_\mathrm{in}|^2 S_{\phi\phi}(\omega_\mathrm{m}) \approx \omega_\mathrm{m}^2 n_\mathrm{c} S_{\phi\phi}(\omega_\mathrm{m})/\kappa_\mathrm{e}$, we obtain the contribution of phase noise to the noise-equivalent phonon number:

$$n_{\mathrm{NEP},\phi}|_{\Delta=\pm\omega_\mathrm{m}} = \left(\frac{\omega_\mathrm{m}\kappa}{2\kappa_\mathrm{e}g_0}\right)^2 S_{\phi\phi}(\omega_\mathrm{m}) \left(\frac{\kappa_\mathrm{f}}{4}\left(1-\frac{2\kappa_\mathrm{e}}{\kappa}\right)^2 + \frac{\kappa_\mathrm{e}\gamma_\mathrm{OM}}{\kappa}\left(\frac{\gamma_\mathrm{OM}\kappa_\mathrm{e}}{\gamma\kappa} \pm \left(1-\frac{2\kappa_\mathrm{e}}{\kappa}\right)\right)\right). \tag{1.56}$$

Like the drive-bleed through, phase noise leads to a constant contribution to n_NEP, and leads to squashing or anti-squashing of the noise depending on detuning and cooperativity, similar to heterodyne detection. Note, however, that in the case $\kappa_\mathrm{e}/\kappa = 0.5$ the contribution of phase noise will not depend upon detuning. Thus, even in the presence of large phase noise it is possible to avoid detuning dependent noise squashing/anti-squashing, though one will still have a large overall phase noise floor.

It is useful for characterization purposes to calculate the phase-noise contribution to the observed count rates when the laser is far-detuned from the cavity resonance ($|\Delta| \gg \omega_\mathrm{m}$). Assuming that the laser-filter detuning is kept constant at $\pm\omega_\mathrm{m}$, the phase-noise count rate in this case is just

$$\begin{aligned}\Gamma_\phi|_{|\Delta|\gg\omega_\mathrm{m}} &= \eta \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} S_{\alpha\alpha}(\omega)|F(\omega,\pm\omega_\mathrm{m})|^2 \\ &= \eta\frac{\kappa_\mathrm{f}}{4} n_\phi,\end{aligned} \tag{1.57}$$

with a corresponding noise-equivalent phonon number

$$n_{\mathrm{NEP},\phi}|_{|\Delta|\gg\omega_\mathrm{m}} = \left(\frac{\omega_\mathrm{m}\kappa}{4\kappa_\mathrm{e}g_0}\right)^2 \kappa_\mathrm{f} S_{\phi\phi}(\omega_\mathrm{m}). \tag{1.58}$$

1.2 Applications of Sideband-Resolved Systems in the Quantum Regime

We now turn our attention to a brief discussion of some concrete applications of the sideband-resolved cavity optomechanical systems considered in the previous section. We will focus on two specific applications, each with potential uses in future quantum networks [28]. The first is coherent conversion between optical and microwave frequencies [29, 30], which allows for an efficient interface between superconducting microwave circuits [31] and optical photons for the purposes of long-distance communication between disparate quantum information processing nodes. The second is the heralded generation of non-classical mechanical states, including single phonon Fock states and entangled states of separate mechanical resonators [32–34], which can be utilized for scalable distribution of entangled quantum states [35, 36].

Though these tasks are by no means the only useful applications of cavity optomechanics, they both share technical challenges that serve to motivate the experimental work performed in this thesis. In particular, both require low thermal occupancy of the mechanical resonator, which may be achieved in optomechanical resonators with sufficiently high mechanical frequency by operating in a dilution refrigerator, as in Ref. [37]. Moreover, interfacing an optomechanical device with a superconducting microwave resonator *requires* operation in a dilution refrigerator, while the low thermal noise provided by the fridge environment allows for greatly enhanced coherence times of any generated non-classical states.

1.2.1 Wavelength Conversion

As the field of experimental quantum information processing advances, a key goal is to develop the ability to efficiently link disparate quantum processing units [28]. Optical signals traveling in fiber, particularly in the telecom band, are typically proposed as A communication channel due to the extremely low losses (~ 0.2 dB/km). Though one would ultimately like to develop efficient optical interfaces for a variety of quantum systems, a particularly relevant case is the problem of conversion between optical and superconducting microwave signals, given the tremendous recent successes of circuit QED systems in the field of quantum computing [38, 39]. As described below, optomechanics provides an efficient means of conversion between different electromagnetic wavelengths, utilizing a common mechanical resonator to connect two optical or microwave cavities. This mechanism has already been used to demonstrate efficient optical-to-optical wavelength conversion [40], as well as reversible optical-to-microwave conversion at the classical level [41]. It remains an open challenge to demonstrate optical-to-microwave conversion at the single photon level.

The basic concept behind the optomechanical wavelength conversion process is illustrated in Fig. 1.3. Two optical cavities, with photon annihilation operators and resonant frequencies $\hat{a}_{\mathrm{j}}$

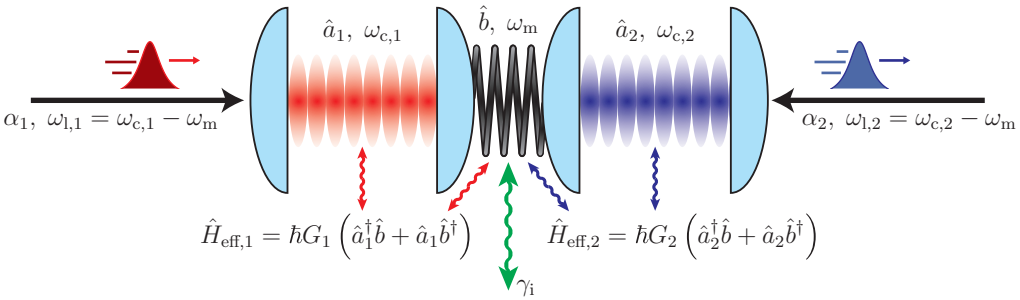

Figure 1.3: **Wavelength conversion.** Schematic of the wavelength conversion process. Two optical cavities, with resonant frequencies $\omega_{c,1}$ and $\omega_{c,2}$, respectively, are coupled to the same mechanical oscillator. Each cavity is driven by a red-detuned coherent tone (black arrows) with amplitude α_j $(j = 1, 2)$ and frequency $\omega_{l,j} = \omega_{c,j} - \omega_m$. In the sideband resolved regime, this results in an effective beam-splitter type interaction in each cavity with linearized coupling rates $G_j = \alpha_j g_{0,j}$. An additional weak optical signal input to cavity 1 at frequency $\sim \omega_{c,1}$ (red pulse) can then be converted into a signal exiting cavity 2 at frequency $\sim \omega_{c,2}$ (blue pulse) via the joint optomechanical interaction. The efficiency of this process depends on matching the effective scattering rates in each cavity and overcoming mechanical noise entering the system at rate γ_i (green arrow).

and $\omega_{c,j}$ $(j = 1, 2)$, respectively, are coupled to the same mechanical oscillator via the standard optomechanical interaction. In general each cavity will have a different coupling rate so that the total interaction Hamiltonian is given by $\hat{H}_{\text{int}} = \sum_j \hbar g_{0,j} \hat{a}_j^\dagger \hat{a}_j \left(\hat{b} + \hat{b}^\dagger\right)$. Both cavities are driven by a coherent tone, with amplitudes α_j and drive frequencies $\omega_{l,j}$. Provided each cavity is in the sideband resolved regime, and each drive tone is red-detuned $(\Delta_j = \omega_{c,j} - \omega_{l,j} = \omega_m)$, the effective interaction in the rotating wave approximation will have the form $\hat{H}_{\text{eff}} = \sum_j \hbar G_j \left(\hat{a}_j^\dagger \hat{b} + \hat{a}_j \hat{b}^\dagger\right)$, with $G_j = \alpha_j g_{0,j}$. If an additional weak optical signal is sent into cavity 1 under these conditions, the effective beam-splitter interaction will convert it into a phononic signal at the optomechanical scattering rate $\gamma_{OM,1} = 4G_1^2/\kappa_1$, which may then be converted back into an optical signal inside cavity 2 at rate $\gamma_{OM,2} = 4G_2^2/\kappa_2$, with a corresponding shift in optical frequency. The overall effectiveness of the conversion will depend on the relative photon-phonon scattering rates $\gamma_{OM,j}$. Additionally, as the beam-splitter interaction will convert *any* phonons into photons, any thermal noise in the mechanical resonator (which enters from the bath at rate γ_i) will appear as optical noise in the output.

One may work out the quantitative details of the conversion process by solving the linearized Heisenberg-Langevin equations of the joint optomechanical system to derive an effective scattering matrix relating the optical input in cavity 1 to the output at cavity 2 [30, 40, 42, 43]. In particular, we consider a weak optical input to cavity 1 (in addition to the red-detuned coherent drive) with spectral density $S_{\text{in},1}(\omega)$ in a frame rotating at $\omega_{l,1}$. The corresponding output spectral density from cavity 2, in a frame rotating at $\omega_{l,2}$, is

$$S_{\text{out},2}(\omega) = \frac{\kappa_{e,2}}{\kappa_2}\gamma_{OM,2}\left(\frac{\gamma_i}{(\omega+\omega_m)^2 + (\gamma/2)^2} n_b + \frac{\kappa_{e,1}}{\kappa_1}\frac{\gamma_{OM,1}}{(\omega+\omega_m)^2 + (\gamma/2)^2} S_{\text{in},1}(\omega)\right), \tag{1.59}$$

where $\gamma = \gamma_i + \gamma_{OM,1} + \gamma_{OM,2}$ is the total mechanical linewidth in the presence of the red-detuned driving of both cavities. The first term in Eq. 1.59 represents the output noise due to conversion of thermal phonons, while the second term represents the desired optomechanically induced wavelength conversion. The bandwidth of the conversion process is equal to the total mechanical linewidth γ, and the conversion efficiency is maximized for frequencies detuned by ω_m from the cavity drives, which corresponds to input and output at $\omega_{c,1}$ and $\omega_{c,2}$, respectively. The peak conversion efficiency is given by

$$\eta_{max} = \frac{\kappa_{e,1}}{\kappa_1}\frac{\kappa_{e,2}}{\kappa_2}\frac{4C_1C_2}{(1+C_1+C_2)^2}, \tag{1.60}$$

where $C_j = \gamma_{OM,j}/\gamma_i$, which makes it clear that we can in principle reach unity conversion efficiency, provided we can obtain $\gamma_{OM,1} = \gamma_{OM,2}$ and $\gamma_{OM,j} \gg \gamma_i$. That is, the photon-phonon conversion rates in each cavity must be equal (this can be thought of as an impedance matching requirement [30]) and must exceed the intrinsic dissipation rate of the mechanics so that phonons are not lost before being converted back to photons. Additionally, Eq. 1.60 makes it clear that we need $\kappa_{e,j}/\kappa_j$ to be as high as possible, as photons lost to intrinsic absorption are useless for our purposes.

In addition to the conversion efficiency, we must also consider the output optical noise added by the presence of thermal noise in the mechanics [7]. This is most conveniently expressed in an effective number of added noise quanta [40]:

$$n_{added} \approx 2\frac{\kappa_1}{\kappa_{e,1}}\frac{\gamma_i n_b}{\gamma}. \tag{1.61}$$

As the above equation makes clear, one necessary requirement for quantum limited wavelength conversion is to achieve a large *effective* cooperativity, $C_{eff} = \gamma/\gamma_i n_b$. Note that in the large cooperativity limit, where conversion is most efficient, this is approximately given by $C_{eff} \sim \langle n \rangle^{-1}$. The necessary low occupancy can in principle be achieved solely via sideband cooling due to the red-detuned drives necessary for the conversion process. However, while sideband cooling has been previously used [45, 46] to cool mechanical resonators to $\langle n \rangle < 1$, occupancies lower than $\langle n \rangle \approx 0.2 - 0.3$ have not been achieved in any system via this method. In particular, the Si optomechanical devices used in this work (detailed in section 1.3.4), suffer from degradation of the optical and mechanical quality factor at high powers due to two-photon absorption [45], providing a limit to the effects of laser cooling. Passive cooling of the device inside a dilution refrigerator offers a route to circumvent this technical roadblock and achieve $C_{eff} \gg 1$, and thus $n_{added} \ll 1$.

Even more importantly, as mentioned above one of the most promising applications of the optomechanical wavelength conversion process is to provide an efficient and reversible transducer between quantum optical and microwave signals. As the two cavities must couple to the same me-

[7] There is also noise in the output due to quantum backaction from the coherent drive. However, these terms are proportional to $(\kappa/4\omega_m)^2$, and thus can typically be neglected for sideband resolved systems until we reach the regime $\gamma_i n_b/\gamma \sim (\kappa/4\omega_m)^2$. For full detail see Refs. [40] and [44].

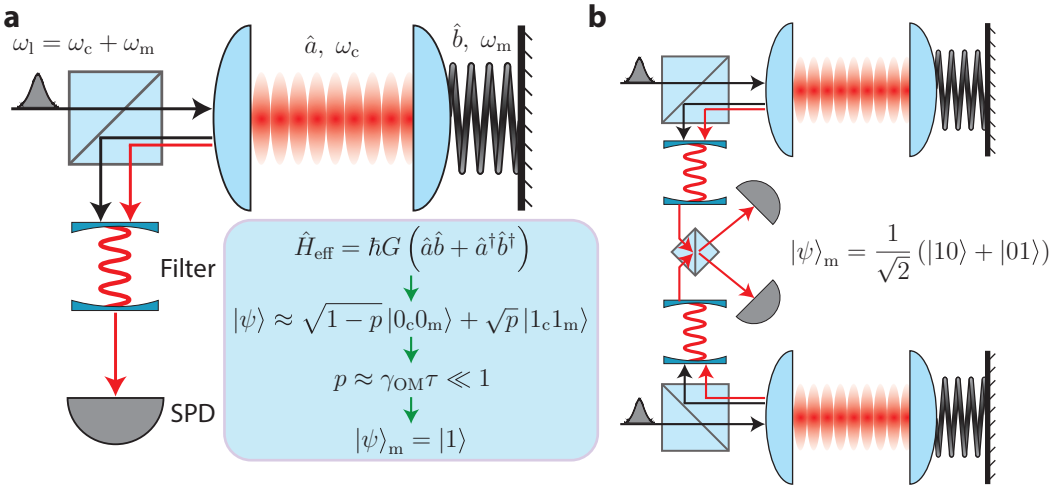

Figure 1.4: **Heralded Fock and entangled states. a**, Schematic of the process used to generate a single phonon Fock state. A short blue-detuned ($\Delta = -\omega_m$) pulse (with pulse width τ such that $\gamma_{OM}\tau \ll 1$) is used to drive the cavity, resulting in an effective two-mode squeezing Hamiltonian in the sideband resolved limit. After interacting with the cavity, there is a probability $p \approx \gamma_{OM}\tau$ that an entangled photon-phonon pair is created. After filtering out the drive frequency ω_l, measurement of a photon using a single photon detector (SPD) projects the mechanical system into a Fock state. **b**, To create an entangled state, two cavities are driven in this manner, and their outputs combined on a 50/50 beam-splitter prior to the SPD. This erases the which-path information such that a single click on the SPD projects the joint system into an entangled state.

chanical oscillator, it is necessary to bring the optical cavity into close physical proximity to the superconducting microwave resonator, which requires operation inside a dilution fridge.

1.2.2 Heralded State Preparation

At its most basic level, the optomechanical interaction allows us to add and subtract phonons from the mechanical resonator. Even in the vacuum weak-coupling regime, this allows for the creation of a variety of non-classical mechanical states using appropriate red- or blue-detuned pulses [32], provided we are content to generate such states in probabilistic, heralded fashion. In particular, recent proposals have explicitly considered heralding of both a single phonon Fock state [34] and an entangled state of two mechanical resonators [33, 47] via application of a short blue-detuned pulse to an optomechanical system in its motional quantum ground state and subsequent detection of the emitted sideband photon.

The basic idea behind heralded generation of a Fock state is illustrated in Fig. 1.4a. A short, blue-detuned drive pulse is sent into the cavity which, as discussed in section 1.1, will result in an effective interaction Hamiltonian (in the sideband-resolved limit) $\hat{H}_{eff} = \hbar G\left(\hat{a}\hat{b} + \hat{a}^\dagger\hat{b}^\dagger\right)$ which creates and annihilates photons and phonons in correlated pairs. If the pulse is sufficiently short (specifically, if the pulse width τ is short enough that $\gamma_{OM}\tau \ll 1$), it will generate a single photon-

phonon pair with a small probability $p \approx \gamma_{\mathrm{OM}}\tau$, with a vanishingly small chance of higher-order (i.e., multi-phonon) events. In this case, provided the mechanical system starts in the ground state, the approximate state of the system after filtering out the drive pulse will be $|\psi\rangle \approx \sqrt{1-p}\,|0_{\mathrm{c}}0_{\mathrm{m}}\rangle + \sqrt{p}\,|1_{\mathrm{c}}1_{\mathrm{m}}\rangle$, where subscripts c and m refer to the optical cavity and mechanical resonator subsystems, respectively. Subsequent detection of the sideband photon on a single photon counter (which occurs with probability ηp, where η is the total detection efficiency) will herald the creation of a single phonon Fock state.

The same protocol can be extended to creating an entangled state of two separate mechanical resonators, as shown in Fig. 1.4b. Two optomechanical cavities are driven by a blue-detuned pulse, as in the case of Fock state generation, and their respective outputs are combined using a 50/50 beam-splitter after filtering out the drive. This erases any information about which cavity emitted a sideband photon, such that detection of a single sideband photon projects the joint mechanical system into an entangled state. Similar protocols have been used successfully in the past to generate heralded entanglement of atomic ensembles [48, 49], trapped ions [50], and even vibrational modes in macroscopic diamonds [51].

This form of heralded entanglement provides the basic building block for a "quantum repeater", enabling the scalable distribution of entangled states utilizing the well known DLCZ protocol [35,36]. Beyond applications in quantum networking, however, the generation of entangled (or other non-classical) states of mesoscopic mechanical objects potentially allows for the testing of novel theories regarding quantum decoherence [52] or the nature of entanglement at the macroscopic scale [33, 51, 53], as well as additional practical applications such as improving the measurement sensitivity of weak forces [54].

The fidelity of the heralding process is directly related to the initial thermal occupancy of the mechanical mode. Unity fidelity can only be achieved in principle by starting fully in the ground state ($\langle n\rangle = 0$), while generally the initial mechanical state will be a thermal state with $\langle n\rangle > 0$. However, high fidelity can still be achieved provided $\langle n\rangle \ll 1$. In particular, the fidelity of heralded Fock state generation will be roughly given by $(\langle n\rangle + 1)^{-1}$, the probability of measuring a thermal state in the mechanical ground state. Again, as with wavelength conversion, the necessary low occupancies can in principle be achieved by using a strong red-detuned pulse to pre-cool the system via the optomechanical interaction, as considered in Ref. [34]. Due to the aforementioned limitations of laser cooling in current realizable systems, however, we are forced to turn to passive cooling in a dilution fridge as a means to obtain $\langle n\rangle \ll 1$.

An additional prerequisite for the heralding protocols discussed here is the ability to count single emitted sideband photons with high efficiency. Indeed, without the effective nonlinearity provided by the photon counter it is impossible to generate non-Gaussian states such as a Fock state from solely Gaussian optical and mechanical inputs. While this type of measurement is conceptually

straightforward, it should be noted that prior to the work performed in this thesis such measurements had not been demonstrated. Though the development of this kind of phonon counting measurement is not the primary focus of this work, it is technically nontrivial as efficient detection requires both strong optomechanical coupling, low detector dark counts, and a high degree of drive suppression, as detailed in section 1.1.5. The measurements performed in Chapter 4, alongside a previous experiment detailed in Refs. [55] and [56], are the first experiments to perform phonon counting in a cavity optomechanical system.

1.3 Optomechanical Crystals

As mentioned previously, the basic physics of cavity optomechanics apply to any system where mechanical motion couples to the frequency of an electromagnetic cavity, and a variety of experimental geometries have been realized in the last several years across an enormous range of length and mass scales, ranging from suspended gram scale mirrors [57], to reflective micromechanical membranes [58–61], to microtoroids, disks, and spheres [18, 44, 62, 63], to superconducting microwave circuits [64–67], all the way down to cold atomic ensembles trapped in an optical cavity [68]. One notices a general trend in these systems towards smaller optical and mechanical resonators, with physical dimensions and motional masses pushed below the micrometer and femtogram scale, respectively. This is no accident, and relates to the strength of the fundamental optomechanical interaction.

In section 1.1 we explicitly considered an optomechanical system composed of a Fabry-Perot optical cavity with one mirror mounted on the spring. In this case the optomechanical coupling rate g_0 is easily shown to be inversely proportional to both the length of the cavity and the square root of the mirror mass (this latter fact comes from the relation $x_{\mathrm{ZPF}} = \sqrt{\frac{\hbar}{2m\omega_{\mathrm{m}}}}$, where m is the mass of the mirror). Thus, the smaller the cavity and motional mass, the larger the fundamental quantum interaction strength. Although we have explicitly shown this to be the case for this simple geometry, in general it is true that the vacuum optomechanical coupling rate will increase as the size of the cavity decreases, much in the same way that the interaction strength between an atom and an optical cavity is inversely proportional to the effective cavity volume. Reducing the motional mass of the oscillator typically has the added benefit of raising the mechanical frequency, making it easier to achieve the sideband-resolved regime and allowing one to reach a low phonon occupancy simply by passive cooling in a cryogenic environment.

However, keep in mind that the parameters which really matter are the ratios g_0/κ (which quantifies the strength of the optomechanical interaction) and $\omega_{\mathrm{m}}/\kappa$ (which quantifies the level of sideband resolution). As volume of the cavity is reduced it is crucial to maintain a low optical decay rate κ, or equivalently a large optical quality factor $Q_{\mathrm{c}} = \omega_{\mathrm{c}}/\kappa$. Thus, roughly speaking, what we would like

to do is maximize the ratio Q/V for both the optical and mechanical resonators, again in analogy to the problem of atom-cavity coupling encountered in cavity quantum electrodynamics. One strategy for achieving this goal is to use photonic crystal cavities, which have demonstrated Q-factors approaching or in excess of 10^6 at telecom frequencies $\lambda \approx 1300 - 1600$ nm in common semiconductor materials such as Si and GaAs [69, 70] while maintaining effective cavity volumes close to the theoretical minimum value of $(\lambda/2n)^3$. Beyond confining photons, with proper engineering such devices can also be made into phononic crystal cavities which simultaneously create a high-quality acoustic mode co-localized with the optical resonance. Such structure are dubbed optomechanical crystals (OMCs).

1.3.1 Photonic Crystals

As the name suggests, the physics of photonic crystals [71] are analogous to those of crystalline solids. That is, a periodic array of dielectric scatterers (in our case, holes etched into a dielectric membrane) will affect the allowed propagating solutions of an electromagnetic wave in much the same way that a periodic electronic potential will modify the quantum mechanical motion of an electron. In particular, for the right periodicity destructive interference between fields scattered from adjacent holes can give rise to bandgaps, regions of frequency where light cannot propagate freely in the structure. Such bandgaps serve as excellent effective mirrors for an optical cavity. We will briefly present the main ideas necessary for understanding the optical structures used in this work. A full treatment of the theory and properties of photonic crystals can be found in Refs. [72] and [73].

Mathematically, we consider a linear, lossless, and isotropic dielectric medium. As we typically only care about solutions to Maxwell's equations within a particular frequency range, it is appropriate to ignore dispersion such that the material is completely characterized by a scalar dielectric function $\epsilon(\boldsymbol{x}) = \epsilon_0\epsilon_{\mathrm{r}}(\boldsymbol{x})$ which is only a function of the position vector $\boldsymbol{x}$. Note that ϵ_0 is the constant vacuum permittivity, and the relevant quantity is the relative permittivity ϵ_{r}, which can be equivalently represented by the index of refraction n defined by $\epsilon_{\mathrm{r}} = n^2$. As most materials of interest are nonmagnetic, we usually set the magnetic permeability equal to the vacuum permeability ($\mu = \mu_0$) [8]. As we have assumed linear materials, an arbitrary solution of Maxwell's equations can be decomposed into a sum of time-harmonic solutions. That is, the total magnetic field may be written as

$$\boldsymbol{H}(\boldsymbol{x}, t) = \sum_{\mathrm{m}} = a_{\mathrm{m}}\boldsymbol{H}_{\mathrm{m}}(\boldsymbol{x})e^{-i\omega_{\mathrm{m}}t}, \tag{1.62}$$

and similarly for the electric field $\boldsymbol{E}$. Finding these time-harmonic solutions can then be cast as a

[8] It is straightforward, though more cumbersome, to deal with anisotropic, magnetic materials, and small amounts of material absorption. Details can be found in Ref. [72].

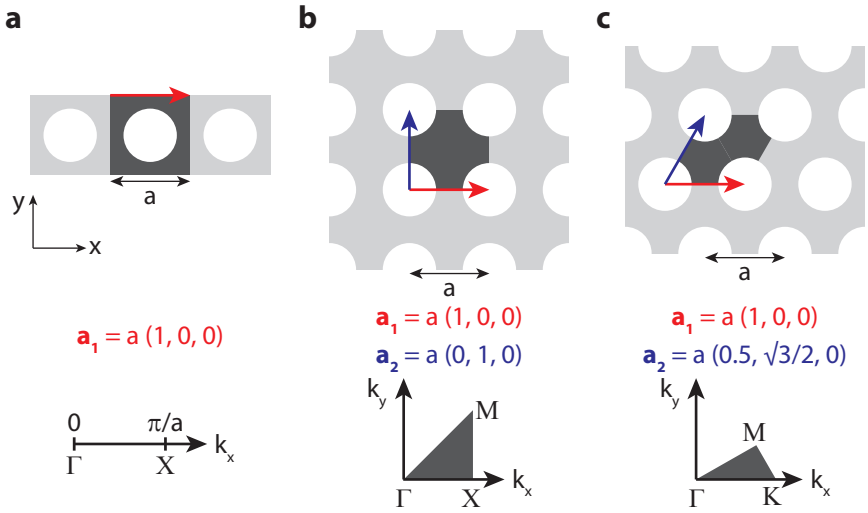

Figure 1.5: **One- and two-dimensional lattices.** Schematics of the different types of crystal lattices relevant to the work in this thesis, including **a**, one-dimensional, **b**, two-dimensional square, and **c**, two-dimensional hexagonal lattices. The primitive lattice vectors and their explicit representations in the xyz coordinate system are shown by red/blue arrows and text, respectively, while the unit cell for each lattice is shown in dark gray. The corresponding first irreducible Brillouin zone is illustrated below each lattice.

Hermitian eigenvalue problem [72]:

$$\hat{\Theta}\boldsymbol{H}_{\mathrm{m}}(\boldsymbol{x}) = \nabla \times \left(\frac{1}{\epsilon_r(\boldsymbol{x})}\nabla \times \boldsymbol{H}_{\mathrm{m}}(\boldsymbol{x})\right) = \left(\frac{\omega_{\mathrm{m}}}{c}\right)^2 \boldsymbol{H}_{\mathrm{m}}(\boldsymbol{x}). \tag{1.63}$$

We are interested specifically in dielectrics possessing discrete translational symmetry. That is,

$$\epsilon_r(\boldsymbol{x}) = \epsilon_r(\boldsymbol{x} + \boldsymbol{R}), \tag{1.64}$$

where $\boldsymbol{R} = \sum_{\mathrm{j}} m_{\mathrm{j}}\boldsymbol{a}_{\mathrm{j}}$, $m_{\mathrm{j}} \in \mathbb{Z}$, and the vectors $\boldsymbol{a}_{\mathrm{j}}$ are the primitive lattice vectors which define the periodic structure of the dielectric. In particular these lattice vectors help define the unit cell, a primitive dielectric structure which is tiled infinitely throughout space to form the crystal, and the lattice constant a, which is the fundamental length scale of the system. Though there are many crystal lattices possessing distinct symmetries, in this work we are concerned only with three cases, in which the dielectric is periodic in either one or two dimensions. These three lattices, one-dimensional, two-dimensional square, and two-dimensional hexagonal, are illustrated in Fig. 1.5.

For each set of primitive lattice vectors we may define a set of reciprocal lattice vectors $\boldsymbol{b}_{\mathrm{j}}$ such that $\boldsymbol{a}_{\mathrm{i}} \cdot \boldsymbol{b}_{\mathrm{j}} = 2\pi\delta_{\mathrm{ij}}$ (δ_{ij} is the Kronecker delta function). The discrete translational symmetry of the lattice guarantees that Bloch's theorem [74] holds, which states that any solution of the wave

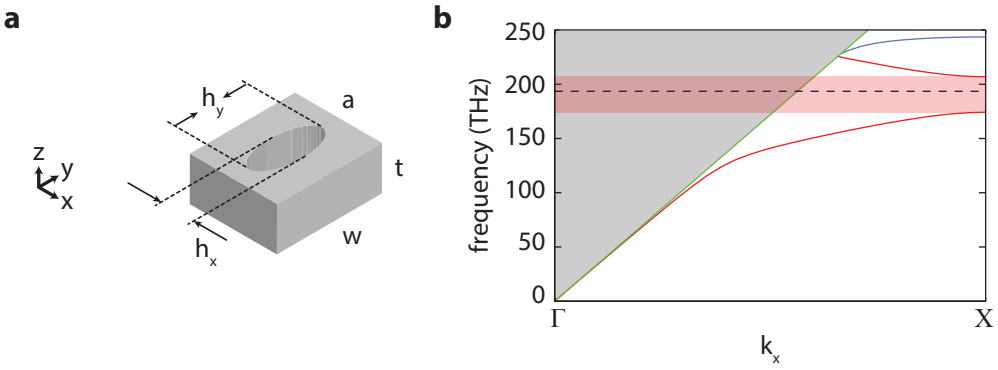

Figure 1.6: **Nanobeam optical bandstructure.** **a**, Three-dimensional unit cell of a nanobeam photonic crystal, with relevant dimensional parameters labeled. This structure has periodicity only along the axial (x) direction. **b**, Computed photonic band structure, showing mode frequency versus the axial component of the wavevector k_{x} for z-symmetric modes, assuming a Si beam (refractive index $n = 3.48$) and specific dimensions $(a, t, w, h_{\mathrm{x}}, h_{\mathrm{y}}) = (436, 220, 529, 165, 366)$ nm. Bands shown in red are modes of odd vector symmetry in y, while blue bands are even. The red shaded region indicates the pseudo-bandgap, while the gray shaded region is the light cone. The frequency of interest ($\nu \sim 194$ THz, corresponding to $\lambda \sim 1550$ nm) is indicated by a black dashed line.

equation can be expressed in the form

$$\boldsymbol{H}_{\boldsymbol{k}}(\boldsymbol{x}) = e^{i\boldsymbol{k}\cdot\boldsymbol{x}}\boldsymbol{h}_{\boldsymbol{k}}(\boldsymbol{x}), \tag{1.65}$$

where the wavevector $\boldsymbol{k} = \sum_{\mathrm{j}} l_{\mathrm{j}}\boldsymbol{b}_{\mathrm{j}}$, $l_{\mathrm{j}} \in \mathbb{Z}$ exists in the reciprocal lattice and $\boldsymbol{h}_{\boldsymbol{k}}(\boldsymbol{x})$ is a periodic function on the lattice (i.e., $\boldsymbol{h}_{\boldsymbol{k}}(\boldsymbol{x}) = \boldsymbol{h}_{\boldsymbol{k}}(\boldsymbol{x} + \boldsymbol{R})$). The eigenvalue problem can then be cast in the form [72]

$$\hat{\Theta}_{\boldsymbol{k}}\boldsymbol{h}_{\boldsymbol{k}}(\boldsymbol{x}) = (i\boldsymbol{k} + \nabla) \times \left(\frac{1}{\epsilon_r(\boldsymbol{x})}(i\boldsymbol{k} + \nabla) \times \boldsymbol{h}_{\boldsymbol{k}}(\boldsymbol{x})\right) \left(\frac{\omega(\boldsymbol{k})}{c}\right)^2 \boldsymbol{h}_{\boldsymbol{k}}(\boldsymbol{x}). \tag{1.66}$$

Due to the periodic boundary conditions imposed upon $\boldsymbol{h}_{\boldsymbol{k}}$, the eigenvalue problem can be restricted to a single unit cell of the crystal, which results in a discrete spectrum of eigenfrequencies $\omega_{\mathrm{j}}(\boldsymbol{k})$ for any given wavevector $\boldsymbol{k}$ (i.e., the band structure of the crystal). Furthermore, due to both the discrete translational symmetry and the other symmetries of the lattice point group [74], we may restrict our attention to a limited region of the reciprocal lattice, known as the first irreducible Brillouin zone (IBZ). Solutions to Eq. 1.66 for any other wavevectors $\boldsymbol{k}$ may then be obtained from this subset of solutions by applying the appropriate symmetry operations. The IBZs are illustrated below their respective lattices in Fig. 1.5. Note that for the one-dimensional lattice in Fig. 1.5a the IBZ is merely the line $k_{\mathrm{x}} \in [0, \pi/a]$. In practice one solves for the band structure of a given dielectric lattice via one of a variety of numerical methods. We use a planewave-expansion based algorithm developed at MIT [75], which is implemented in the freely available MIT Photonic Bands (MPB) software package [76].

To illustrate the properties of photonic band structures, we will focus on the structure of primary interest in this work: the nanobeam photonic crystal, whose full three-dimensional unit cell is shown in Fig. 1.6a. This structure, like the schematic shown in Fig. 1.5a, has periodicity only along the axis of the beam, denoted here as the x-direction, while the other dimensions of the beam are finite in extent (the beam is assumed to be surrounded by vacuum with $n = 1$). A typical computed band structure for this class of photonic crystal is shown in Fig. 1.6b, where one observes a bandgap (red shaded region) in which no guided modes of the structure exist. However, as the crystal is finite in extent in y and z there exist different classes of modes. The modes we care about are the "guided" modes, for which the electromagnetic fields are confined via index-guiding (i.e., total internal reflection) to a small region around the nanobeam. In addition to these modes there exists a continuum of modes which can carry energy infinitely far away from the nanobeam, the so-called "radiation" modes. As these modes must satisfy the free-space condition $\omega^2 = c^2|\boldsymbol{k}|^2$, they can only exist in a particular region of frequency-momentum space known as the "light cone" (gray shaded region) defined for a nanobeam by $\omega > ck_x$. Finally, there exist "leaky" modes, guided modes which, while localized near the waveguide, exist in the light cone region and can strongly couple to radiation modes. Due to the existence of radiation modes, we can never have true photonic bandgaps in such partially index-guided structures, but rather pseudo-gaps in which no guided or leaky modes exist.

There are a few additional points worth making before moving on. The first is that, despite the lack of translational symmetry in y or z, there is *reflection* symmetry in both directions about the center of the beam. This allows us to classify the solutions of Maxwell's equations by their vector symmetry with respect to reflection about the $y = 0$ and $z = 0$ planes (assuming the origin is placed along the beam axis). As we will see in section 1.3.3, in practice we are only concerned with the lowest frequency optical band. For the optically thin structures considered in this work this mode will always be of even vector symmetry in z (these are the so-called "TE-like" modes, whose electric field lies predominantly in the $z = 0$ plane). Moreover, for the nanobeam the lowest lying mode is always of even vector symmetry in y, and thus when computing photonic band structures it is sufficient to compute only modes which are even (odd) symmetry in z (y), which allows us to simulate only one-quarter of the full unit cell provided the appropriate boundary conditions are applied at the symmetry planes.

The second point is that, due to the scale invariance of Maxwell's equations the frequencies of the optical bands can be shifted arbitrarily without changing the band structure itself by making appropriate transformations to the dielectric function and structural dimensions [72]. In particular, uniformly scaling the dimensions of the crystal by some factor s (formally making the transformation $\epsilon'(\boldsymbol{x}) = \epsilon(\boldsymbol{x}/s)$) will simply result in the same band structure with scaled frequencies and mode profiles $\omega' = \omega/s$ and $\boldsymbol{u}'_{\boldsymbol{k}}(\boldsymbol{x}) = \boldsymbol{u}_{\boldsymbol{k}}(\boldsymbol{x}/s)$, respectively. Once a suitable band structure is obtained it is trivial to adjust the design parameters such that the particular frequency of interest lies in

the bandgap [9]. When performing band structure simulations the dimensions of the structure are typically expressed in units of the fundamental lattice constant a, which can then be selected to yield the desired mode frequencies. Similarly, scaling the dielectric function according to $\epsilon'(\boldsymbol{x}) = s^2\epsilon(\boldsymbol{x})$ scales the frequencies as $\omega' = s\omega$ with no change in the band structure or mode profiles. This is only true if one scales the *entire* dielectric function (i.e., that of the cladding as well as the nanobeam). If one only changes the nanobeam material the band structure can and will change, with smaller index contrasts yielding smaller bandgaps, all else being equal.

One final point is that, while there seem to be many free parameters in even this simple, quasi-one-dimensional problem, there are a number of rules of thumb which can provide a rough starting point for the design. In this case we want to maximize the size of the bandgap and place our desired cavity frequency roughly in the center of the gap (the attenuation and corresponding confinement of a mode in the bandgap region is higher the further the mode frequency is from the band edges). Note that in Fig. 1.6b the bandgap is located around a frequency of $\nu \sim 200$ THz, where the refractive index of Si is roughly $n = 3.48$. The corresponding wavelength is thus $\lambda = c/n\nu \sim 430$ nm $\sim a$. In general one finds that any bandgap will tend to occur at frequencies such that the lattice constant a is comparable to one wavelength in the material (λ/n), as this maximizes interference between the scattered fields from adjacent unit cells. Furthermore the size of the bandgap, typically measured by the dimensionless constant $\Delta\nu/\nu_0$, where $\Delta\nu$ is the frequency span of the gap and ν_0 is the mid-gap frequency, is maximized for a particular thickness and hole size, which depends on the index contrast between the beam and the cladding [72]. In our case the index contrast is about 3.5, and optimal thickness tends to be around $t \sim (0.5 - 1)a$, while the optimal hole size is typically such that the filling fraction (the ratio between the hole area and the total unit cell area) is roughly $0.2 - 0.25$. The bandgap isn't terribly sensitive to the hole *shape* (e.g., whether the hole is circular or oblate) as long as the filling fraction is kept fixed. This will prove useful when designing a simultaneous optical and mechanical cavity.

1.3.2 Phononic Crystals

The propagation of acoustic waves through a periodic elastic medium can be treated in a similar way to the propagation of electromagnetic fields [77, 78]. Thus, any photonic crystal will also serve as a *phononic* crystal [10] and, with the proper design, will possess bandgaps for certain frequencies of sound as well as light [79, 80]. Due to the large difference between the speeds of sound and light

[9] In practice this can be complicated by other restrictions on the structural dimensions. For example, one might only have high-quality material with a particular device layer thickness, fixing the thickness of the fabricated nanobeam and preventing arbitrary uniform scaling of the structure.

[10] Actually, this is not strictly true. As light can propagate in vacuum it is possible to have a photonic crystal composed of an array of disconnected dielectric objects. As phonons cannot propagate through vacuum any acoustic waves would be unaffected by the periodicity. Rather, each object would behave as an independent mechanical oscillator. In this work we are only concerned with suspended, connected membranes with periodic holes and we do not need to worry about this other type of photonic crystal.

such bandgaps will occur at frequencies orders of magnitude lower than the optical bandgaps, in the GHz range.

The mathematics of phononic crystals are slightly more complicated than those of photonic crystals, as the transverse and longitudinal components of the acoustic displacement field can mix. Considering a linear, lossless elastic medium, we may again represent an arbitrary solution of the acoustic wave equations as a linear superposition of time-harmonic solutions $\boldsymbol{Q}(\boldsymbol{x},\omega) = \boldsymbol{Q}(\boldsymbol{x})e^{-i\omega t}$, with the displacement field vector $\boldsymbol{Q}$ given by a solution to the eigenvalue problem [81]:

$$\hat{\Phi}\boldsymbol{Q}(\boldsymbol{x}) = \frac{1}{\rho(\boldsymbol{x})}\nabla^{\mathrm{T}}\left(\boldsymbol{C}(\boldsymbol{x})\nabla\boldsymbol{Q}(\boldsymbol{x})\right) = -\omega^2\boldsymbol{Q}(\boldsymbol{x}), \tag{1.67}$$

where ρ and $\boldsymbol{C}$ are the scalar material density and rank four elasticity tensor, respectively, and where ∇^{T} denotes the transpose of the nabla operator. This can be equivalently written in component form as

$$\frac{1}{\rho(\boldsymbol{x})}\partial_{\mathrm{j}}\left(C_{\mathrm{ijkl}}(\boldsymbol{x})\partial_{\mathrm{l}}Q_{\mathrm{k}}(\boldsymbol{x})\right) = -\omega^2 Q_{\mathrm{i}}(\boldsymbol{x}), \tag{1.68}$$

where repeated subscript indices are summed over. In contrast to the electromagnetic eigenvalue equation we have treated the elastic material as explicitly anisotropic. This is in general necessary even if the dielectric tensor is isotropic, as is the case with Si. As before, in the presence of discrete translational symmetry we may express the solutions to the acoustic wave equation as Bloch modes and numerically solve for the phononic band structure. In this case we use COMSOL [82], a commercial finite-element-method (FEM) solver, using the appropriate periodic boundary conditions to solve for the desired number of bands at any desired point within the IBZ [11].

The computed phononic band structure for the nanobeam unit cell is shown in Fig. 1.7, assuming the elastic properties of Si and the same unit cell dimensions used to compute the photonic band structure in Fig. 1.6b. As with the electromagnetic solutions we may classify the modes by their vector symmetries with respect to reflection in the y and z directions. In this case we care about the acoustic modes with even vector symmetry in both directions, for reasons which will be made clear in section 1.3.3. These modes are shown as red lines, while other symmetries are displayed as blue lines. The most striking difference from the photonic band structure is the absence of a light cone, as phonons cannot propagate in the cladding (vacuum) and thus radiation transverse to the beam is impossible. This in principle allows for the formation of complete phononic bandgaps. In this structure, however, we have no true gaps but rather only gaps for particular symmetries. The relevant pseudo-gap (sometimes called a symmetry gap) for the even symmetry modes is shaded in red, and contains several bands of differing symmetry. In a perfectly symmetric structure this is of no

[11] One might reasonably ask why we don't use FEM to solve for the photonic band structure. The reason is that, due to the condition of transversality imposed by Maxwell's equations ($\nabla\cdot\boldsymbol{H} = 0$, or for a Bloch mode $(i\boldsymbol{k}+\nabla)\cdot\boldsymbol{h}_{\boldsymbol{k}} = 0$) it is more efficient to use a planewave basis for finding physical solutions to the eigenvalue equation [75]. As acoustic waves do not have this constraint, it is more efficient to solve them using FEM. This also affords several advantages such as the use of a nonuniform mesh, allowing us to resolve fine features in the geometry more effectively.

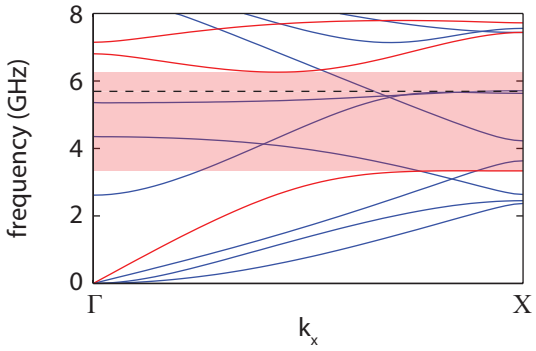

Figure 1.7: **Nanobeam mechanical bandstructure.** Computed acoustic band structure for a nanobeam with a unit cell like that of Fig. 1.6a, showing mode frequency versus the axial component of the wavevector k_x for z-symmetric modes, assuming a Si beam with the same dimensions as Fig. 1.6b. Bands shown in red are modes of even vector symmetry in both y and z, while blue bands are other symmetries. The red shaded region indicates the bandgap, while the frequency of interest ($\nu \sim 5$ GHz) is indicated by a black dashed line.

consequence, as different symmetries are uncoupled from each other. In a real, fabricated structure, on the other hand, imperfections will always exist which break the perfect mirror symmetries of the structure and couple modes of different symmetry. While this coupling is typically not strong enough to prevent localization of a mechanical mode, it has important ramifications for the quality of the mechanical cavities formed from such a gap, discussed further below.

1.3.3 Cavity Design

Once a unit cell has been found with appropriate photonic and phononic bandgaps, we may form a resonant cavity by creating a defect in the lattice which locally perturbs the otherwise perfect crystal symmetry. For example, one might remove a hole from an otherwise uniform crystal structure, or even just change the dimensions of one unit cell slightly. The breaking of translational symmetry then allows for the existence of modes at frequencies within the bandgap, which are necessarily localized in space around the defect region.

However, due to the possibility of radiating light into the cladding we must take care in how we design the defect in order to assure high optical quality. In particular, unlike a guided mode of the ideal, infinite photonic crystal, which exists at a single, well-defined point in $\boldsymbol{k}$-space, a spatially localized cavity mode will necessarily be composed of many individual momentum components. As with electronic defect states [83], the photonic defect modes can be approximately written in terms of the Bloch functions of particular band edge, $\boldsymbol{H}_{\mathrm{d}}(\boldsymbol{x}) \approx \sum_j f_{\mathrm{j}}(\boldsymbol{x})\ e^{-i\boldsymbol{k}_{\mathrm{j}}\cdot\boldsymbol{x}}\ \boldsymbol{h}_{\boldsymbol{k}_{\mathrm{j}}}$, where $\boldsymbol{k}_{\mathrm{j}}$ are the degenerate $\boldsymbol{k}$ points corresponding to the band edge (i.e., the star of $\boldsymbol{k}$ [74]), and $f_{\mathrm{j}}(\boldsymbol{x})$ are a set of envelope functions which can be described in terms of a Wannier-like equation dependent on the dielectric perturbation which forms the defect [84–87]. Thus, the decomposition of the defect mode

into its Fourier components will be peaked near the band edges, with a distribution dependent on the Fourier transforms of the envelope functions. As any momentum components located within the light cone may couple to radiation modes and contribute to optical loss, it is crucial to tailor the defect so that the envelope functions are as localized as possible in the Fourier domain [88, 89].

In a quasi-one-dimensional waveguide such as the nanobeam, this is straightforwardly accomplished by creating a smooth variation of the unit cell along the direction of propagation such that the frequency of one of the optical band edges at the X point of the IBZ ($k_{\mathrm{x}} = \pi/a$) tunes into the bandgap, creating a localized defect mode drawn from that band edge. For a perturbation which varies roughly quadratically the envelope function will be approximately Gaussian, with a width in $\boldsymbol{k}$-space inversely related to the spatial extent of the defect region. As discussed in section 1.3.1, we focus on the lowest lying band edge (the "valence" band), as a lower frequency corresponds to a larger distance in $\boldsymbol{k}$-space between the X-point and the light cone. Due to the aforementioned scaling properties of Maxwell's equations, the frequency of the lowest band can be tuned into the bandgap by reducing the local lattice constant a. At the same time the size of the hole is also reduced in an attempt to keep the local band structure constant aside from an overall frequency shift (this is not strictly possible, of course, without also scaling beam width and thickness). This results in a defect unit cell shown in Fig. 1.8b, which can be compared to the nominal unit cell in Fig. 1.8a. Fig. 1.8c shows the band edge from which we draw our mode, and the tuning of the band edge frequency as the nominal unit cell is deformed into the defect unit cell.

The behavior of the optical mode in the defect region is not the only concern. While the localized mechanical mode cannot suffer transverse radiation loss in the same manner as the optical mode, not all localized mechanical modes will couple strongly to the optical cavity mode. We must determine from which mechanical band we want to draw our localized mode to maximize the optomechanical coupling strength, and ensure the defect is designed in such a way as to tune this mode into the bandgap. In order to do so we must first derive a formula for computing g_0, given the unperturbed optical and mechanical modes.

There are two components in the computation of g_0. The first is determining the zero-point amplitude of the mechanical motion, given by $x_{\mathrm{ZPF}} = \sqrt{\frac{\hbar}{2m_{\mathrm{eff}}\omega_{\mathrm{m}}}}$, where the effective mass is defined [12] via an integral over the mode volume as

$$m_{\mathrm{eff}} = \rho \int d^3x |\boldsymbol{q}(\boldsymbol{x})|^2, \tag{1.69}$$

where $\boldsymbol{q}(\boldsymbol{x})$ is the displacement vector field of the mechanical mode, normalized such that $\max(|\boldsymbol{q}|) = 1$. The zero-point amplitude can thus be easily determined from the vector field profile of a mechan-

[12] This definition is determined by choosing the mechanical mode to be parametrized by an amplitude α such that $\boldsymbol{Q} = \alpha\boldsymbol{q}$, and requiring that the parametrized potential energy equal the true potential energy. In principle we could choose a different definition of mode amplitude α, which would result in a different definition for effective mass, or vice versa. We may choose one freely, but not both.

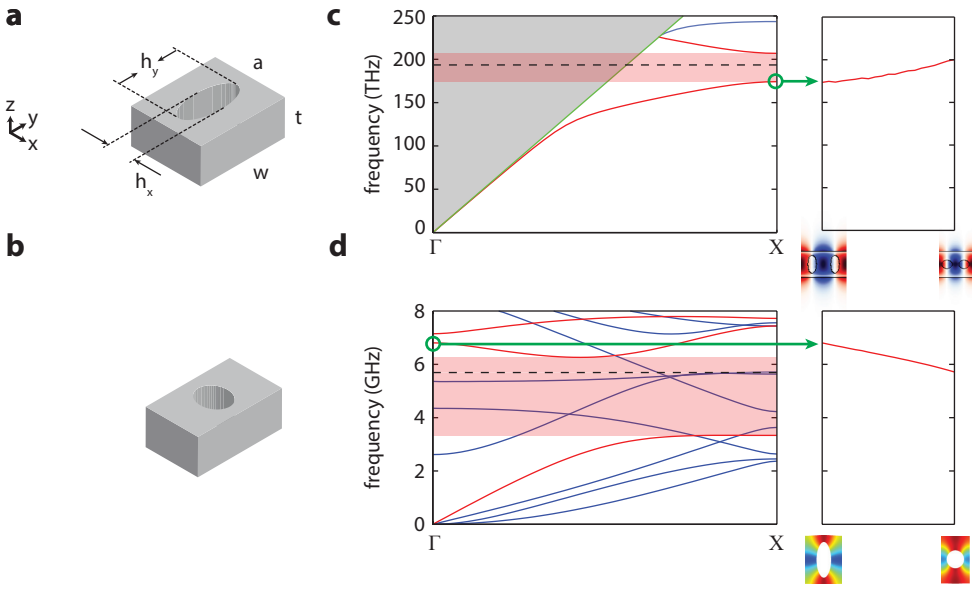

Figure 1.8: **Nominal and defect unit cells.** **a**, Nominal unit cell of the OMC nanobeam cavity, representing the mirror region. The dimensions are $(a, t, w, h_{\mathrm{x}}, h_{\mathrm{y}}) = (436, 220, 529, 165, 366)$. **b**, Unit cell in the center of the defect region, with dimensions $(a, t, w, h_{\mathrm{x}}, h_{\mathrm{y}}) = (327, 220, 529, 199, 170)$. **c**, Optical bandgap of the nominal unit cell, with the X-point band edge mode of interest circled in green. The plot on the side shows the tuning of the band edge frequency as the unit cell geometry is varied between the nominal and defect dimensions, along with corresponding plots of the dominant electric field component E_{y}. **d**, Acoustic bandgap of the nominal unit cell, with the Γ-point band edge mode circled in green. The side plot shows tuning of the mode in the defect region along with corresponding plots of the displacement field $\boldsymbol{Q}$.

ical eigenmode obtained via FEM simulation. The second component is determining the quantity $g_{\mathrm{OM}} = \partial\omega_{\mathrm{c}}/\partial\alpha$, the change in optical mode energy for an infinitesimal mechanical displacement amplitude α. From standard perturbation theory we know that the first order change in energy will depend on the unperturbed eigenvectors (in this case, the electric field mode). Since the energy density for an electromagnetic wave depends only on $\boldsymbol{E}$ and the dielectric tensor $\boldsymbol{\epsilon}$ this leads naturally to the expression

$$g_{\mathrm{OM}} = \frac{\partial\omega_{\mathrm{c}}}{\partial\alpha} = -\frac{\omega_{\mathrm{c}}}{2}\frac{\int d^3x\ \boldsymbol{E}^*(\boldsymbol{x})\cdot\frac{\partial\boldsymbol{\epsilon}(\boldsymbol{x})}{\partial\alpha}\boldsymbol{E}(\boldsymbol{x})}{\int d^3x\ \boldsymbol{E}^*(\boldsymbol{x})\cdot\boldsymbol{\epsilon}(\boldsymbol{x})\boldsymbol{E}(\boldsymbol{x})}, \tag{1.70}$$

which is recognized as a version of the standard Hellman-Feynman perturbation formula [90], and where the electric field vector is the unperturbed eigenmode. To first order, there are two fundamental ways in which the dielectric tensor can be perturbed by small displacements of the geometry. The first is that the material's internal dielectric properties remain unchanged while the boundaries of the object distort, altering the spatial overlap of the dielectric and electric field. This effect is often termed the "boundary contribution". Computation of this figure requires careful consideration of the boundary conditions imposed by Maxwell's equations in order to ensure the result is expressed in terms of functions which are continuous across the boundary. For an isotropic dielectric function it is given in terms of the optical and mechanical mode profiles as [91]

$$g_{\mathrm{OM,bnd}} = -\frac{\omega_{\mathrm{c}}}{2}\frac{\int dA\ \boldsymbol{q}(\boldsymbol{x})\cdot\hat{\boldsymbol{n}}(\boldsymbol{x})\left(\Delta\epsilon|\boldsymbol{E}_{\parallel}(\boldsymbol{x})|^2 - \Delta\epsilon^{-1}|\boldsymbol{D}_{\perp}(\boldsymbol{x})|^2\right)}{\int d^3x\ \epsilon(\boldsymbol{x})|\boldsymbol{E}(\boldsymbol{x})|^2}, \tag{1.71}$$

where the top integral is over all boundaries of the crystal while the bottom is over all space, $\hat{\boldsymbol{n}}$ is the surface normal vector at any boundary, $\boldsymbol{q}$ is the normalized mechanical displacement field, $\Delta\epsilon$ is the difference between the dielectric constant of the crystal and that of the cladding (here, $\Delta\epsilon = \epsilon_{\mathrm{Si}} - 1$), and $\boldsymbol{D} = \epsilon\boldsymbol{E}$ is the electric displacement field. The subscripts $\parallel$ and $\perp$ denote vector components parallel and perpendicular to the boundary, respectively. The second way mechanical motion can change $\boldsymbol{\epsilon}$ is via the photoelastic effect, where the internal strain caused by the deformation causes a shift in the internal dielectric properties of the material. This shift is given by [92]

$$\frac{d\boldsymbol{\epsilon}}{d\alpha} = \boldsymbol{\epsilon}\left(\frac{\boldsymbol{pS}}{\epsilon_0}\right)\boldsymbol{\epsilon}, \tag{1.72}$$

where $\boldsymbol{p}$ is the rank-four photoelastic tensor of the material and $\boldsymbol{S}$ is the strain tensor defined in terms of the displacement field as $S_{\mathrm{ij}} = \left(\frac{\partial q_{\mathrm{i}}}{\partial x_{\mathrm{j}}} + \frac{\partial q_{\mathrm{j}}}{\partial x_{\mathrm{i}}}\right)$. For an initially isotropic material this simplifies to

$$\frac{d\epsilon_{\mathrm{ij}}}{d\alpha} = -\epsilon_0\epsilon_{\mathrm{r}}^2 p_{\mathrm{ijkl}} S_{\mathrm{kl}}. \tag{1.73}$$

Note that the shifted dielectric tensor is in general anisotropic even if the unperturbed material is isotropic. This expression can be inserted into Eq. 1.70 to obtain a straightforward, if somewhat

cumbersome, expression for the photoelastic contribution

$$g_{\mathrm{OM,pe}} = \frac{\epsilon_0 \epsilon_{\mathrm{r}}^2 \omega_{\mathrm{c}}}{2} \frac{\int dV\ E_{\mathrm{i}}^*(\boldsymbol{x}) E_{\mathrm{j}}(\boldsymbol{x}) p_{\mathrm{ijkl}} S_{\mathrm{kl}}(\boldsymbol{x})}{\int d^3x\ \epsilon(\boldsymbol{x}) |\boldsymbol{E}(\boldsymbol{x})|^2}, \tag{1.74}$$

where the top integral is over the volume of the crystal, and all repeated subscript indices are summed over. In practice this is greatly simplified by the underlying microscopic crystal symmetries of the dielectric, which require most elements of the photoelastic tensor to vanish [92]. In particular, for Si there exist only three independent, non-zero elements of $\boldsymbol{p}$, which can be obtained at the relevant optical frequencies from Ref. [93]. Putting it all together, the vacuum optomechanical coupling rate is given by the expression

$$g_0 = x_{\mathrm{ZPF}} \left(g_{\mathrm{OM,bnd}} + g_{\mathrm{OM,pe}} \right), \tag{1.75}$$

which can be calculated using only the unperturbed mechanical and optical eigenmodes (obtained via FEM simulation) and the known elastic and optical properties of the material.

With an explicit expression for g_0 in hand we can immediately eliminate most mechanical modes from consideration. First, notice that g_0 is quadratic in the electric field but linear in either the displacement $\boldsymbol{q}$ or the strain tensor $\boldsymbol{S}$. Consequently, it follows that any mechanical modes possessing odd vector symmetry with respect to either mirror symmetry plane of the structure will have a vanishing contribution to g_0 when integrated over all space. Moreover, recall that at the X-point the guided mode will have an overall phase factor $e^{i\pi x/a}$. That is, the guided mode profiles in adjacent unit cells will be π out of phase with each other. Thus, $\boldsymbol{q}$ and $\boldsymbol{S}$ will flip sign while $|\boldsymbol{E}|^2$ stays constant. This implies that mechanical modes drawn from the X-point will have negligible g_0 due to a high degree of cancellation between adjacent unit cells in the cavity region. The only remaining mechanical modes which are expected to yield high g_0 are those at the Γ-point ($k_{\mathrm{x}} = 0$) with completely even vector symmetry, from which we must draw our mechanical cavity mode.

We can estimate the expected relative optomechanical coupling strengths of each mode by computing g_0 per unit cell using the guided mode optical and mechanical profiles. We find that the optimal mechanical mode is the second even band, circled in green in the mechanical band structure shown in Fig. 1.8d. Note that this mode is higher in frequency than the bandgap, and our defect must tune the Γ-point frequency down into the gap. Unfortunately, as the elastic wave equations possess the same type of scale invariance as Maxwell's equations, the reduction of lattice constant a and hole area will tend to raise the frequency of the mechanical bands as well as the optical bands. Unlike the optical bands, however, the mechanical bands are highly sensitive to the hole shape. In particular, by increasing the hole aspect ratio $h_{\mathrm{x}}/h_{\mathrm{y}}$ while keeping the total hole area constant we effectively raise the motional mass of the mechanical mode of interest, reducing its frequency, while minimally perturbing the optical modes. Adding this effect yields the full defect cell shown in Fig. 1.8b, and tunes the mechanical mode in the desired direction, as shown in Fig. 1.8d.

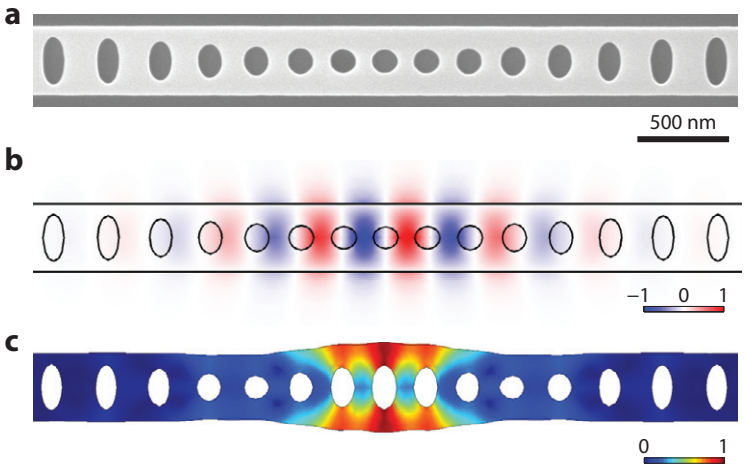

Figure 1.9: **Optimized nanobeam OMC cavity.** **a**, SEM image of a fabricated nanobeam cavity, showing the defect cavity region. In addition to this, nominal ("mirror") unit cells extend for several periods on either side. **b**, Normalized electric field E_y component of the fundamental optical mode at $\omega_c/2\pi \approx 194$ THz. **c**, Normalized displacement profile of the mechanical "breathing" mode at $\omega_m/2\pi \approx 5.1$ GHz.

With the optimal modes identified, and a suitable defect perturbation in hand, creating a good optomechanical cavity now comes down to choosing a gradual, smooth transition between the nominal and defect unit cells in order to reduce optical scattering and match the spatial extents of the optical and mechanical defect modes (as can be inferred from the above expressions for g_0, roughly speaking we need the extent of $\boldsymbol{q}$ to be comparable to that of $|\boldsymbol{E}|^2$). Moreover, there is still a fair amount of optimization of the specific dimensions of both the nominal and defect unit cell, which will provide small but important changes to the final value of g_0 for a given cavity. At this point, it proves more efficient to rely on numerical optimization, and we use the Nelder-Mead simplex algorithm [94] in order to optimize the cavity over a range of dimensional parameters. At each step the localized optical and mechanical eigenmodes are computed using COMSOL and the resulting calculated value of g_0 is used as a fitness function for the optimization, with an additional penalty applied if the optical Q drops too low. The full details of the optimization procedure can be found in Ref. [1].

1.3.4 Nanobeams

A scanning electron microscope (SEM) image of a fabricated, fully optimized Si nanobeam OMC is shown in Fig. 1.9, with the corresponding FEM-simulated electric field and displacement mode profiles shown in Fig. 1.9b and c, respectively. For this design, the nominal unit cell (which acts as the cavity mirror) has the dimensions specified in the caption of Fig. 1.8a, while the unit cell at the center of the defect region has the dimensions specified in the caption of Fig. 1.8b. In

between these two points the lattice constant and hole area are varied quadratically over several holes, with a smooth variation of the hole aspect ratio as described in the previous section. Though the image displays only the cavity region and a single mirror period on either end, in practice several more periods of the mirror region are present on either end of the cavity in order to provide good confinement of the modes ($6-7$ periods on either side is usually more than sufficient).

The as-designed optical and mechanical mode frequencies are $\omega_{\rm c}/2\pi \approx 194$ THz and $\omega_{\rm m}/2\pi \approx 5.1$ GHz, respectively, while in practice small fabrication errors can result in small variations in the measured frequencies. The optimized optomechanical coupling for this structure is $g_0 = 860$ kHz, of which 950 kHz is due to the photoelastic effect and -90 kHz is due to the moving boundary effect. While this is particular design has the maximum as-designed g_0, it is possible to find other designs with only slightly lower g_0. In particular, a design with $\omega_{\rm m}/2\pi \approx 3.6$ GHz, but a simulated g_0 which is still on the order of $700-800$ kHz, is used in the measurements of Chapter 3 as it allows us to avoid a prominent phase noise peak of the laser at ~ 5 GHz.

The optical quality factor $Q_{\rm c} = \omega_{\rm c}/\kappa$, as calculated in COMSOL, reaches roughly 10^8. However, this only accounts for coupling due to radiation losses in an ideal structure. In a typical fabricated structure the real measured intrinsic optical quality factor lies in the range $5\times10^5 - 10^6$, limited by a number of factors including disorder, surface roughness, and optical absorption within the Si. The highest intrinsic Q measured in such a nanobeam is 1.22×10^6 ($\kappa_{\rm i}/2\pi \approx 160$ MHz) [95], which is within an order of magnitude of the highest recorded optical Q in *any* Si photonic crystal cavity [70].

The mechanical quality factor $Q_{\rm m} = \omega_{\rm m}/\gamma_{\rm i}$ is another story, as it depends strongly on temperature. At room temperature the mechanical Q is typically $Q_{\rm m} \approx 2000-3000$. This is limited by material damping processes such as phonon-phonon interactions [96] and thermoelastic damping [97], and shows little dependence on the quality of the acoustic bandgap in the mirror section or on any additional acoustic shielding as described below. Detailed modeling of the mechanical damping at high temperature can be found in Ref. [1]. As the temperature is lowered, the effects of these intrinsic damping mechanisms are expected to reduce and radiation of acoustic energy through the beam (sometimes called "clamping loss") should begin to dominate the mechanical damping. However, one finds $Q_{\rm m} \approx 10^4$ even at low temperatures ($T_{\rm b} \lesssim 10$ K), with little increase as the number of mirror periods is increased. The culprit is the lack of a complete phononic bandgap in the nanobeam, as discussed in section 1.3.2. As the mirror region only possesses a symmetry bandgap, small amounts of fabrication-induced disorder can couple the mechanical cavity mode to nearby modes of different symmetry for which no bandgap exists, limiting the ability of the mirror region to suppress phonon radiation. This problem, first studied for nanobeams in Ref. [98], can be remedied by surrounding the cavity with an additional phononic crystal possessing a complete bandgap.

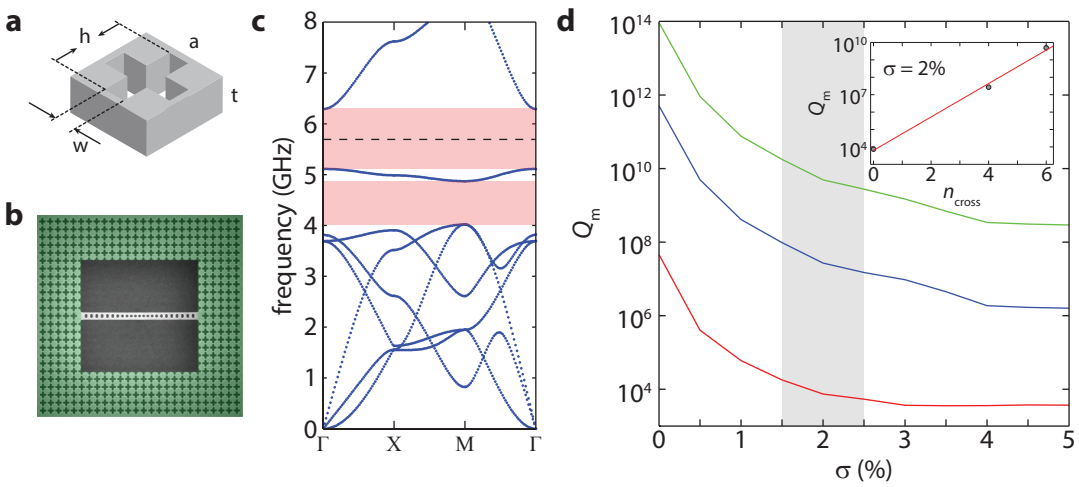

Figure 1.10: **Cross phononic crystal.** **a**, Unit cell of the 2D "cross" phononic crystal, with relevant dimensional parameters labeled. This crystal is periodic in both x and y, forming a square lattice (see Fig. 1.5b). **b**, SEM image of a fabricated nanobeam cavity with a surrounding cross pattern (green) serving as an acoustic radiation shield. **c**, Complete phononic band structure in Si for $(a, t, h, w) = (534, 220, 454, 134)$ nm. The bandgaps (red shaded region) are complete phononic bandgaps, and the nanobeam mechanical mode frequency is indicated by a black dashed line. **d**, Radiation limited mechanical quality factor Q_{m} obtained from FEM simulations as a function of disorder parameter σ, for a nanobeam without (red) a phononic radiation shield and with a shield consisting of four (blue) and six (green) cross periods in all directions. The expected range of disorder estimated from SEM images of fabricated devices is represented by the gray shaded region. The inset shows the dependence of Q_{m} on the number of cross periods n_{cross} for $\sigma = 2\%$.

1.3.5 Crosses

Generally speaking, the necessary requirement for obtaining a complete phononic bandgap is to find a crystal comprised of large masses connected by thin sections of material [80]. For low frequency modes, where the phonon wavelength is large compared to the lattice constant, the crystal can be approximated as a homogeneous effective medium, the stiffness of which is greatly reduced compared to the unpatterned material due to the weak links between the masses. For high frequency modes, on the other hand, the crystal may be described by a tight binding model where the interaction between the masses is reduced by the small links. The net result is that the low lying modes are pushed down in frequency while the dispersion of the high frequency modes flattens out, creating a bandgap for frequencies where the phonon wavelength is comparable to the lattice constant.

One of the simplest and most effective crystals of this type is the so-called "cross" crystal, the unit cell of which is illustrated in Fig. 1.10a. This crystal, which consists of cross-shaped holes arrayed in a two-dimensional square lattice (see Fig. 1.5b) is placed around the entire nanobeam cavity, providing an additional shield for acoustic radiation. A simulated phononic band structure of a cross crystal is shown in Fig. 1.10c, in which we observe a large (~ 1 GHz) complete phononic gap centered around the mechanical resonance frequency of the nanobeam.

While this structure is still affected by disorder, the large gap between the frequency of interest and the band edges and the complete nature of the bandgap minimize the effects of small amounts of disorder. The radiation limited Q_{m} can be simulated in COMSOL by placing an additional region around the shielded cavity, in which the elasticity tensor has a small imaginary component which mimics acoustic absorption. By tuning the size and position dependence of this fictitious absorption this layer becomes the acoustic equivalent of a perfectly matched layer [99], which provides an absorbing boundary condition while minimizing spurious reflections of mechanical energy. The full details of how the PML is tuned can be found in Ref. [98]. To introduce disorder into the simulation, we randomize the dimensions and positions of the holes about the nominal parameters according to a normal distribution with a specified standard deviation. For simplicity, the standard deviation for each parameter is chosen to be a fixed percentage σ of its nominal value. In Fig. 1.10d we show the simulated Q_{m} for a nanobeam without a cross shield (red) and with four (blue) and six (green) periods of the cross shield in all directions, as a function of the disorder parameter σ, with the gray shaded region indicating the range of disorder ($\sigma \sim 2\%$) expected in actual fabricated devices based on SEM images. A plot of Q_{m} versus the number of cross periods n_{cross} for $\sigma = 2\%$ is shown in the inset, showing the expected exponential dependence with $Q_{\mathrm{m}} \approx 10^{3.8+0.96 n_{\mathrm{cross}}}$. As we shall see in Chapters 3 and 4, the observed mechanical Q at millikelvin temperatures agrees well with the simulated radiation-limited values in the presence of moderate amounts of disorder, suggesting we can reach the regime in which the mechanical lifetime is determined solely by the phononic engineering of the structure.

1.3.6 Snowflakes

One might reasonably ask why we don't just make an OMC defect cavity using the 2D cross structure, as its complete phononic bandgap is superior in terms of acoustic confinement when compared to the symmetry gap of the nanobeam. Unfortunately, the cross structure has no photonic pseudo-gap for the TE-like modes due to the presence of leaky modes, and does not even possess a guided mode gap for the TM-like modes [80]. This arises from the low symmetry of the cross structure's square lattice, which leads to very different behavior for plane waves propagating along different directions of the crystal (e.g., the X and M directions, as seen in the IBZ diagram of Fig. 1.5b), resulting in a correspondingly smaller photonic bandgap.

This problem can be remedied by moving to a lattice with a higher degree of symmetry, such as the hexagonal lattice illustrated in Fig. 1.5c. Extending the ideas of the cross structure to this lattice symmetry results in the "snowflake" crystal [80], the unit cell of which is shown in Fig. 1.11a. With proper tuning of the dimensional parameters, we can obtain simultaneous photonic pseudo-gaps and complete phononic bandgaps, as illustrated in Figs. 1.11b and c.

Given the 2D nature of the lattice, there is more freedom in how the defect cavity is formed.

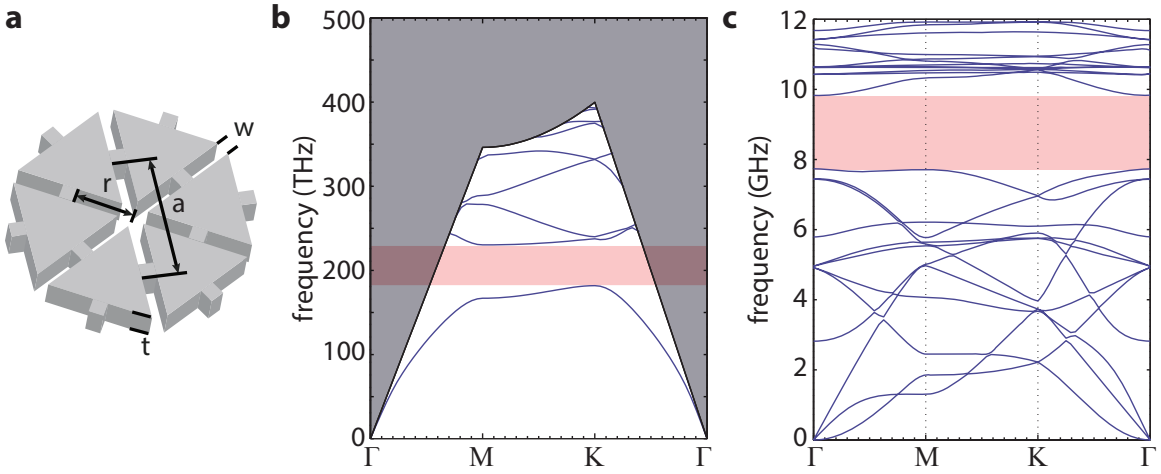

Figure 1.11: **Snowflake optomechanical crystal. a**, Unit cell of the 2D "snowflake" crystal, with relevant dimensional parameters labeled. **b**, Computed photonic band structure for the TE-like modes in Si, for $(a, r, w, t) = (500, 200, 75, 220)$ nm, with the light cone (gray) and pseudo-gap (red) regions indicated. **c**, Computed phononic band structure for the same dimensions as **b**, with the complete phononic gap indicated by the red shaded region.

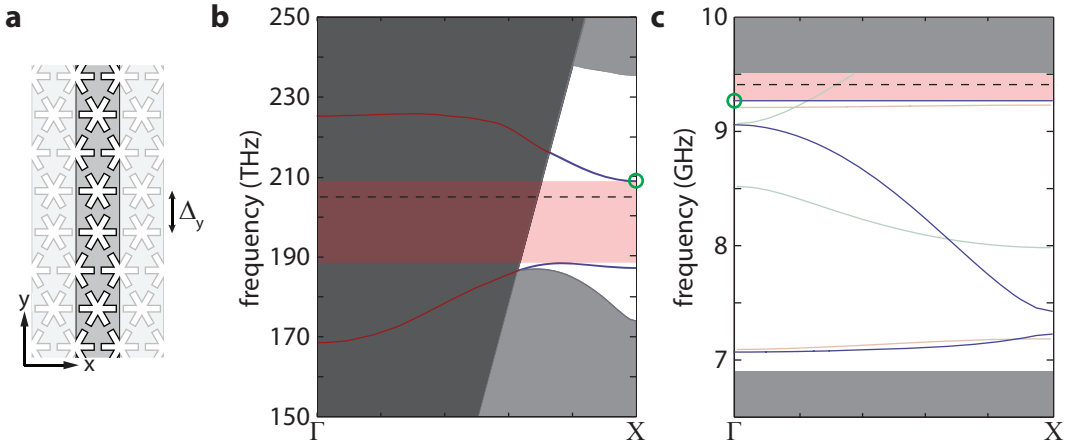

Figure 1.12: **Snowflake $W1$ waveguide. a**, Unit cell of the $W1$ waveguide formed from a snowflake lattice. One row of holes is removed and the remaining holes are shifted inward by an amount W, yielding a waveguide width of $\Delta_y = \sqrt{3}a - 2W$. Photonic and phononic band structures for the effective 1D waveguide are shown in **b** and **c**, respectively, assuming the lattice parameters detailed in Fig. 1.11 and $W = 200$ nm ($\Delta_y = 400$ nm). The bandgaps are shaded in red and the optical and mechanical cavity frequencies indicated by black dashed lines. Green circles indicate the optical X-point and mechanical Γ-point modes from which the cavity modes are drawn. The dark gray region of the photonic band structure indicates the light cone (radiation out of plane), while the light gray regions in both band structures indicate frequencies above and below the in-plane bandgap defining the waveguide (radiation in-plane).

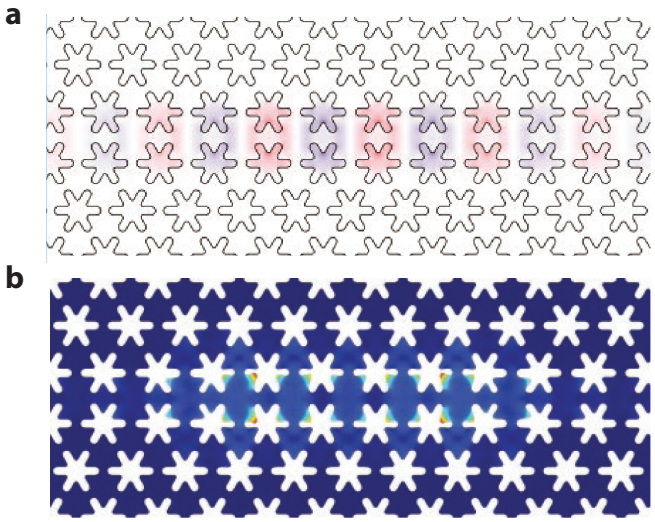

Figure 1.13: **Heterostructure snowflake OMC cavity.** FEM simulated optical (**a**) and mechanical (**b**) mode profiles of the $W1$ snowflake cavity.

For our purposes we choose to form a double heterostructure cavity, which has been demonstrated to yield extremely high optical Qs in similar hexagonal lattice photonic crystals [100, 101]. In this type of cavity we first remove a single row of holes from the lattice, forming a line defect which serves as a waveguide [102, 103]. This reduces the problem to that of a 1D waveguide, similar to the nanobeam, with an effective unit cell shown in Fig. 1.12a. The corresponding band structures for this waveguide are shown in Fig. 1.12b and c. In this case the remaining holes in the lattice are shifted inward towards the center of the line defect in order to tune the optical and mechanical waveguide bands and obtain the desired band structures [80].

Forming a defect cavity now becomes very similar to the case of the nanobeam, with a gradual modulation of the geometry introduced along the waveguide axis in order to locally tune the desired optical X-point and mechanical Γ-point modes into the bandgap. In this case, a small ($\sim 3\%$) reduction in the hole radius is gradually introduced, leading to the localized optical and mechanical modes shown in Fig. 1.13a and b, respectively. The optical and mechanical resonance frequencies of this cavity are $\omega_c/2\pi = 205$ THz and $\omega_m/2\pi = 9.5$ GHz, respectively, with a vacuum optomechanical coupling rate of $g_0/2\pi \approx 300$ kHz.

Though snowflake cavities are not be used in the main experimental work presented in this thesis, their two-dimensional nature presents an attractive solution the problem of excess heating at sub-kelvin temperatures, as discussed further in Chapter 5.

1.3.7 Fabrication Overview

The Si devices studied in this work are fabricated from a silicon-on-insulator (SOI) wafer (SOITEC, 220 nm device layer, 3 μm buried oxide) using electron beam lithography followed by reactive ion etching (RIE/ICP) with a C_4F_8/SF_6 chemistry.

For the end-fire coupled devices detailed in section 2.5 and measured in chapters 3 and 4, the Si device layer is then masked using ProTEK PSB photoresist to define a mesa region of the chip. Outside of the protected mesa region, the buried oxide is removed with a plasma etch and a trench is formed in the underlying silicon substrate using tetramethylammonium hydroxide (TMAH), allowing a lensed fiber tip to be brought into the near-field of the end-fire couplers.

Finally the devices are released in hydrofluoric acid (49 % aqueous HF solution) and cleaned in a piranha solution (3-to-1 H_2SO_4:H_2O_2) before a final hydrogen termination in diluted HF in order to passivate the surface [104]. During the initial lithography step, arrays of the nominal design are scaled by ± 2 % to account for frequency shifts due to fabrication imperfections and disorder.

Further details about the fabrication process can be found in Refs. [44] and [55].

Chapter 2

Optical Coupling in a Dilution Refrigerator

2.1 Motivation

To couple light in and out of the optical cavity, the initial nanobeam OMC experiments [13,40,45,105] utilized an evanescently coupled, dimpled fiber taper probe developed within the Painter group [106]. This coupling scheme is useful for rapid development and testing of devices, as it allows one to quickly test hundreds of devices spread across a single chip. However, it suffers from a number of drawbacks. The coupling scheme is inherently two-sided, meaning that light can couple from the cavity into the forwards and backwards propagating waveguide modes with equal efficiency. In practice we only detect light at one end of the optical fiber, and thus this coupling scheme limits us to a coupling efficiency of 50% even in the ideal case of no additional intrinsic loss in the cavity ($\kappa_\mathrm{i} = 0$). In a realistic cavity the actual coupling efficiency is typically much less than 50% due to finite intrinsic loss from fabrication imperfections, absorption, etc. Moreover, the exact extrinsic coupling rate κ_e depends very sensitively on the positioning of the fiber relative to the cavity, which can be difficult to achieve consistently. The taper itself can in some cases cause additional scattering of the optical cavity mode, increasing κ_i without increasing κ_e, reducing the overall coupling efficiency.

In principle these problems can be overcome by utilizing a slightly different coupling scheme, where the dimpled taper is coupled to an intermediate tapered waveguide which subsequently directs light to the optical cavity [107–109]. By separating the fiber-to-chip coupling efficiency and the cavity coupling efficiency, this method allows for single-sided coupling where in principle 100% of the cavity light can be directed to the output channel and the cavity coupling rate κ_e can be precisely controlled independent of the fiber coupling efficiency. However, both this scheme and the direct fiber-to-cavity coupling require the careful positioning of the fiber taper in the near-field of a nanostructure. As no signal from the cavity is observed until the fiber has been brought to within a few microns of the optimal position, this coupling scheme requires real-time imaging of

the device during the coupling process. This is difficult to achieve in a dilution refrigerator due to issues with heating of the environment by the imaging system and alignment issues owing to thermal contraction at low temperatures of the various materials comprising the imaging setup. It is possible to construct imaging systems suitable for real-time imaging at mK temperatures utilizing optical fiber bundles to transmit images to a room-temperature camera [110], and coupling of a fiber taper to an optomechanical bottle resonator inside a dilution fridge has recently been demonstrated using this technique [111]. However, a simpler approach is to develop a coupling scheme that does not utilize a fiber taper at all and can be aligned without direct imaging of the sample.

One alternative approach is to use an on-chip grating coupler [112–118] to couple a near-normal incidence free-space beam (e.g., launched from a lensed optical fiber) to an on-chip waveguide. While in principle such couplers can reach efficiencies of greater than 90% [113], typical experimentally demonstrated values are closer to $20-30\%$, though this can be improved to $60-70\%$ via nonuniform etching or modification of the substrate [115–117]. More importantly, such grating couplers typically require leaving the buried oxide substrate intact, which complicates integration with the fully undercut OMC structures used in this work. While progress has been made in developing high efficiency, fully suspended grating couplers [118], they have only demonstrated efficiencies of $\sim 25\%$ for the TE-like modes used in our OMC designs. There were some initial efforts to design and fabricate novel suspended free-space couplers within the Painter group [1], but the fabricated couplers possessed rather low coupling efficiencies ($\sim 20\%$) and were difficult to characterize owing to multiple reflections within the coupler structure that greatly complicated the optical reflection spectrum.

Ultimately, the approach which proved most successful was to design couplers consisting of either a lensed or cleaved optical fiber aligned on-axis to a narrow, mode-matched on-chip waveguide, which tapers up to the appropriate width for coupling into the optical cavity. Couplers of this type have been widely used for efficient fiber-to-chip coupling and on-chip optical interconnects [119–122], and allow for single-sided coupling with demonstrated fiber-to-cavity efficiencies as high as 75%. Most importantly, it is straightforward to couple such devices without imaging the sample, and in some cases without any in-situ alignment at all, as we will discuss below.

2.2 Adiabatic Mode Coupling Theory

In this section, we will present an overview of the mode coupling theory underlying the design of the fiber-to-waveguide couplers used in this work. A more rigorous and complete treatment of the theory presented here can be found in a number of textbooks on optical waveguide theory, particularly Refs. [123] and [124].

2.2.1 Waveguide Modes

We begin by considering a waveguide formed from lossless, isotropic dielectric materials, where the axis of the waveguide will be taken as the z-axis, with an arbitrary cross-sectional refractive index profile $n(x, y)$ but homogeneous along the z-axis. We consider only monochromatic fields, so that all electric and magnetic fields have an assumed time-dependence of $e^{i\omega t}$. Assuming nonmagnetic materials, Maxwell's wave equations for the electric and magnetic field vectors $\boldsymbol{E}$ and $\boldsymbol{H}$ are

$$\nabla \times \boldsymbol{E} = -i\omega\mu_0 \boldsymbol{H}, \tag{2.1}$$

$$\nabla \times \boldsymbol{H} = i\omega\epsilon_0 n^2 \boldsymbol{E}. \tag{2.2}$$

In general, the longitudinal components of the fields $\boldsymbol{E}_\mathrm{z}$ and $\boldsymbol{H}_\mathrm{z}$ can be determined from the transverse components. In particular, it is straightforward to show using Eqs. 2.1 and 2.2 that

$$\boldsymbol{E}_\mathrm{z} = \left(i\omega\epsilon_0 n^2\right)^{-1} \nabla_\mathrm{t} \times \boldsymbol{H}_\mathrm{t}, \tag{2.3}$$

$$\boldsymbol{H}_\mathrm{z} = -\left(i\omega\mu_0\right)^{-1} \nabla_\mathrm{t} \times \boldsymbol{E}_\mathrm{t}, \tag{2.4}$$

where $\nabla_\mathrm{t} = \nabla - \hat{\boldsymbol{z}}\partial/\partial z$ and $\hat{\boldsymbol{z}}$ is the unit vector in the z-direction. This simple relation allows us to focus solely on the transverse field components, which is advantageous when considering junctions between two waveguides along the waveguide axis as the transverse components are continuous across the boundary.

Due to the assumed continuous translational symmetry of the system along the z-axis, the solutions of Maxwell's equations can be expressed as $\boldsymbol{E} = \boldsymbol{E}(x, y)e^{i\beta z}$ for some propagation constant β (similarly for $\boldsymbol{H}$). The propagation constant is often written in terms of an effective refractive index, $\beta = (\omega/c)n_\mathrm{eff}$. Any such solution $\boldsymbol{E}_\mathrm{j}$ will satisfy the following wave equations for the transverse field vectors:

$$\frac{-1}{i\omega\mu_0}\nabla_\mathrm{t} \times (\nabla_\mathrm{t} \times \boldsymbol{E}_\mathrm{j}) + i\beta\,(\hat{\boldsymbol{z}} \times \boldsymbol{H}_\mathrm{j}) = i\omega\epsilon_0 n^2 \boldsymbol{E}_\mathrm{j}, \tag{2.5}$$

$$\frac{1}{i\omega\epsilon_0}\nabla_\mathrm{t} \times \left(\frac{1}{n^2}\nabla_\mathrm{t} \times \boldsymbol{H}_\mathrm{j}\right) + i\beta\,(\hat{\boldsymbol{z}} \times \boldsymbol{E}_\mathrm{j}) = -i\omega\mu_0 \boldsymbol{H}_\mathrm{j}. \tag{2.6}$$

These solutions can be further divided into guided modes and radiation modes. Guided modes are defined as those modes for which energy is confined to a finite region in the xy-plane and does not propagate freely to infinity. There exists a discrete set of these modes, with real effective indices satisfying

$$n^2(\infty) < n_\mathrm{eff}^2 < \max\left(n^2(x, y)\right). \tag{2.7}$$

The specific number of guided modes depends on the specific refractive index distribution, but there

will always be at least one guided mode whenever $n(x,y) > n(\infty)$ for some x and y in the plane. There will also exist an infinite continuum of unconfined modes, the so-called radiation modes, which carry energy away to infinity and for which Eq. 2.7 need not be satisfied (and n_{eff} need not be real). The guided and radiation modes together form a complete set of modes, so that an arbitrary transverse field can be expanded in terms of the guided modes and the radiation field [123]:

$$\boldsymbol{E}_{\text{t}}(x,y,z) = \sum_{\text{j}} \left(a_{\text{j}}^{+} e^{i\beta_{\text{j}} z} + a_{\text{j}}^{-} e^{-i\beta_{\text{j}} z} \right) \boldsymbol{E}_{\text{j}}(x,y) + \int d\rho \left(a_{\rho}^{+} e^{i\beta_{\rho} z} + a_{\rho}^{-} e^{-i\beta_{\rho} z} \right) \boldsymbol{E}_{\rho}(x,y), \tag{2.8}$$

$$\boldsymbol{H}_{\text{t}}(x,y,z) = \sum_{\text{j}} \left(a_{\text{j}}^{+} e^{i\beta_{\text{j}} z} - a_{\text{j}}^{-} e^{-i\beta_{\text{j}} z} \right) \boldsymbol{H}_{\text{j}}(x,y) + \int d\rho \left(a_{\rho}^{+} e^{i\beta_{\rho} z} - a_{\rho}^{-} e^{-i\beta_{\rho} z} \right) \boldsymbol{H}_{\rho}(x,y), \tag{2.9}$$

where the summation over j extends over all possible guided modes, and the integral over ρ over the continuum of radiation modes. Note that we have explicitly separated forward traveling guided modes (superscript + on the expansion coefficients) from backward traveling guided modes (superscript − on the expansion coefficients) due to the different symmetry properties of the electric and magnetic transverse fields upon reflection in the z direction.

In addition to forming a complete set, the transverse guided and radiation modes of the waveguide are orthogonal with respect to power. That is, for a lossless, source-free (i.e., no free charge or current) waveguide it can be shown that any two distinct guided or radiation modes of the waveguide (i.e., modes with distinct propagation constants) with vector field profiles $\boldsymbol{E}_{\text{j}}$ and $\boldsymbol{E}_{\text{k}}$ obey the orthogonality relation [123]:

$$\int dA \left(\boldsymbol{E}_{\text{j}} \times \boldsymbol{H}_{\text{k}}^{*} \right) \cdot \hat{\boldsymbol{z}} = \int dA \left(\boldsymbol{E}_{\text{j}}^{*} \times \boldsymbol{H}_{\text{k}} \right) \cdot \hat{\boldsymbol{z}} = 0, \quad j \neq k, \tag{2.10}$$

where the integration is over the entire xy-plane. For $j = k$ the real part of the integrals in Eq. 2.10 are proportional to the modal power flow in the z-direction. We also note that the orthogonality relation only involves the transverse field components.

Strictly speaking, waveguide modes as we have described above only exist for an ideal infinite waveguide (i.e., a waveguide with perfect translational invariance). In real waveguides this assumption will always be violated to some degree, and in particular the couplers we will discuss in this section involve waveguides where the cross-sectional refractive index profile $n(x,y)$ is intentionally varied along the axis of the waveguide. Such waveguides cannot support the modes described above. However, as the ideal waveguide modes form a complete set, at any point $z = z_0$ along the waveguide we may expand the real propagating field in terms of the modes of the corresponding infinite waveguide (that is, the waveguide where $n = n(x,y,z_0)\ \forall\ z$). If the index variation is slow enough along the z-axis it is intuitively obvious that these modes will provide an approximation to the real solutions of Maxwell's equations in some local region about z_0. As we will see below this local mode formalism may be used to determine how slowly a waveguide must change in order to transmit light

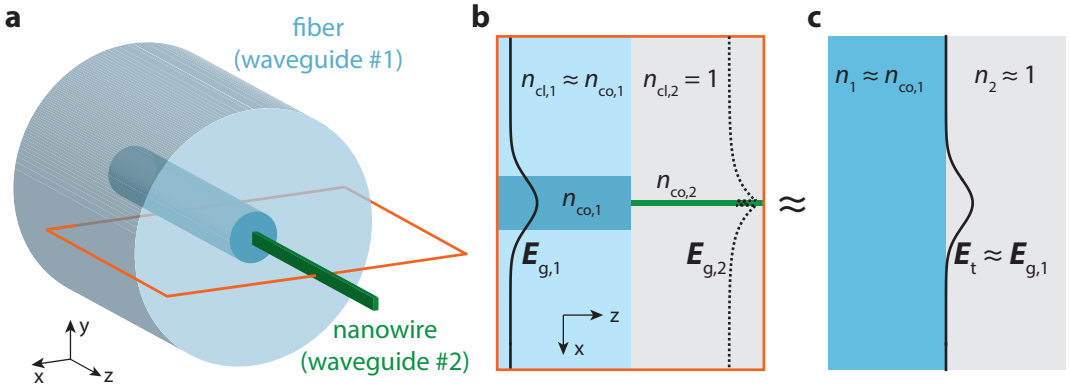

Figure 2.1: **Fiber-to-waveguide coupling.** **a**, Optical fiber aligned with a nanowire waveguide (not to scale). **b**, Cross-sectional detail in the xz-plane, corresponding to the orange boxed region in **a**. The core and cladding regions of the fiber with refractive indices $n_{\mathrm{co},1}$ and $n_{\mathrm{cl},1}$, are indicated by dark and light blue shading, respectively, while the core and air-cladding of the nanowire with refractive indices $n_{\mathrm{co},2}$ and $n_{\mathrm{cl},2} = 1$, are indicated by green and gray shading. The solid black line represents the fundamental Gaussian guided mode of the fiber, $\boldsymbol{E}_{\mathrm{g},1}$, while the dashed black line represents the guided mode of the nanowire, $\boldsymbol{E}_{\mathrm{g},2}$, the energy of which is primarily concentrated in the exponentially decaying evanescent field. **c** As explained in the main text, the physical structure shown in **b** may be approximated by two infinite, uniform dielectric half-spaces for the purposes of determining the reflected fields at the interface. In the case that $n_1 \approx n_{\mathrm{co},1}$ is sufficiently small, the reflected field may be neglected and the transverse field at the interface may be taken to be $\boldsymbol{E}_{\mathrm{t}} \approx \boldsymbol{E}_{\mathrm{g},1}$.

efficiently in a single mode of the waveguide.

2.2.2 Fiber-to-Nanowire Coupling

We now turn our attention to the problem of coupling light from a single mode optical fiber to a thin, suspended, on-chip dielectric waveguide with a rectangular cross-section, hereafter referred to as a nanowire waveguide, as shown in Fig. 2.1a. In both waveguides the dimensions and refractive index profile are assumed to be such that they are single mode. That is, there is only one guided mode of each orthogonal transverse polarization. In practice, the photonic crystal structures to which we wish to couple only possess bandgaps for the TE-like polarizations where the electric field lies predominantly in the plane of the device layer (the $y = 0$ plane, as shown in the orange boxed region of Fig. 2.1a and Fig. 2.1b). Thus, we need only consider a single bound mode of each waveguide, with transverse field vectors $\boldsymbol{E}_{\mathrm{g},1}$ and $\boldsymbol{E}_{\mathrm{g},2}$ for the input (fiber) and output (nanowire) waveguides, respectively.

At the junction between the two waveguides, the total transverse field $\boldsymbol{E}_{\mathrm{t}}$ may be expanded in

terms of the local waveguide modes of waveguide 2, as in Eqs. 2.8 and 2.9:

$$\boldsymbol{E}_{\mathrm{t}} = a_{\mathrm{g}}\boldsymbol{E}_{\mathrm{g},2} + \boldsymbol{E}_{\mathrm{other},2}, \tag{2.11}$$

$$\boldsymbol{H}_{\mathrm{t}} = a_{\mathrm{g}}\boldsymbol{H}_{\mathrm{g},2} + \boldsymbol{H}_{\mathrm{other},2}, \tag{2.12}$$

where $\boldsymbol{E}_{\mathrm{other},2}$ denotes the sum over all other guided and radiation modes of waveguide 2 (similarly for $\boldsymbol{H}_{\mathrm{other},2}$). Physically it is clear that only the forward traveling guided mode can be excited in this configuration, so we do not need to distinguish between forward and backwards traveling guided modes at this stage. The total power propagating in the output guided mode is given by

$$P_{\mathrm{out},2} = \frac{1}{2}\mathrm{Re}\left[|a_{\mathrm{g}}|^2 \int dA \left(\boldsymbol{E}_{\mathrm{g},2} \times \boldsymbol{H}^*_{\mathrm{g},2}\right) \cdot \hat{\boldsymbol{z}}\right], \tag{2.13}$$

while the mode coefficient a_{g} can be determined using the orthogonality relation Eq. 2.10:

$$a_{\mathrm{g}} = \frac{\int dA \left(\boldsymbol{E}_{\mathrm{t}} \times \boldsymbol{H}^*_{\mathrm{g},2}\right) \cdot \hat{\boldsymbol{z}}}{\int dA \left(\boldsymbol{E}_{\mathrm{g},2} \times \boldsymbol{H}^*_{\mathrm{g},2}\right) \cdot \hat{\boldsymbol{z}}}, \tag{2.14}$$

We would like to relate this output power to the total power in the specified input mode, $\boldsymbol{E}_{\mathrm{g},1}$. However, the expressions above involve the *total* transverse field $\boldsymbol{E}_{\mathrm{t}}$, which is the sum of the known incident field and any fields reflected from the interface between the two waveguides. In general, the problem of determining the reflected fields is quite complicated and often has no exact analytical solution. However, in this specific case we can make a few reasonable approximations to obtain $\boldsymbol{E}_{\mathrm{t}}$ in terms of $\boldsymbol{E}_{\mathrm{g},1}$.

The refractive index profile of both waveguides, shown in the xz-plane in Fig. 2.1b, consists of a high-index core with refractive index n_{co} near the z-axis surrounded by a lower index cladding with refractive index $n_{\mathrm{cl}} < n_{\mathrm{co}}$. In both cases the waveguide modes of interest are sufficiently localized about the z-axis that the uniform cladding can be approximated as extending infinitely far in all transverse directions. In the case of a typical single mode fiber, $n_{\mathrm{co}} \approx n_{\mathrm{cl}}$ (e.g., in our work we use SMF-28 optical fiber at a vacuum wavelength of $\lambda \approx 1550$ nm, where $n_{\mathrm{co}}/n_{\mathrm{cl}} \approx 1.004$), and for the purposes of determining the reflected fields, the half-space containing the fiber may be approximated as a uniform dielectric with index $n_1 \approx n_{\mathrm{co}} \approx 1.5$. In the case of the nanowire, the cladding is vacuum ($n_{\mathrm{cl}} = 1$) while the core is typically has a relatively large refractive index (e.g., $n_{\mathrm{co}} \approx 2$ for Si_3N_4 and 3.5 for Si at $\lambda = 1550$ nm), so in general this half-space may not be approximated as having a single refractive index. In practice, however, the cross-sectional dimensions of the nanowire are small ($\sim$ 100s of nm) compared to the dimensions of the fiber core ($\sim 10\ \mu$m), and high coupling efficiency requires the spatial extents of the guided modes (i.e., the mode-field diameters) to be comparable. Consequently, the nanowire mode must be largely evanescent, such that almost all of the energy is carried in the cladding. In this case, we can approximate this region as a uniform dielectric with

$n \approx n_{\mathrm{cl}} = 1$, and so for the purposes of determining the reflected field the whole system may be roughly approximated as two infinite dielectric half-spaces, as shown in Fig. 2.1c.

Furthermore, when the guided mode diameter is large compared to the free-space wavelength of electromagnetic field (here, $\sim$ 10 μm and $\sim$ 1.5 μm, respectively), we may approximate the reflection by the usual expressions for reflection of a normal incidence plane wave at the interface between two uniform dielectrics. In this approximation the transverse field at the interface will be approximately [123]:

$$\boldsymbol{E}_{\mathrm{t}} \approx t\boldsymbol{E}^{\mathrm{f}}_{\mathrm{g},1} + r\boldsymbol{E}^{\mathrm{b}}_{\mathrm{g},1}, \tag{2.15}$$

where $\boldsymbol{E}^{\mathrm{f}}_{\mathrm{g},1}$ and $\boldsymbol{E}^{\mathrm{b}}_{\mathrm{g},1}$ denote the forward and backward propagating guided modes of waveguide 1, respectively. Clearly, only the forward propagating wave should contribute to the forward propagating power in waveguide 2. The prefactors t and r are the usual Fresnel reflection and transmission coefficients. The corresponding power transmission and reflection coefficients, T and R, are given by

$$T = |t|^2 = \frac{4n_1 n_2}{((n_1 + n_2)^2}, \tag{2.16}$$

$$R = |r|^2 = \frac{(n_1 - n_2)^2}{((n_1 + n_2)^2}, \tag{2.17}$$

where n_1 and n_2 are the approximate uniform refractive indices in the half-spaces corresponding to waveguide 1 and 2, respectively. Using $n \approx 1.5$ for the fiber and $n \approx 1$ for the on-chip waveguide, we find $T \approx 0.96$, and so neglecting for the moment the small amount of reflection loss we may approximate $\boldsymbol{E}_{\mathrm{t}} \approx \boldsymbol{E}_{\mathrm{g},1}$. Inserting this into Eq. 2.14, using Eq. 2.13 and normalizing by the power in the incident guided mode, we thus arrive at the approximate expression for the fiber-to-nanowire coupling efficiency $\eta_{\mathrm{overlap}} = P_{\mathrm{out},2}/P_{\mathrm{in},1}$ in terms of the guided modes

$$\eta_{\mathrm{overlap}} \approx \mathrm{Re}\left[\frac{\left|\int dA \left(\boldsymbol{E}_{\mathrm{g},1} \times \boldsymbol{H}^*_{\mathrm{g},2}\right) \cdot \hat{\boldsymbol{z}}\right|^2}{\int dA \left(\boldsymbol{E}_{\mathrm{g},2} \times \boldsymbol{H}^*_{\mathrm{g},2}\right) \cdot \hat{\boldsymbol{z}}}\right] \frac{1}{\mathrm{Re}\left[\int dA \left(\boldsymbol{E}_{\mathrm{g},1} \times \boldsymbol{H}^*_{\mathrm{g},1}\right) \cdot \hat{\boldsymbol{z}}\right]}. \tag{2.18}$$

The guided modes may be computed readily via a number of numerical methods, after which the vector fields can be numerically integrated according to Eq. 2.18. For any given fabricated structure, the device layer has a fixed thickness, and thus optimization of the coupling efficiency simply involves computed η_{overlap} as a function of nanowire width. As we will see below, the efficiency of this junction (neglecting the small amount of reflection) can approach $\eta_{\mathrm{overlap}} \approx 0.95$.

2.2.3 Adiabatic Waveguide Taper

Once the light has been coupled onto the chip, the width of the waveguide (and the confinement of the guided mode) must subsequently be increased to an appropriate width for coupling into the

optical cavity. As the width of the waveguide changes, so too do the corresponding local modes. Since the local fundamental guided mode at any given point along the waveguide is not in general orthogonal to the higher order guided modes and radiation modes at a different point, care must be taken in how the waveguide dimensions are altered so as to avoid coupling significant amounts of power out of the mode of interest.

This problem is formally similar to the problem of energy transfer between coupled modes in quantum mechanics, first considered by Landau and Zener [125] in the context of molecular energy level crossings. In the case of a time-dependent Hamiltonian $\hat{H}(t)$, which varies smoothly between $\hat{H}_0$ at time $t = t_0$ and $\hat{H}_1$ at $t = t_1$, the adiabatic theorem [126] states that a system initially prepared in the ground state of $\hat{H}_0$ at $t = t_0$ will be found in the ground state of $\hat{H}_1$ at $t = t_1$ provided that the change in the Hamiltonian is sufficiently slow. In particular, if the minimum energy between the ground state and any excited state of the system during the evolution of the Hamiltonian is Δ, "sufficiently slow" is taken to mean that the evolution time $\tau = t_1 - t_0 \gg \hbar\Delta^{-1}$.

This concept of adiabatic passage can be readily adapted to the problem of efficient eigenmode conversion in fiber-coupled optical waveguides [119,127], where variation of the waveguide dimensions along the propagation axis substitutes for the variation in the Hamiltonian in time, and difference in the modal effective refractive index substitutes for difference in the eigenstate energies. The relevant adiabatic condition can be readily derived from a coupled mode theory involving the local modes at different points along the waveguide.

We may write the total transverse electric field at any point z along the length of the waveguide as a sum over the complete set of local modes [123]:

$$\boldsymbol{E}_{\mathrm{t}} = \sum_{\mathrm{j}} \left(b_{\mathrm{j}}^{+}(z) + b_{\mathrm{j}}^{-}(z) \right) \boldsymbol{E}_{\mathrm{j}}(z), \tag{2.19}$$

$$\boldsymbol{H}_{\mathrm{t}} = \sum_{\mathrm{j}} \left(b_{\mathrm{j}}^{+}(z) - b_{\mathrm{j}}^{-}(z) \right) \boldsymbol{H}_{\mathrm{j}}(z), \tag{2.20}$$

$$b_{\mathrm{j}}^{\pm}(z) = a_{\mathrm{j}}^{\pm}(z) e^{\pm i \int_{z_0}^{z} dz' \beta_{\mathrm{j}}(z')}, \tag{2.21}$$

where both the modal amplitudes and the modal vector fields are functions of z, $\pm$ once again denotes forward or backward propagating waves, and the expansion coefficients $b_{\mathrm{j}}^{\pm}$ explicitly incorporate both the modal amplitudes $a_{\mathrm{j}}^{\pm}$ and the changing phase due to the z-dependence of the propagation constants. Note that, for convenience of notation, the sums in Eqs. 2.19 and 2.20 implicitly include both a sum over guided modes and an integral over radiation modes. Due to the lack of translation symmetry in z, the total transverse fields $\boldsymbol{E}_{\mathrm{t}}$ and $\boldsymbol{H}_{\mathrm{t}}$ now satisfy the more general, globally valid

transverse wave equations:

$$\frac{-1}{i\omega\mu_0}\nabla_{\mathrm{t}}\times(\nabla_{\mathrm{t}}\times\boldsymbol{E}_{\mathrm{t}})+\left(\hat{\boldsymbol{z}}\times\frac{\partial\boldsymbol{H}_{\mathrm{t}}}{\partial z}\right)=i\omega\epsilon_0 n^2\boldsymbol{E}_{\mathrm{t}}, \tag{2.22}$$

$$\frac{1}{i\omega\epsilon_0}\nabla_{\mathrm{t}}\times\left(\frac{1}{n^2}\nabla_{\mathrm{t}}\times\boldsymbol{H}_{\mathrm{t}}\right)+\left(\hat{\boldsymbol{z}}\times\frac{\partial\boldsymbol{E}_{\mathrm{t}}}{\partial z}\right)=-i\omega\mu_0\boldsymbol{H}_{\mathrm{t}}. \tag{2.23}$$

Inserting Eqs. 2.19 and 2.20 into Eqs. 2.22 and 2.23, and making use of the fact that the local modes at any given point $z=z_0$ ($\boldsymbol{E}_{\mathrm{j}}(z_0)$ and $\boldsymbol{H}_{\mathrm{j}}(z_0)$, and their corresponding propagation constants) are the complete set of solutions to Eqs. 2.5 and 2.6 for $n=n(x,y,z_0)$, we obtain

$$\sum_{\mathrm{j}}\left[\left(\frac{db_{\mathrm{j}}^{+}}{dz}-\frac{db_{\mathrm{j}}^{-}}{dz}-i\beta_{\mathrm{j}}\left(b_{\mathrm{j}}^{+}+b_{\mathrm{j}}^{-}\right)\right)(\hat{\boldsymbol{z}}\times\boldsymbol{H}_{\mathrm{j}})+\left(b_{\mathrm{j}}^{+}-b_{\mathrm{j}}^{-}\right)\left(\hat{\boldsymbol{z}}\times\frac{\partial\boldsymbol{H}_{\mathrm{j}}}{\partial z}\right)\right]=0, \tag{2.24}$$

$$\sum_{\mathrm{j}}\left[\left(\frac{db_{\mathrm{j}}^{+}}{dz}+\frac{db_{\mathrm{j}}^{-}}{dz}-i\beta_{\mathrm{j}}\left(b_{\mathrm{j}}^{+}-b_{\mathrm{j}}^{-}\right)\right)(\hat{\boldsymbol{z}}\times\boldsymbol{E}_{\mathrm{j}})+\left(b_{\mathrm{j}}^{+}+b_{\mathrm{j}}^{-}\right)\left(\hat{\boldsymbol{z}}\times\frac{\partial\boldsymbol{E}_{\mathrm{j}}}{\partial z}\right)\right]=0, \tag{2.25}$$

which is valid for any set of local modes defined at a given z_0.

We would like to now obtain the equation detailing the change in the mode amplitude of the fundamental forward propagating guided mode. If the waveguide variation is sufficiently slow (i.e., the local modes at $z=z_0$ are good approximations to the real solutions of Maxwell's equations in some small region about z_0), the fundamental guided mode can be well approximated at any given point along the waveguide by one of the local modes, which we will denote by the vector fields $\boldsymbol{E}_0$ and $\boldsymbol{H}_0$ with mode amplitude a_0^{+}. In the absence of absorption, we may choose the field components such that the transverse components are purely real. Then, taking the dot product of Eqs. 2.24 and 2.25 with $\boldsymbol{E}_0$ and $\boldsymbol{H}_0$, respectively, using the vector relation $\boldsymbol{A}\cdot(\boldsymbol{B}\times\boldsymbol{C})=\boldsymbol{B}\cdot(\boldsymbol{C}\times\boldsymbol{A})$, using the orthogonality relation Eq. 2.10, and subtracting the results, we arrive at the equations:

$$\frac{db_0^{+}}{dz}=i\beta_0 b_0^{+}+\sum_{\mathrm{j}}\left(C_{\mathrm{j}}^{+}b_{\mathrm{j}}^{+}+C_{\mathrm{j}}^{-}b_{\mathrm{j}}^{-}\right), \tag{2.26}$$

$$C_{\mathrm{j}}^{\pm}=-\frac{1}{2}\frac{\int dA\left(\boldsymbol{E}_0\times\frac{\partial\boldsymbol{H}_{\mathrm{j}}}{\partial z}\pm\frac{\partial\boldsymbol{E}_{\mathrm{j}}}{\partial z}\times\boldsymbol{H}_0\right)\cdot\hat{\boldsymbol{z}}}{\int dA\,(\boldsymbol{E}_0\times\boldsymbol{H}_0)\cdot\hat{\boldsymbol{z}}},\ j\neq 0;\quad C_0=0. \tag{2.27}$$

Using the orthogonality of the local modes, as well as the fact that $\boldsymbol{E}_0$, $\boldsymbol{H}_0$ and their derivatives go to zero exponentially as x and y go to infinity, it can be shown that [123, 128]

$$\left(\boldsymbol{E}_0\times\frac{\partial\boldsymbol{H}_{\mathrm{j}}}{\partial z}\pm\frac{\partial\boldsymbol{E}_{\mathrm{j}}}{\partial z}\times\boldsymbol{H}_0\right)\cdot\hat{\boldsymbol{z}}=\pm\sqrt{\frac{\epsilon_0}{\mu_0}}\frac{1}{n_{\mathrm{eff},0}-n_{\mathrm{eff,j}}}\boldsymbol{E}_0\cdot\boldsymbol{E}_{\mathrm{j}}\frac{\partial n^2}{\partial z}, \tag{2.28}$$

leading to

$$C_{\mathrm{j}}^{\pm}=\mp\sqrt{\frac{\epsilon_0}{\mu_0}}\frac{1}{4P_0\left(n_{\mathrm{eff},0}-n_{\mathrm{eff,j}}\right)}\int dA\boldsymbol{E}_0\cdot\boldsymbol{E}_{\mathrm{j}}\frac{\partial n^2}{\partial z}, \tag{2.29}$$

where P_0 is the power propagating in mode $\boldsymbol{E}_0$. For the tapered waveguides considered here, n^2 has the explicit expression $n^2 = 1 + \left(n_{\text{co}}^2 - 1\right) \text{Rect}(x/w(z), y/t)$, where $\text{Rect}(x, y)$ is the 2D rectangle function, $w(z)$ is the z-dependent width of the tapered waveguide, and t its (fixed) thickness. Taking the derivative with respect to z, we find

$$\frac{\partial n^2}{\partial z} = \left(n_{\text{co}}^2 - 1\right) \text{Rect}(y/t) \left(\delta(x/w - 1/2) - \delta(x/w + 1/2)\right) \frac{x}{w(z)^2} \frac{dw}{dz}. \tag{2.30}$$

Inserting this into Eq. 2.29 yields

$$C_{\text{j}}^{\pm} = \mp \sqrt{\frac{\epsilon_0}{\mu_0}} \frac{\left(n_{\text{co}}^2 - 1\right)}{8P_0 \left(n_{\text{eff},0} - n_{\text{eff,j}}\right)} \frac{1}{w} \frac{dw}{dz} \int_{-t/2}^{t/2} dy \left(\left(\boldsymbol{E}_0 \cdot \boldsymbol{E}_{\text{j}}\right)|_{x=w/2} + \left(\boldsymbol{E}_0 \cdot \boldsymbol{E}_{\text{j}}\right)|_{x=-w/2} \right). \tag{2.31}$$

From Eq. 2.31 it is clear that, regardless of any particulars of the specific local modes at a given point z_0 along the taper, the adiabatic transfer of power within the fundamental mode (i.e., $|C_{\text{j}}^{\pm}| \ll 1$) is possible provided the waveguide taper satisfies

$$\frac{dw}{dz} \ll |n_{\text{eff},0} - n_{\text{eff,j}}| \quad \forall j \neq 0, \tag{2.32}$$

at every point along the taper length. This adiabatic condition directly shows that efficient transmission requires that the rate of change of the waveguide dimensions be small relative to the gap in eigenvalue (here, the effective refractive index) between the mode of interest and any other modes of the system, and is directly analogous to the traditional adiabatic theorem of quantum mechanics.

In principle, Eqs. 2.26 and 2.31 can be used to directly compute the efficiency of power transfer in the fundamental guided mode over the length of the taper by computing a sufficient number of local modes at various points along the taper length and using them to make discrete approximations to $\partial b_0^+ / \partial z$. However, this method is both cumbersome and tricky. In particular, most of the time the waveguide will be single mode for most (if not all) of the length of the taper, and can only lose power via coupling to radiation modes. Thus, accurate computation of the transmission loss requires accurate determination of a large number of radiation modes. Due to the extended character of such modes, their simulation is highly dependent on the size and boundary conditions of the simulation region. A simpler, more efficient method for determining the transmission loss is to simulate propagation of the fundamental mode along the taper directly in the time-domain, using the finite-difference-time-domain method (FDTD), and utilize a mode coupling integral analogous to that used to compute η_{overlap} in Eq. 2.18 in order to determine the power in the local fundamental mode at the output of the taper. However, the adiabatic condition in Eq. 2.32 still proves a useful guide for choosing an appropriate taper shape and length.

2.3 Waveguide-to-Cavity Coupling

Once the nanowire waveguide has been tapered up to the appropriate size, we typically want to accomplish two tasks: couple resonant light into the optical cavity in a controllable way, and reflect off-resonant light back towards the fiber. The latter goal is not strictly necessary, but proves useful for calibrating the reflection efficiency and, in some cases, aligning the fiber. The most effective way of reflecting light over the wavelength band of interest is to simply terminate the waveguide in a broadband photonic crystal mirror, of the same type used to confine light in the optical cavity itself. To avoid excess scattering in the mirror region it is necessary to introduce the mirror gradually, by slowly increasing the hole size (with fixed period) over a number of periods n_{trans}, up the full hole dimensions required by the mirror. This is followed by a variable number of mirror periods n_{mirr}, which may be either as long as desired to give near-unity reflection, or varied to adjust the waveguide cavity coupling rate in the case of end-coupled cavities, as discussed below. In principle one can rigorously derive an adiabatic theorem for such a photonic crystal taper [129], similar to that derived in section 2.2.3, in order to estimate the adiabaticity of a given taper based on a set of local guided modes. In practice, however, when one is coupling into a bandgap region (rather than a guided mode of the crystal) the taper need not be long or especially complex, and so it is easier to simply run FDTD simulations to determine the scattering loss and reflectivity of a given tapered mirror and iteratively find an optimal design. Generally a linear tapering of the hole size over $n_{\mathrm{trans}} \approx 8 - 10$ periods is sufficient to nearly eliminate scattering loss. Regardless of the taper design, the scattering loss in this region is practically limited by how small of a hole one can fully etch (i.e., how gradually one can actually increase the hole size), which is dependent on the thickness and material of the device layer. Since this is finite, even the best designed tapered mirror typically has a real scattering loss of several percent.

2.3.1 End-coupling

To couple light into the cavity, two approaches exist. The most straightforward is to end-couple the cavity, as shown in Fig. 2.2a. By simply reducing n_{mirr} to make a partially transparent mirror which serves as one of the cavity's end-mirrors, a controllable amount of the incident light is permitted to leak through to the cavity region. As the defect region of the cavity is highly reflective for light detuned from the cavity resonance frequencies, this does not compromise the overall off-resonant reflectivity. The measured quality factors of a series of identical cavities are shown in Fig. 2.2b, with extrinsic (Q_{e}), intrinsic (Q_{i}), and total (Q_{t}) quality factor shown as red, blue, and green circles, respectively. Though this coupling method can be generally applied to any type of photonic crystal cavity, the specific devices measured here are the so-called "zipper" double nanobeam cavities [130, 131] shown below in section 2.4, where the partial mirror exists on only one of the two beams that

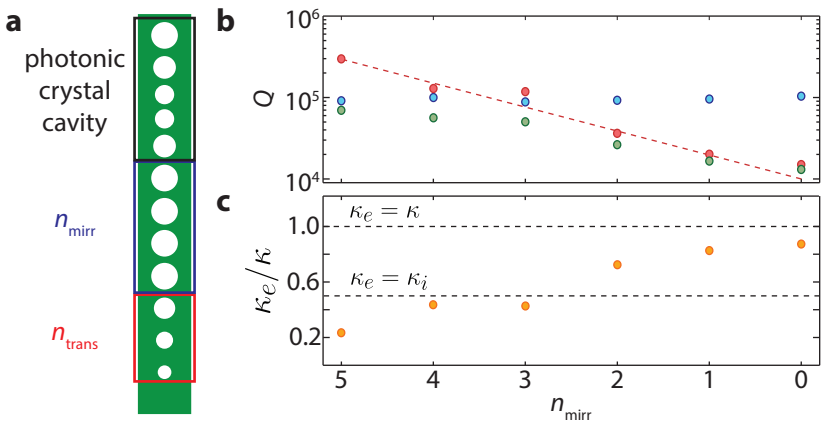

Figure 2.2: **End-coupled cavity.** **a**, Schematic of the end-coupled cavity. The sizes of the photonic crystal holes gradually taper up over a number of periods n_{trans} to minimize scattering loss, after which the mirror section extends for a number of periods n_{mirr}, which controls the coupling rate between the waveguide and the photonic crystal cavity. **b**, Measured quality factor Q, for a series of identical end-coupled cavities with varying n_{mirr}. Red circles show the extrinsic quality factor Q_{e}, blue circles show the intrinsic quality factor Q_{i}, and green circles illustrate the total quality factor. The red dashed line indicates the theoretical extrinsic quality factor determined from FEM simulations. **c**, Waveguide-cavity coupling ratio $\kappa_{\mathrm{e}}/\kappa$ corresponding to the points shown in **b**, showing the transition from under- to overcoupling. Dashed horizontal lines correspond to critical coupling ($\kappa_{\mathrm{e}} = \kappa_{\mathrm{i}} = \kappa/2$) and perfect overcoupling ($\kappa_{\mathrm{e}} = \kappa$).

comprise the cavity. The extrinsic coupling is seen to match well with a theoretical estimate based comparing the power outflow in the waveguide to the total power outflow in a FEM simulation of the waveguide coupled cavity. Due to the exponential attenuation of light in a photonic crystal bandgap the coupling rate between the waveguide and cavity is exponentially dependent on n_{mirr}. For the modest Q_{i} values measured in these devices, variation of the coupling rate over a wide range is not observed to affect the intrinsic scattering rate of the cavity, though in general there can be a small modification which is apparent once the intrinsic scattering rate becomes low enough. For optomechanical measurements the relevant figure of merit is the cavity-to-waveguide coupling efficiency $\kappa_{\mathrm{e}}/\kappa = Q_{\mathrm{t}}/Q_{\mathrm{e}}$, which is displayed in Fig. 2.2c. As this coupling geometry is single-sided, we can reach the overcoupled regime $\kappa_{\mathrm{e}} > \kappa_{\mathrm{i}}$, and can in principle approach the perfectly overcoupled regime $\kappa_{\mathrm{e}} = \kappa$.

2.3.2 Side-coupling

While the end-coupled scheme is conceptually simple, it has a number of drawbacks. The first is that we are limited in how finely we can control the extrinsic coupling rate, as we can only add or remove discrete numbers of holes which cause relatively large jumps in the waveguide-cavity coupling due to the exponential dependence on n_{mirr}. More importantly, the end-coupled scheme removes the ability to clamp one end of the cavity in a specific way. In particular, for the nanobeam OMCs studied

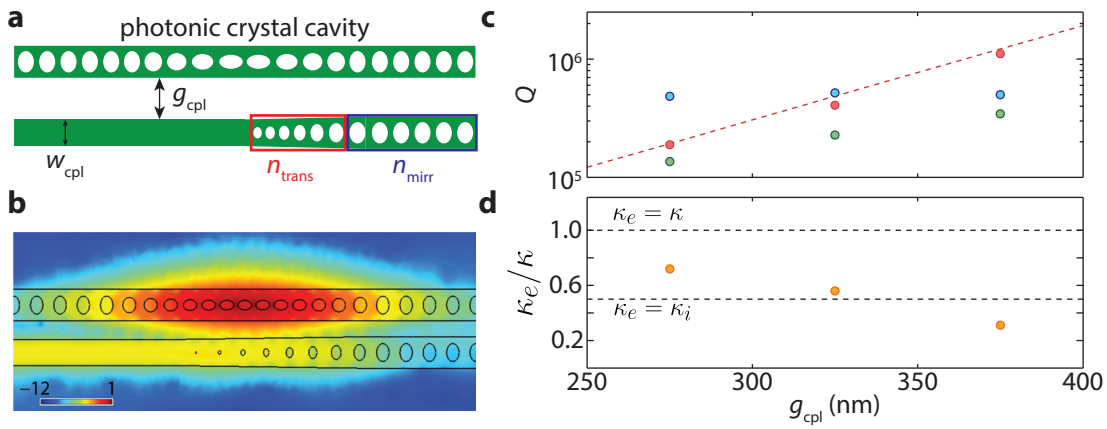

Figure 2.3: **Side-coupled cavity.** **a**, Schematic of the side-coupled cavity. The input waveguide transitions into a photonic crystal mirror similarly to the end-coupled case, but the width of the waveguide prior to the mirror, w_{cpl}, can be smaller than the width of the mirror section in order to maximize the evanescent field overlap with the cavity. The overall coupling rate is then primarily determined by the gap between the waveguide and the cavity, g_{cpl}. **b**, Typical FEM simulation of the optical field intensity for a side-coupled nanobeam cavity, normalized to its maximum and displayed on a logarithmic scale. **c**, Measured quality factors Q, for a series of identical side-coupled cavities with varying g_{cpl}, for fixed w_{cpl}. Red circles show the extrinsic quality factor Q_{e}, blue circles show the intrinsic quality factor Q_{i}, and green circles illustrate the total quality factor. The red dashed line indicates the theoretical extrinsic quality factor determined from FEM simulations. **d**, Waveguide-cavity coupling ratio $\kappa_{\mathrm{e}}/\kappa$ corresponding to the points shown in **b**, showing the transition from under- to overcoupling. Dashed horizontal lines correspond to critical coupling ($\kappa_{\mathrm{e}} = \kappa_{\mathrm{i}} = \kappa/2$) and perfect overcoupling ($\kappa_{\mathrm{e}} = \kappa$).

in this work we typically want to place acoustic shielding on either end of the cavity, as described in section 1.3.5. A solution to both of these problems is to use the side-coupling scheme illustrated in Fig. 2.3a. In this scheme, the waveguide is terminated in a tapered photonic crystal mirror, but now n_{mirr} can be made arbitrarily large. Placing the waveguide in the near-field of the cavity allows for coupling via the evanescent fields of the cavity and waveguide, where the coupling rate may be continuously varied by changing the waveguide-cavity gap g_{cpl}. In addition, the waveguide width just prior to the photonic crystal mirror, w_{cpl}, is typically fabricated smaller than the width of the mirror itself (with the width tapering along with the hole size in the transition region), in order to maximize the evanescent field overlap. In general, there is an optimal value of w_{cpl} where the waveguide mode is not too confined to overlap with the cavity mode profile, but not so extended that such overlap is minimal. For the Si nanobeam design used in this work, the optimal value is $w_{\mathrm{cpl}} \approx 500$ nm. A FEM simulation of the coupled waveguide-cavity system is displayed in Fig. 2.3b. As with the end-coupled devices, the measured quality factors and cavity-to-waveguide coupling efficiencies are shown as a function of g_{cpl} in Fig. 2.3c and d, respectively. For these measurements the cavity is the Si nanobeam design discussed in section 1.3.4 and used for the experiments in chapters 3 and 4. Similarly to the end-coupled mirror, the extrinsic coupling rate has an exponential

dependence on g_{cpl}, which matches well our theoretical estimates based on FEM simulations.

2.4 Butt-Coupling

The simplest approach to building a fiber-to-waveguide coupler is to cleave a standard single mode fiber and place the endface in direct contact (or, more practically, very nearly in contact) with the endface of the on-chip nanowire waveguide, exactly as shown in Fig. 2.1a. This approach is often called "butt-coupling". For the initial demonstration of these couplers, we chose to work with devices fabricated from thin-film, stoichiometric Si_3N_4 deposited on top of Si. Though the lower refractive index and smaller photoelastic effect of Si_3N_4 compared to Si results in reduced optomechanical coupling, there are a number of advantages which led to the selection of this material for the initial coupler work. At the time a primary reason to select Si_3N_4 over Si was due to its much larger bandgap as compared to Si. Two-photon absorption limited earlier work in Si nanobeam OMCs at higher input optical powers by causing degradation of the optical and mechanical quality factors [45], and moving to a wider bandgap material mitigates this effect. As we will see in subsequent chapters, the dominant source of absorption-induced heating at sub-kelvin temperatures in Si is found to be linear absorption due to mid-gap defect states rather than two-photon absorption. Nevertheless, the material properties of stoichiometric Si_3N_4 do provide a number of advantages when fabricating butt-couplers. In particular, the large tensile stress (~ 1 GPa) of the device layer causes the suspended waveguide to pull taut when it is undercut, allowing fabrication of extremely long suspended waveguides with minimal bending out of plane. This allows for nearly perfect adiabatic waveguide tapers, as the change in waveguide width can be made extremely slow as a function of length along the waveguide. Additionally, the Si substrate etches anisotropically in KOH, while Si_3N_4 remains unetched, allowing for easy fabrication of v-grooves that enable robust self-alignment of the fiber.

The greatly reduced optomechanical coupling rate in Si_3N_4 compared to Si makes it difficult to measure the high-frequency mechanical modes of the nanobeam OMCs. Consequently, the initial experiments involving these couplers made use of an alternative optomechanical cavity design comprised of two parallel nanobeams, the so-called "zipper" cavity [130, 131], which strongly couples the optical cavity to the low-frequency, fundamental in-plane mechanical modes of the beams. These structures are in the sideband unresolved regime, where $\omega_{\mathrm{m}} \ll \kappa$, and as such are better suited for applications relying on fast measurement of the mechanical oscillator's position, such as precision force sensing [132, 133], and feedback cooling of mechanical motion [134, 135], rather than the quantum state-engineering applications which are the ultimate goal of this work. As a result, our presentation of these devices will focus only on the design and optical characterization of the coupler efficiency, while details about the associated optomechanical measurements may be found

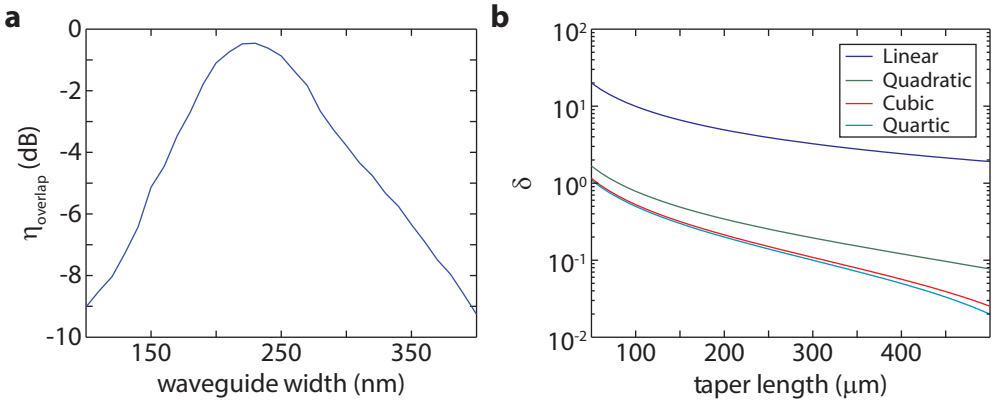

Figure 2.4: **Butt-coupler design.** **a**, Fiber-to-nanowire waveguide coupling efficiency η_{overlap} as a function of the nanowire width. **b**, Adiabaticity parameter δ as a function of taper length for different polynomial taper shapes, assuming initial and final waveguide widths of 230 nm and 850 nm, respectively. For both plots the free-space optical wavelength is $\lambda = 1550$ nm and the thickness of the nanowire waveguide is fixed at 400 nm.

in Refs. [136] and [55].

2.4.1 Design

The zipper cavities used in this work are designed for operation at an optical frequency of $\omega_c/2\pi \approx 193$ THz, corresponding to a free-space wavelength of $\lambda \approx 1550$ nm. As designed, the cavities have a thickness of 400 nm and each nanobeam possesses a width of 850 nm. This fixes both the thickness and the final width of the nanowire waveguide, and optimization of the coupler involves determining the optimal initial width at the fiber-to-nanowire junction as well as the optimal taper length and shape. The optical fiber used is standard SMF-28, which has a single guided mode with a mode-field diameter of 10.4 μm at $\lambda = 1550$ nm [137]. The guided field of the fiber can be obtained numerically via FEM simulation using the known core diameter and core-cladding refractive index difference, or more simply by using the well-known approximate analytical solutions for weakly guiding fibers [124], adjusting the parameters to obtain the appropriate mode-field diameter. Once the guided mode of the fiber is specified, we simulate the fundamental TE-like guided mode profile of the nanowire waveguide as a function of waveguide width using the commercial FEM package COMSOL [82]. We can then compute the coupling efficiency η_{overlap} by numerically integrating the two mode profiles according to Eq. 2.18, producing the plot shown in Fig. 2.4a, where the optimal width of ~ 230 nm is seen to yield a coupling efficiency of $\eta_{\text{overlap}} \approx 95\%$ (neglecting the small amount of reflection from the interface, as discussed in section 2.2.2).

Once the initial waveguide width is fixed, we would like to determine the optimal tapering function between the initial and final waveguide widths. As discussed in section 2.2.3, the most efficient method of computing the transmission efficiency of the tapered waveguide is to use FDTD

to propagate the initial guided mode along the waveguide. Accurate computations tend to be rather time-consuming due to the necessity of a large simulation volume combined with a sub-wavelength grid spacing, and in order to reduce the overall number of full propagation simulations we must perform it is useful to use the general adiabatic condition given in Eq. 2.32 to roughly define the taper shapes and lengths for which we expect high efficiency. We use COMSOL to compute the local guided modes and corresponding effective indices at a variety of waveguide widths w in between the two fixed endpoints. Once we select a specific taper function $w(z)$ describing the width as a function of position along the waveguide, we may use these results to fit an approximate functional form to $n_{\text{eff}}(w(z))$ for each guided mode. This allows us to compute both dw/dz and n_{eff} for any arbitrary point z along the length of the waveguide. We define the adiabatic parameter

$$\delta = \max \left[\left(|n_{\text{eff,0}} - n_{\text{eff,j}}| \right)^{-1} \frac{dw}{dz} \right]_{\text{z,j}}, \tag{2.33}$$

where as in section 2.2.3 the subscripts 0 and j refer to the fundamental local guided mode and any other local transverse modes, respectively, and where as indicated the maximization is performed over all points z along the waveguide and over all local modes other than the fundamental at each point z. According to Eq. 2.32 a necessary condition for an efficient taper is that $\delta \ll 1$. After the initial computation of several sets of local guided modes, we may rapidly compute δ for any number of trial taper functions $w(z)$ to perform an initial assessment of their adiabaticity.

Over the majority of the taper, the width is small enough that the waveguide is single-mode. Therefore, the fundamental mode can only lose energy via coupling to radiation modes. These modes all have effective indices $n_{\text{eff}} \leq 1$, so the index difference in Eq. 2.33 may be approximated as $|n_{\text{eff,0}} - n_{\text{eff,j}}| \approx n_{\text{eff,0}} - 1$. As discussed in section 2.2.2, near the fiber-nanowire coupling junction, the fundamental guided mode of the nanowire must be almost entirely evanescent in order to match the fiber mode's large mode-field diameter, and thus at this point the effective index of the mode is $n_{\text{eff,0}} \approx 1$. Therefore, dw/dz must be very small close to the nanowire tip in order to avoid scattering into radiation modes. As the width increases and the mode becomes more confined, $n_{\text{eff,0}}$ becomes larger and the width may increase at a faster rate. This suggests that using higher order polynomials (i.e., $w(z) \propto z^a$ for large a) will result in higher transmission efficiency. Fig. 2.4b shows the adiabatic parameter δ computed for a variety of polynomial shapes, from linear ($a = 1$) to quartic ($a = 4$), with lengths ranging up to several hundred microns. As we can see, the necessary condition $\delta \ll 1$ is only obtained for higher degree polynomials and even then only for lengths over ~ 300 μm. We can also see that going to larger and larger degree polynomials offers diminishing returns, as the difference between cubic and quartic taper shapes is minimal. Guided by this approximate analysis, we select a cubic, 400 μm long taper for our design. Using the commercial FDTD solver Lumerical [138] to compute the full transmission efficiency η_{taper}, we find that this design yields $\eta_{\text{taper}} \approx 98\%$.

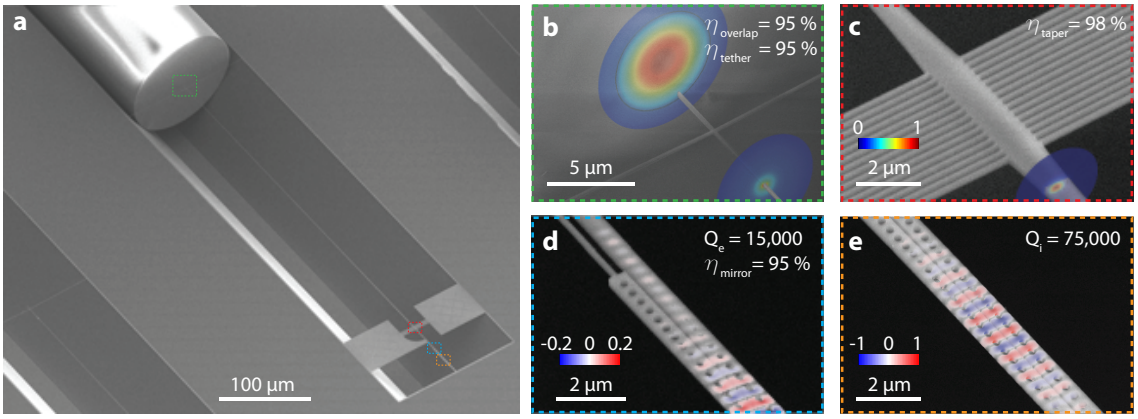

Figure 2.5: **Fabricated Si_3N_4 butt-coupler.** SEM images illustrating the optical coupling scheme, with FEM-simulated mode profiles of optical power in **b** and **c**, and electric field in **d** and **e**. **a**, Fabricated device after fiber coupling via self-aligned v-groove placement. **b**, The fiber-to-nanowire junction. **c**, The waveguide with supporting tethers after tapering. **d**, Photonic crystal mirror. **e**, Photonic crystal cavity.

A series of SEM images of a fabricated butt-coupler device are shown in Fig. 2.5. Releasing the structure in KOH selectively etches the Si substrate down to the $\langle 111 \rangle$ crystal planes, forming a v-groove in which the cleaved fiber rests, as shown in Fig. 2.5a. As the fiber diameter and v-groove angles are known, this allows us to precisely define the vertical displacement of the fiber relative to the device layer simply by changing the width of the groove during the initial lithography. In this way, the v-groove allows for robust vertical and lateral alignment with the nanowire waveguide. Once the fiber is firmly set in the groove it may be permanently affixed with epoxy, at which point no further alignment is necessary. We have verified that no appreciable change in the coupling efficiency is observed when such a device is cooled down to sub-kelvin temperatures inside a dilution refrigerator, indicating that misalignment due to thermal contraction is not a significant issue. Although the calculation of the theoretical fiber-to-nanowire coupling efficiency assumes they are in direct contact, in practice it is only necessary to bring the two close enough so that the field launched from the cleaved fiber facet does not expand appreciably before reaching the nanowire. As the mode of the fiber is approximately Gaussian, this distance can be characterized by the Rayleigh range, the distance at which the beam radius increases by a factor of $\sqrt{2}$, which in this case should be roughly 50 μm. Consequently, it is sufficient to only bring the fiber within $5 - 10$ μm of the nanowire tip, to avoid the possibility of damaging the nanowire waveguide upon contact.

As shown in Fig. 2.5b, a single tether is necessary to support the nanowire tip. Due to the tensile stress of the device layer, in principle no other tethers are needed to support the structure, though in this particular device additional tethers are placed near the optomechanical cavity, as shown in Fig. 2.5c, in order to isolate the zipper cavity from the low-frequency mechanical vibrations

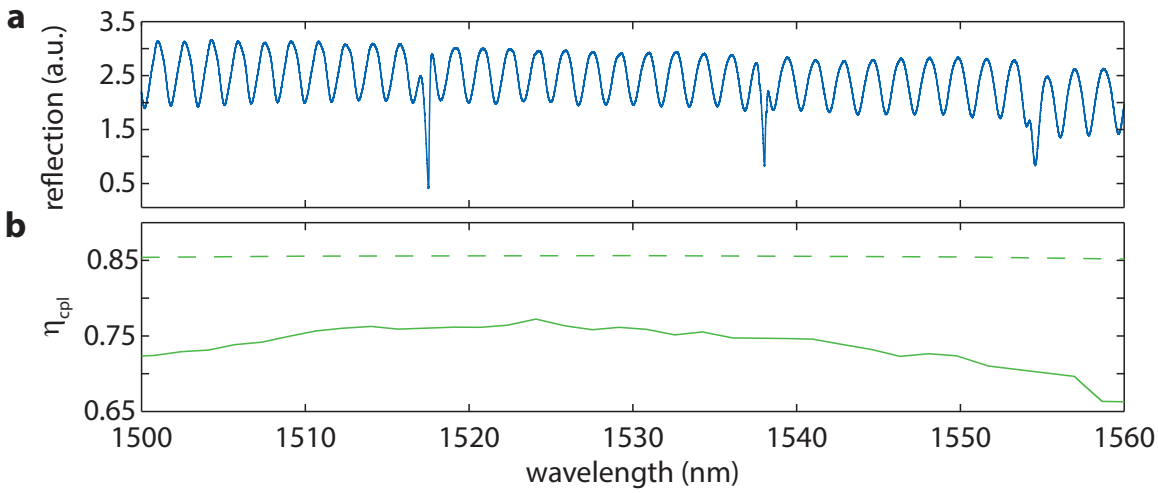

Figure 2.6: **Si_3N_4 butt-coupler optical response.** **a**, Wide-band optical reflection spectrum of a butt-coupled device. **b**, The solid line shows the estimated single-pass transmission efficiency η_{cpl} extracted from the fringe visibility in **a**. The as-designed ideal value of η_{cpl} is shown as a dashed line.

of the taper itself. The additional scattering loss of the tethers may be computed directly via FDTD propagation simulations. For a 70 nm wide tether (the smallest tether size which may be reliably fabricated) near the nanowire tip the transmission is computed to be $\eta_{\mathrm{tether}} \approx 95\%$. The array of tethers further back along the waveguide has a negligible contribution to the overall transmission efficiency due to the high confinement of the guided mode at this point. After tapering up to the full cavity nanobeam width of 850 nm the waveguide is terminated in a photonic crystal mirror, shown in Fig. 2.5d, and end-coupled to the optomechanical cavity shown in Fig. 2.5e, as described in section 2.3. Due to the transition between nanowire waveguide and photonic crystal mirror, the reflection efficiency of this section is $\eta_{\mathrm{mirror}} \approx 95\%$. Overall, the as-designed round-trip reflection efficiency of the coupler is $\eta_{\mathrm{rt}} = \left(\eta_{\mathrm{overlap}}\eta_{\mathrm{tether}}\eta_{\mathrm{taper}}\right)^2 \eta_{\mathrm{mirror}} = 74\%$. For the purposes of optomechanical transduction, we typically are only concerned with the single-pass transmission efficiency (that is, the transmission efficiency from the cavity output into the fiber), as we are rarely limited by our input power and only the detection efficiency matters. For the presented coupler design, this figure is ideally $\eta_{\mathrm{cpl}} = \sqrt{\eta_{\mathrm{rt}}} = 86\%$.

2.4.2 Characterization

The broadband optical reflection spectrum of a butt-coupled device in the 1500 nm wavelength band is displayed in Fig. 2.6a. The three sharp dips around $\lambda = 1515$, 1535, and 1555 nm correspond to resonant modes of the optical cavity, while the broad off-resonant fringes arise from a low-finesse Fabry-Perot etalon formed between the near-perfect photonic crystal mirror and the weakly reflective cleaved fiber facet. The exact reflectivity of the fiber can be calculated via FDTD and is found to

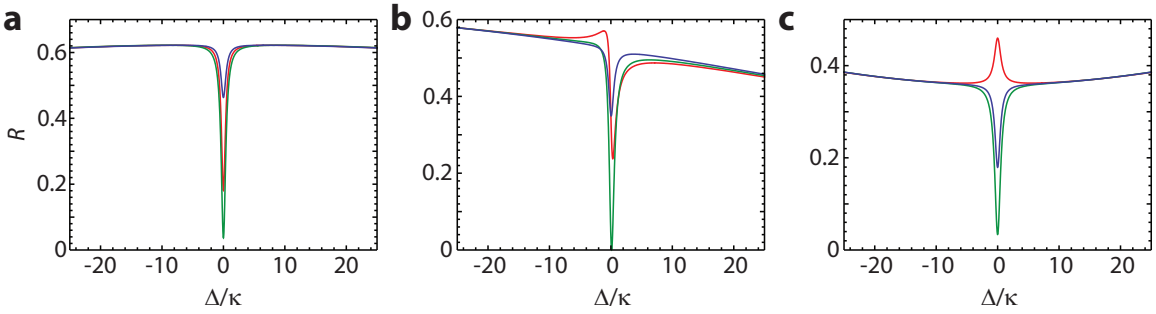

Figure 2.7: **Under- and overcoupled butt-coupled cavities.** Simulated normalized reflection R for undercoupled (blue, $\kappa_e/\kappa = 0.1$), critically coupled (green, $\kappa_e/\kappa = 0.5$) and overcoupled (red, $\kappa_e/\kappa = 0.9$) cavities, showing the different lineshapes which result when the cavity resonance frequency is **a**, at the top of a fringe, **b**, on the side of a fringe, and **c**, at the bottom of a fringe. For all of these plots we assume $R_{fib} = 3.5\%$, $\eta_{cpl} \approx 70\%$, and κ equal to roughly 1% of the etalon's free-spectral range.

be $R_{fib} = 3.5\%$, in good agreement with the rough prediction based on Fresnel coefficients made in section 2.2.2. While the presence of this additional etalon can complicate certain applications, it does provide a convenient way of calibrating the coupler efficiency without calibrating the losses in the detection train. In particular, given that the reflectivity of the fiber endface is known, and approximating that the photonic crystal mirror has unity reflectivity, a simple Fabry-Perot model including internal loss shows that the visibility of the fringes is related to the single-pass efficiency by

$$V = \frac{R_{max} - R_{min}}{R_{max} + R_{min}} = \frac{2\sqrt{R_{fib}}\,(1 - R_{fib})\,\eta_{cpl}\left(1 - \eta_{cpl}^2\right)}{R_{fib} + \eta_{cpl}^2 + R_{fib}\eta_{cpl}^2\left(R_{fib} + \eta_{cpl}^2 - 4\right)}, \tag{2.34}$$

where R_{max} and R_{min} are the maximum and minimum reflection of a fringe, respectively. The single-pass efficiency, determined using Eq. 2.34, is shown as a solid line in Fig. 2.2.2b, with the theoretical maximum single-pass efficiency shown as a dashed line. This device is observed to have a measured single-pass efficiency of $\eta_{cpl} = 75\%$. The discrepancy with the ideal value of 86% is attributed to small differences in the sizes of the waveguide dimensions (e.g., the nanowire tip width, the tether width, etc.) due to fabrication error.

The presence of the fringes also simplifies the determination of κ_e. Typically in cavity input-output theory the off-resonance reflection is assumed flat as a function of frequency. In that case, the cavity appears in the reflection scan as a Lorentzian dip with a linewidth equal to the total cavity loss rate κ and a minimum reflection level (normalized to the off-resonance reflection level) given by $R_0 = (1 - 2\kappa_e/\kappa)^2$. In the case of double-sided coupling, where $\kappa_e < \kappa/2$, this minimum reflection level can be used to directly determine the extrinsic coupling rate κ_e. In single-sided coupling, however, this is no longer sufficient, as there exist two possible values of κ_e/κ which yield the same value of R_0. In particular, the transformation $\kappa_e/\kappa \to 1 - \kappa_e/\kappa$ leaves R_0 unchanged. In general, one

will need to look at both the amplitude *and* phase response of the optical cavity in order to determine the cavity-to-waveguide coupling efficiency, as described below in section 2.5.2. In the presence of an intermediate low-finesse etalon, however, the optical cavity can have highly non-Lorentzian shapes depending on κ_e/κ and where the cavity resonance sits on a fringe. A simple coupled-cavity model, illustrated in Fig. 2.7, shows that an overcoupled cavity can possess a strongly asymmetric, Fano-like lineshape when positioned on the side of a fringe, or even become a peak rather than a dip when positioned at the bottom of a fringe. These lineshapes arise from interference between the circulating etalon field and the cavity reflection due to the strong, frequency-dependent phase shift upon reflection from an overcoupled cavity. In most cases, this model allows us to determine κ_e/κ directly from the amplitude response of the cavity. Further details of the model used to derive Eq. 2.34 and produce Fig. 2.7 are given in Appendix B.

2.5 End-fire Coupling

The butt-couplers presented in section 2.4 possess many desirable characteristics: high efficiency, controllable single-sided coupling, and alignment-free coupling (after the initial placement and gluing of the fiber). However, there are a number of drawbacks that make them unsuitable for our purposes. In particular, while the permanent placement of a fiber in a v-groove is convenient, it also limits us to measuring a small number of devices per fiber. This is undesirable for measurements inside a dilution fridge, as we are often limited in the number of fibers which can be sent into the fridge. Most importantly, utilizing a cleaved fiber requires a large evanescent field at the fiber-to-nanowire junction which in turn requires a waveguide taper with a length on the order of several hundred microns, as explained in section 2.4.1. While this is possible with high stress Si_3N_4, a free-standing Si waveguide will begin to sag if it is made much longer than $15 - 20$ μm, precluding the use of such long waveguide tapers without an excessive number of supporting tethers (which in turn greatly increase the scattering loss). As we would like to use Si devices to take advantage of the stronger optomechanical coupling it is necessary to substantially reduce the taper length, which in turn necessitates a larger initial nanowire width and correspondingly smaller mode-field diameter.

To address this issue, the tapered coupler may be modified to couple to a lensed optical fiber rather than a cleaved optical fiber. Commercial lensed optical fibers can possess mode-field diameters at the beam waist of only a couple microns, allowing for a much larger optimal nanowire tip width. Couplers of this type are often referred to as "end-fire" couplers, though the design principles are identical to that used for "butt-couplers". While it is possible to mount such lensed fibers in v-grooves similar to those used for the butt-couplers, we opt for a design where the lensed fiber is mounted on a set of nanopositioning stages for alignment with the waveguide tapers, to avoid the drawback associated with permanently coupled fibers. As we will explain below, the reflection from

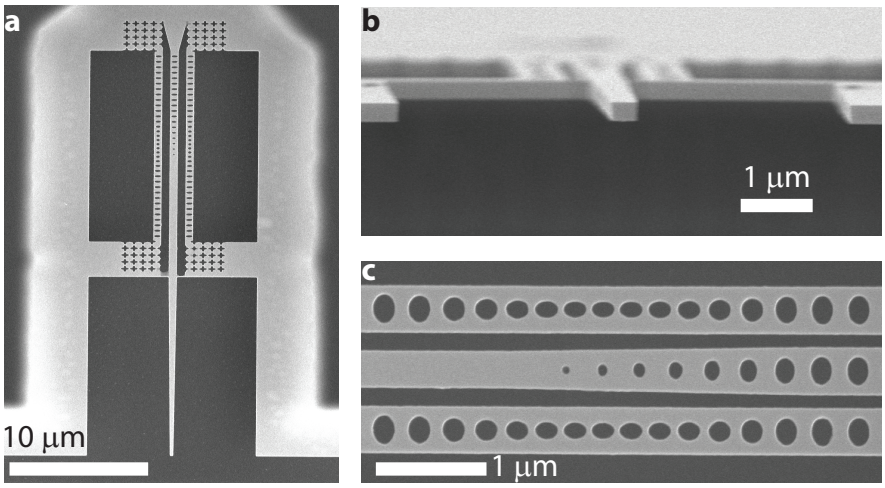

Figure 2.8: **Fabricated Si end-fire coupler.** **a**, SEM image of an end-fire coupled nanobeam OMC. **b**, Edge view SEM image of the central waveguide tip. **c**, Zoom-in SEM image showing details of the waveguide-cavity coupling region.

the chip proves sufficient to allow precise alignment without imaging the sample, and thus allows for simple, high-efficiency coupling even inside a dilution fridge.

2.5.1 Design

The design process for the end-fire coupler is identical to that described for the butt-couplers in section 2.4.1. The lensed fibers used in this work produce a focused Gaussian beam with a mode-field diameter at the beam waist of ~ 2.5 μm, resulting in a substantially larger optimal beam width (in this case, the coupling is optimized assuming the nanowire tip is placed in the focal plane of the lensed fiber) but a comparable fiber-to-nanowire coupling efficiency of $\eta_{\text{overlap}} \approx 97\%$. Due to the limitations on maximum taper length in Si, the optimal tapering strategy is found to be a linear taper with a length of ~ 13 μm; a longer taper would require additional support tethers to avoid sagging of the waveguide, which in turn would add increased scattering loss and negate any efficiency gained by utilizing a longer taper. The estimated transmission efficiency of such a short taper is only $\eta_{\text{taper}} \approx 80\%$, in contrast to the near-unity efficiency achieved with the much longer taper in the case of Si_3N_4. Including scattering loss from the single support tether, the ideal single-pass loss of the coupler is estimated to be $\eta_{\text{cpl}} = 70\%$.

A sample fabricated coupler is shown in both a top-down and edge-on view in Fig. 2.8a and b. To couple to the waveguides, we utilize the side-coupling geometry detailed in section 2.3.2 and shown in Fig. 2.8c, which allows us to clamp the nanobeam OMCs on either side with a "cross"-type acoustic radiation shield, as seen in Fig. 2.8a. To maximize the number of devices which can be measured with a single coupler we couple each waveguide to two cavities (one on either side) that

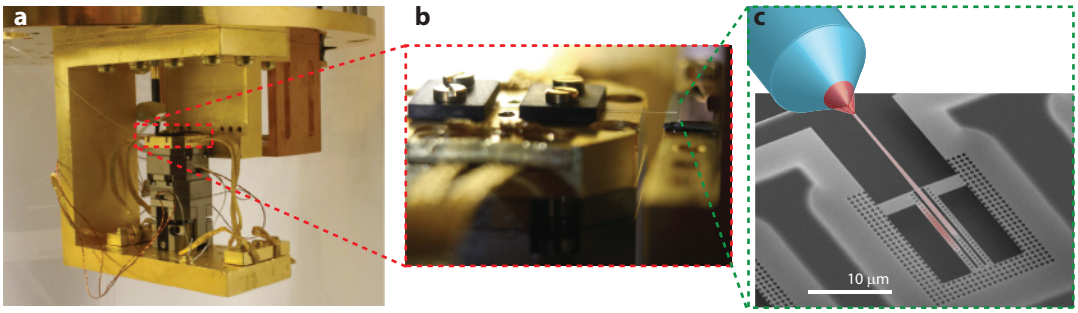

Figure 2.9: **Fiber coupling in the dilution refrigerator. a**, Photo of the attocube nanopositioner stages and sample mount affixed to the bottom of the dilution refrigerator mixing plate. **b**, Photo of the lensed fiber mounted on the positioning stages, showing alignment to a Si test chip. **c**, Diagram showing the lensed fiber (not-to-scale) coupling to an end-fire coupled Si OMC, with a superimposed FEM simulation of optical intensity.

are fabricated to have slightly different resonance frequencies. As discussed in section 1.3.7, during fabrication a deep trench is etched in the Si substrate after etching the device layer, creating an exposed mesa with the couplers sitting at the very edge. This allows us to bring the lensed fiber close enough to the couplers that the nanowire tip lies in the fiber's focal plane (about ~ 14 μm from the fiber tip).

The microchip sample is mounted to bottom of the mixing chamber plate in the dilution refrigerator, and the lensed fiber tip is clamped down on a stack of position encoded piezo xyz-stages (attocube models ANPx101 and ANPz101) (Fig. 2.9a), which allow nanopositioning with respect to the sample. When mounting the sample, the fiber is roughly aligned to within a few millimeters, as illustrated in Fig. 2.9b, prior to closing up the fridge and cooling down. Once the sample is cold, the fiber tip can be carefully positioned to maximize the reflected signal. This may seem to require imaging of the sample, just like fiber-taper coupling. However, in contrast to the case of an evanescently coupled fiber-taper, a measurable reflection from the edge of the chip is observed even when the lensed fiber is far from alignment with an actual device, allowing us to bring the fiber into alignment along one axis at a time. First, the fiber is slowly lowered while monitoring the reflected power. Due to both the high refractive index of the device layer and the fact that light launched from the fiber hits the device layer at roughly normal incidence (as opposed to the underlying silicon substrate which is etched at an angle), the reflected power is maximized when the fiber tip is vertically aligned with the device layer, even if it is not laterally aligned with a coupler. Once rough vertical alignment is achieved, the fiber is scanned laterally until a weak increase in the signal is observed due to the presence of a waveguide coupler. At this point the fiber may be slowly moved towards the chip, readjusting vertical and lateral alignment at each step to maximize the reflected signal, until the fiber is well aligned with the coupler. Once the reflected signal is

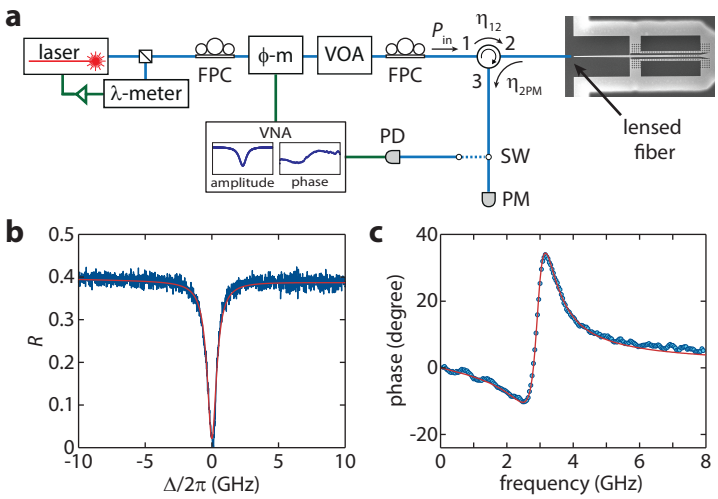

Figure 2.10: **End-Fire coupler optical characterization.** **a**, The experimental setup used to optically characterize the end-fire couplers. λ-meter: wavemeter, FPC: fiber polarization controller, ϕ-m: electro-optic phase modulator, VOA: variable optical attenuator, SW: optical switch, PM: optical power meter, PD: high-speed photodetector, VNA: vector network analyzer. **b**, Calibrated optical reflection spectrum of a cavity resonance versus cavity detuning $\Delta = \omega_c - \omega_l$ (blue) with a Lorentzian fit (red), yielding $Q = 236,000$. **c**, Phase response of the optical resonance shown in **b**, yielding $\kappa_e/\kappa = 0.52$.

globally maximized, it is a straightforward task to scan laterally across the sample to test different devices. When the desired device is found the stages may be grounded, at which point we observe no measurable drift in the reflection efficiency (and thus the fiber position), even over many days of testing.

2.5.2 Characterization

The lensed fibers used for end-fire devices are antireflection coated, thus the only substantial reflections along the taper are due to weak (0.5% or less) reflections from the waveguide-air interface. While this leads to some weak fringing effects similar to those observed in the butt-couplers (see section 2.4.2), the specific reflectivity is highly dependent on fabrication error and is hard to simulate accurately. This, combined with the much weaker visibility of the fringes themselves, prohibits us from accurately determining the coupling efficiency solely from fringe visibility. Instead, we calibrate both the input power and the excess losses in the detection train.

The basic setup used to characterize the coupler efficiency is illustrated in Fig. 2.10a. A tunable laser is used to scan the devices in the $1520 - 1570$ nm wavelength range, with a small amount ($\sim 1\%$) of the output power being sent to a wavemeter in order to provide accurate calibration and stabilization of the laser wavelength. The power input to the device is controlled using a variable optical attenuator and the input polarization is adjusted with a fiber polarization controller in order

to maximize reflection from the sample. Power directed via an optical circulator to the lensed fiber tip is reflected by the device and subsequently detected on an optical power meter. Calibration of the input power to the circulator, P_{in}, the circulator input loss, η_{12}, and the loss between port 2 of the circulator and the power meter, $\eta_{2\mathrm{PM}}$, allows us to compare the reflected power detected on the power meter PM directly to the input power and extract the device's reflection efficiency. We typically measure a single-pass efficiency of $\eta_{\mathrm{cpl}} \approx 50\%$, though in several devices we have measured efficiencies as high as $\eta_{\mathrm{cpl}} = 68\%$, very nearly reaching the ideal efficiency possible with this coupler design.

A representative scan of the calibrated reflection from an end-fire coupled cavity is shown in Fig. 2.10b. Due to the relatively flat off-resonance reflection spectrum in these devices, the cavity does appear as an ideal Lorentzian dip, and as such it is impossible to determine $\kappa_{\mathrm{e}}/\kappa$ simply from the reflected power due to the single-sided nature of the coupling scheme, as explained in section 2.4.2. Instead, the laser is locked several GHz off-resonance with the cavity using the wavemeter and sent through an electro-optic phase modulator. Driving the phase modulator with a vector network analyzer (VNA) creates an optical sideband which can be swept from $0 - 8$ GHz across the optical resonance. For this measurement the reflected optical signal is switched to a high-speed photodetector connected to the VNA, rather than the optical power meter used to calibrate reflection efficiency. The optical sideband beats against the carrier frequency to generate an RF signal measured by the high-speed detector and fed back into the VNA. By calibrating this signal against a sweep taken very far-detuned from the cavity (such that the cavity is not present within the 8 GHz range of the sideband sweep), the phase and amplitude response of the optical detection train can be measured and used to normalize the signal, yielding the amplitude and phase response of the cavity. After determining the cavity resonance frequency and total linewidth κ from the optical reflection spectrum, the phase response (Fig. 2.10c) may be fit to obtain $\kappa_{\mathrm{e}}/\kappa$. For the device shown here, the overall decay rate was $\kappa \approx 820$ MHz ($Q \approx 230,000$) and $\kappa_{\mathrm{e}}/\kappa \approx 0.5$. Typically, once we determine a typical as-fabricated intrinsic decay rate κ_{i}, an appropriate cavity-waveguide gap g_{cpl} (see section 2.3.2) is chosen to aim for critical coupling ($\kappa_{\mathrm{e}}/\kappa = 0.5$), which offers a good compromise between cavity-waveguide coupling efficiency and overall cavity decay rate.

Chapter 3

Optomechanical Spectroscopy in a Dilution Refrigerator

3.1 Balanced Heterodyne Spectroscopy

As discussed in section 1.1, driving the optomechanical cavity with a coherent laser tone at frequency ω_{l} will generate mechanically induced optical noise sidebands at frequencies $\omega_{\mathrm{l}} \pm \omega_{\mathrm{m}}$. In principle one can measure these sidebands by direct detection of the reflected optical signal with a sufficiently high-speed photodetector. Upon detection, the reflected carrier acts as a local oscillator, beating against the noise sidebands to produce an RF electrical signal oscillating at ω_{m} that can be detected on an electronic spectrum analyzer. While this method is technically simple and was successfully utilized in the initial experiments involving the nanobeam OMCs within the Painter group [13, 40, 45, 105], it does possess serious limitations from a signal-to-noise (SNR) standpoint. In particular, the shot noise power spectral density scales with the power of the optical carrier, and the drive must be sufficiently intense in order to amplify this optical noise above the electronic noise floor of the detector. For the high speed photodetectors utilized in previous experiments, a carrier power on the order of 1 mW is necessary in order for the optical shot noise to exceed the electronic noise floor. As the nanoscale optical cavities measured in this work cannot handle this amount of input power, post-cavity amplification using an erbium-doped fiber amplifier (EDFA) is necessary. While the amplifier allows us to beat the electronic noise in the setup, it also adds at least 3 dB of extraneous optical noise (and in practice much more) to the signal [139], greatly reducing the SNR. As the mechanically induced sidebands scale in amplitude with both average phonon number and intracavity photon number, it is crucial to improve the detection sensitivity if we hope to make low power measurements in or near the mechanical ground state.

Obviously, the solution to this problem is to split off most of the laser power prior to the cavity input, recombining it later with the reflected cavity signal before detection. In this way, the local oscillator power can be made as high as necessary in order to amplify the optical shot noise without

increasing the input power to the optical cavity. For maximum sensitivity one typically performs a balanced measurement, where the local oscillator and signal are combined with a 50/50 beamsplitter and the outputs are directed to a pair of identical photodetectors. By measuring the difference of the two resulting photocurrents, the amplified noise of the local oscillator is rejected so that only the amplified noise of the reflected cavity signal is observed [11, 12]. This allows for an ideal shot-noise limited measurement of the kind considered in section 1.1.4.

In general the local oscillator frequency can be shifted from that of the drive tone, which results in the two mechanical sidebands being mixed down to different frequencies. If the anti-Stokes and Stokes sidebands are at frequencies $\omega_\mathrm{l} \pm \omega_\mathrm{m}$, and the local oscillator has frequency $\omega_\mathrm{l} + \omega_\mathrm{m} + \delta\omega$, one finds that the anti-Stokes sideband will mix down to $\delta\omega$ while the Stokes sideband will be mixed to $2\omega_\mathrm{m} + \delta\omega$, and vice versa if the local oscillator has frequency $\omega_\mathrm{l} - (\omega_\mathrm{m} + \delta\omega)$. If we select $\delta\omega \ll \omega_\mathrm{m}$ we may then use a much lower bandwidth amplifier, with correspondingly larger gain and lower noise, in order to measure the signal. This also allows us to utilize commercially available balanced photodetector units, which typically have bandwidths on the order of 100 MHz. The one caveat is that we must select the appropriate local oscillator frequency shift in order to mix down the sideband of interest (depending on whether we are driving with blue or red detuning), but in practice this is simple, as only one sideband will be resonantly enhanced for a detuned, sideband-resolved measurement.

3.2 Experimental Setup

The full experimental setup for balanced heterodyne mechanical spectroscopy and thermometry is shown in Fig. 3.1. A fiber-coupled, wavelength-tunable external cavity diode laser is used as the light source, with a small percentage (1%) of the laser output sent to a wavemeter (λ-meter) for accurate determination of laser frequency. The wavemeter may also be used to provide feedback to the laser controller, allowing us to lock the laser at a specified wavelength to within roughly 0.5 pm ($\sim$ 60 MHz). The remaining laser power is split into two paths: a low-power ($\lesssim$ 20 μW) path to be used to drive the OMC cavity, and a high-power ($\sim$ 1 mW) path to be used as the local oscillator (LO) in a balanced optical heterodyne measurement.

The drive light is sent through an electro-optic amplitude modulator (a-m), which is used both to stabilize the signal intensity and for calibration purposes, as described below in section 3.2.1, and a variable optical attenuator (VOA), which allows control of the drive power sent to the cavity over a 60 dB range. The drive is subsequently sent into port 1 of an optical circulator that directs light to the dilution refrigerator, where the fiber terminates with a lensed tip for end-fire coupling to the device as described in section 2.5. The cavity reflection then circulates to a variable coupler (VC) where it is recombined with the LO path before being directed to a pair of balanced photodiodes

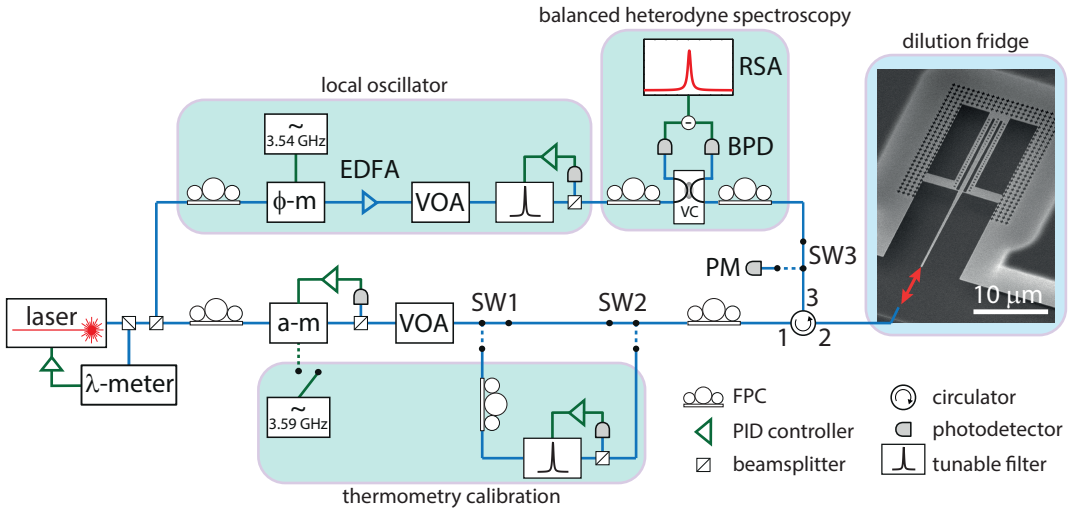

Figure 3.1: **Experimental setup.** Full setup for performing calibrated optical heterodyne spectroscopy and thermometry of the OMC. λ-meter: wavemeter, a-m: electro-optic amplitude modulator, ϕ-m: electro-optic phase modulator, VOA: variable optical attenuator, SW: optical switch, EDFA: erbium-doped fiber amplifier, PM: optical power meter, VC: variable coupler, BPD: balanced photodiode pair, RSA: real-time spectrum analyzer, FPC: fiber polarization controller.

(BPD). The difference photocurrent of the two inputs is then amplified and the NPSD of the resulting electrical signal is measured on a real-time spectrum analyzer (RSA).

The LO is first sent through an electro-optic phase modulator (ϕ-m), which is driven by an RF signal generator in order to produce optical sidebands detuned by $\pm(\omega_{\mathrm{m}}/2\pi - 50\ \mathrm{MHz})$ from the laser frequency for the purposes of mixing the mechanically induced signal modulation down to within the detector bandwidth. An EDFA is used to amplify the LO, while the final LO power reaching the BPD is controlled via another VOA. Finally, a high-finesse tunable Fabry-Perot filter is used to select the appropriate LO sideband and reject the carrier.

Throughout the setup, fiber polarization controllers (FPCs) are utilized to optimize the input polarization to various polarization sensitive components (e.g., the electro-optic modulators, the OMC cavity, etc.). To stabilize the amplitude modulator and lock the tunable Fabry-Perot cavity, inline fiber optic power monitors (represented in Fig. 3.1 by beamsplitters combined with a photodetector) are used to detect the transmitted power, and the resulting signals are fed back to the devices via PID controllers.

In order to reject the noise of the LO (which includes not only shot noise but excess technical noise added by the EDFA), it is crucial to keep the BPD inputs balanced. This is done by adjusting the variable coupler while monitoring the DC output of each photodiode until they give identical readings at full LO power. If the coupler is initially very unbalanced the LO power must be slowly ramped up while adjusting the coupling ratio, in order to avoid a large power imbalance on the two photodiodes, which can overload the difference current amplifier. We may verify that the system

is shot-noise limited by measuring the noise floor versus LO power. In the absence of excess noise, the optical noise floor (i.e., the noise floor with the background electronic noise subtracted) should scale linearly with LO power. The LO power may then be adjusted as high as necessary to elevate the optical vacuum noise above the electronic noise floor. For this setup, an LO power of roughly 1 mW is sufficient to amplify the shot-noise at least 10 dB above the electronic noise.

3.2.1 Thermometry Calibration

As shown in section 1.1.4, in the case of a shot-noise limited measurement the single-sided NPSD of the noise photocurrent for a drive placed on the red-sideband ($\Delta = \omega_{\mathrm{m}}$) is given by

$$S_{\mathrm{II}}(\omega) = S_{\mathrm{dark}} + \frac{G_{\mathrm{e}}^2}{R_{\mathrm{L}}} S_{\mathrm{vac}}^2 \left[1 + 2\eta \frac{\kappa_{\mathrm{e}}}{\kappa} \gamma_{\mathrm{OM}} \bar{S}_{\mathrm{bb}}(\omega)\right], \tag{3.1}$$

where S_{dark} is the electronic NPSD of the detector and RSA, $S_{\mathrm{vac}} = \sqrt{2\hbar\omega_{\mathrm{o}} P_{\mathrm{LO}}}$ is the optical vacuum noise NPSD arising from P_{LO} of LO optical power at optical frequency ω_{o}, G_{e} is the conversion gain between optical power and voltage, R_{L} is the input impedance of the RSA, η is the total detection efficiency, κ_{e} and κ are the extrinsic and total decay rates of the optical cavity, γ_{OM} is the optomechanical scattering rate defined in section 1.1.1, and $\bar{S}_{\mathrm{bb}}$ is the symmetrized mechanical spectral density defined in section A.3 (which in this case will be downshifted in frequency due to heterodyning so that the peak is located at $\omega \approx 2\pi\times(50\text{ MHz})$ rather than at $\omega = \omega_{\mathrm{m}}$). The total noise floor $S_{\mathrm{noise}} = S_{\mathrm{dark}} + \left(G_{\mathrm{e}}^2/R_{\mathrm{L}}\right) S_{\mathrm{vac}}^2$ may be measured by blocking the signal path, while S_{dark} may be measured by blocking both the signal and LO, allowing us to extract the optical NPSD normalized to the shot-noise level. We may then determine the average phonon occupancy of the mechanical resonator $\langle n \rangle$ by integrating this normalized NPSD and using the fact that $\int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \bar{S}_{\mathrm{bb}}(\omega) = \langle n \rangle$, provided we know η, $\kappa_{\mathrm{e}}/\kappa$ and γ_{OM} (see Eq.1.34).

The waveguide cavity coupling ratio $\kappa_{\mathrm{e}}/\kappa$ is determined from the phase and amplitude response of the cavity as described in section 2.5.2, and γ_{OM} may be determined by observing the optomechanical backaction as a function of intracavity photon number, as described below in section 3.3, leaving only the detection efficiency of the setup to be determined. The total detection efficiency is given by $\eta = \eta_{\mathrm{cpl}}\eta_{\mathrm{2det}}\eta_{\mathrm{det}}$, where η_{cpl} is the single-pass transmission efficiency of the end-fire coupler, η_{2det} is the transmission efficiency between port 2 of the circulator and the heterodyne detector, and η_{det} is the efficiency of the detector itself. The determination of the single-pass end-fire efficiency is detailed in section 2.5.2, and in this case was found to be $\eta_{\mathrm{cpl}} = 34\%$, which is a rather low value compared to the typical values measured at room temperature ($\eta_{\mathrm{cpl}} \approx 50\%$ or higher). It was later determined that additional loss was incurred inside the fridge when cold. During this experiment the fibers inside the fridge were tightly wrapped around a series of mandrels in order to keep the excess fiber length neatly organized. While the diameter of each mandrel is large enough (~ 63.5 mm diameter) to avoid

significant macrobending loss in the spooled fiber, this setup induces an additional single-pass loss of roughly 1.5 dB at low temperature. The exact cause of this has not been definitively determined, though our hypothesis is that thermal contraction of the fibers combined with roughness on the surface of the mandrels induces microbending loss in the fibers [140], which can become significant due to the large amount of excess fiber in the fridge. When this extra loss is accounted for, the single-pass end-fire coupling efficiency is around 50%, as expected. The microbending loss can be avoided by spooling the fibers without using mandrels.

The circulator-to-detector efficiency η_{2det} may be measured simply by sending a known optical power into port 2 of the circulator and measuring the transmitted power at the input to the variable coupler, and is found to be $\eta_{\mathrm{2det}} = 84\%$. This leaves only the efficiency of the heterodyne detection itself, which as defined here includes excess insertion loss in the variable coupler, intrinsic quantum efficiency of the BPD, and the alignment of the polarization between the LO and the signal. This is measured by tuning the laser far off-resonance from the optical mode, so that the reflection is only determined by the flat response of the optical coupler, and driving the amplitude modulator to create optical sidebands detuned from the signal by ω_{m}. The optical switches SW1 and SW2 are used to route the signal through a second tunable filter to select a single sideband that is reflected off the device and directed onto the BPD. The power P_{cal} in this sideband is directly measured on a power meter (PM) at SW3, and the photocurrent NPSD as transduced on the RSA is given, in analogy with Eq. 3.1, by

$$S_{\mathrm{II}}(\omega) = S_{\mathrm{dark}} + \frac{G_{\mathrm{e}}^2}{R_{\mathrm{L}}} S_{\mathrm{vac}}^2 \left[1 + \frac{\eta_{\mathrm{det}} S_{\mathrm{cal}}(\omega)}{\hbar\omega_{\mathrm{o}}}\right], \tag{3.2}$$

where S_{cal} is the NPSD of the calibration signal, and $\int_{-\infty}^{\infty} S_{\mathrm{cal}}(\omega) \frac{d\omega}{2\pi} = P_{\mathrm{cal}}$. The detector efficiency may then by extracted via

$$\eta_{\mathrm{det}} = \frac{\hbar\omega_{\mathrm{o}}}{P_{\mathrm{cal}}} \int_{-\infty}^{\infty} \frac{d\omega}{2\pi} \frac{S_{\mathrm{II}}(\omega) - S_{\mathrm{noise}}}{S_{\mathrm{noise}} - S_{\mathrm{dark}}} = 72\%. \tag{3.3}$$

This, combined with the measured device coupling efficiency, yields the overall measurement efficiency η used for calibrated mechanical thermometry

$$\eta = \eta_{\mathrm{cpl}} \eta_{\mathrm{2det}} \eta_{\mathrm{det}} = 20\%. \tag{3.4}$$

One detail that requires caution when calibrating the setup is the power meter used to measure P_{cal}. Most of the efficiency calibrations performed in this work involve measuring the input power to a device and measuring the output power using the same power meter. As such, it is not typically necessary to have an absolutely accurate power meter. Rather, it is sufficient merely for the power meter reading to be consistent and linear with actual measured power. However, calibration of η_{det}

requires comparing the power reading on the BPD unit to the power reading on a separate power meter, which does require an absolute accuracy. Many fiber-coupled power meters are inadequate for this task, as they consist of a fiber-to-free space coupler attached to a free-space power meter. It is relatively easy for the lens inside the fiber-to-free space coupler to become slightly misaligned, leading to inaccurate absolute power readings even if the underlying free-space detector is properly calibrated. We found it necessary to employ a fiber-coupled integrating sphere power meter, which does not suffer from this limitation. This power meter is also used to calibrate the input power to the device, as it is necessary to know the absolute input power in order to accurately determine the intracavity photon number.

3.2.2 Effects of Phase Noise

As discussed in section 1.1.6.1, the presence of laser phase noise can cause substantial systematic errors in calibrated heterodyne thermometry due to noise squashing and anti-squashing effects. However, previous measurements of the phase noise of the laser used in this experiment [23] (a model 6728 New Focus Velocity laser) have shown that the phase noise spectral density at the mechanical frequency $\omega_{\mathrm{m}}/2\pi = 3.6$ GHz is approximately $S_{\phi\phi}(\omega_{\mathrm{m}}) \approx 7.8 \times 10^{-17}$ Hz^{-1}. From this, we can calculate the effective phase noise occupancy n_ϕ as defined in section 1.1.6. The input laser amplitude (in units of $\sqrt{\text{photons}}$/ s) α_{in} can be related to the intracavity photon number n_{c} by the relation $|\alpha_{\mathrm{in}}|^2 \approx \omega_{\mathrm{m}}^2 n_{\mathrm{c}}/\kappa_{\mathrm{e}}$, assuming sideband resolution ($\omega_{\mathrm{m}} \gg \kappa/2$) and a laser detuning of $\Delta = \pm\omega_{\mathrm{m}}$. Then, the phase noise occupancy is given by $n_\phi = |\alpha_{\mathrm{in}}|^2 S_{\phi\phi}(\omega_{\mathrm{m}}) \approx \omega_{\mathrm{m}}^2 S_{\phi\phi}(\omega_{\mathrm{m}}) n_{\mathrm{c}}/\kappa_{\mathrm{e}}$. For the largest intracavity photon number used in this experiment, $n_{\mathrm{c}} \approx 200$, this yields $n_\phi \approx 8 \times 10^{-3}$. The resultant heating of the mechanics and noise squashing/anti-squashing of the detected heterodyne signal can be calculated using the formulae given in section 1.1.6.1, and is found to be negligible on the level of the occupancies measured here ($\langle n \rangle \sim 1$).

3.2.3 Dilution Refrigerator Details

The dilution refrigerator used in this work is a commercial cryogen-free system (BlueFors BF-LD250), which utilizes a pulse tube cryocooler to pre-cool to a fridge base temperature of $T_{\mathrm{f}} \approx 4$ K before adding the ^{3}He/^{4}He mixture to achieve a base temperature of $T_{\mathrm{f}} \approx 10$ mK. While the mixture is circulating, a set of gas-gap heat switches between the mixing chamber and still stages allow the base temperature to be raised from 10 mK up to $T_{\mathrm{f}} \approx 700$ mK (and points in-between). As shown in Fig. 2.9a and b, the microchip sample is mounted to the mixing chamber of the dilution refrigerator on a copper bracket assembly. The bracket is located 10 cm from a calibrated ruthenium oxide resistive temperature sensor. The temperature readings from this sensor are used for comparison and validation of our calibrated optomechanical thermometry measurements, with accuracies of

50 mK near temperatures of 4 K, and 0.06 mK near 10 mK.

3.3 Device Characterization at 4 K

The first step with any optomechanical cavity is a full characterization of the relevant cavity parameters. As many of the cavity's material properties, such as refractive index and photoelastic constants, can shift as the temperature is lowered from room temperature, the initial characterization is performed at 4 K, utilizing only the pulse tube to cool the sample. At 4 K, the rate of change of most material parameters becomes negligible, so the cavity parameters extracted at this temperature will remain valid as the system cools to sub-kelvin temperatures (with the exception of the intrinsic mechanical decay rate γ_i, as we will see). Basic optical characterization of the end-fire coupled cavities is described in section 2.5.2, from which we extract the optical resonance frequency ω_c and the extrinsic and total optical decay rates κ_e and κ. For the device measured in this experiment, the optical cavity parameters are $\omega_\mathrm{c}/2\pi = 194$ THz ($\lambda_\mathrm{c} \approx 1545$ nm), $\kappa_\mathrm{e}/2\pi = 153$ MHz, and $\kappa/2\pi = 529$ MHz. These parameters, along with η_cpl, allow us to determine the intracavity photon number n_c for a known input power and detuning, according to Eq. 1.11.

Once the optical parameters are characterized, we use the optical heterodyne detection scheme discussed in section 3.2 to measure the motion of the localized "breathing" mode at frequency $\omega_\mathrm{m}/2\pi = 3.6$ GHz. As discussed in section 1.1.4, in the case of a drive red-detuned from cavity resonance by $\Delta = \omega_\mathrm{m}$, the resultant NPSD as transduced on a spectrum analyzer yields a Lorentzian signal proportional to the average phonon occupancy $\langle n \rangle$, with a linewidth equal to the total mechanical damping rate $\gamma = \gamma_\mathrm{i} + \gamma_\mathrm{OM}$, where $\gamma_\mathrm{OM} = 4g_0^2 n_\mathrm{c}/\kappa$ is the optomechanically induced damping rate (see section 1.1.3). For a blue-detuned drive ($\Delta = -\omega_\mathrm{m}$), the resonant component of the NPSD is proportional to $\langle n \rangle + 1$ and $\gamma = \gamma_\mathrm{i} - \gamma_\mathrm{OM}$.

The coupling rate g_0 can be determined by observing the dependence of the mechanical linewidth on n_c for both red ($\Delta = \omega_\mathrm{m}$) and blue ($\Delta = -\omega_\mathrm{m}$) drive detunings, as shown in Fig. 3.2a. Above a threshold value, $n_\mathrm{c} > n_\mathrm{thr} \approx 1.5$, optomechanically driven self-oscillation of the mechanical resonator occurs for blue detuning. Below this value, the optomechanical damping γ_OM can be found from the difference between the red- and blue-detuned linewidths. In particular, due to the equal and opposite contributions of γ_OM to the total mechanical linewidth we can average the measured linewidths for red- and blue-detuned drives to obtain γ_i, which in this case has a minimum value of $\gamma_\mathrm{i}/2\pi \approx 13$ kHz. Subtracting this result from either the red- or blue-detuned linewidth to obtain γ_OM. A linear fit of the derived γ_OM versus n_c yields a coupling rate of $g_0/2\pi = 840$ kHz. Using this value of g_0 along with the calibrated optical detection efficiency, the mechanical mode occupancy versus n_c is extracted from the area under the resonant part of the measured NPSD (see Fig. 3.2b and Fig. 3.2c). At high n_c the cooling of the mechanical mode is observed, whereas at low n_c the calibrated occupancy

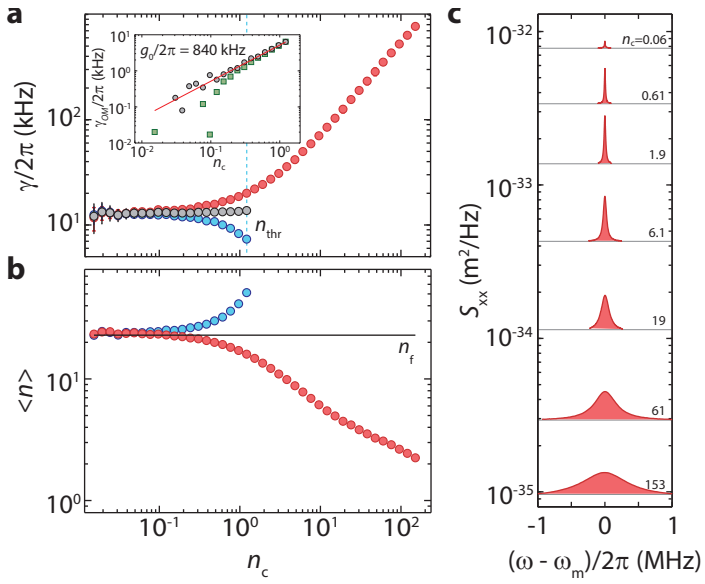

Figure 3.2: **Mechanical characterization at $T_\mathrm{f} = 4$ K. a**, Measured mechanical linewidth γ for $\Delta = \omega_\mathrm{m}$ (red) and $-\omega_\mathrm{m}$ (blue) at a fridge temperature of $T_\mathrm{f} = 4$ K. The vertical blue dashed line indicates the threshold n_c beyond which the mechanical resonance self-oscillates for $\Delta = -\omega_\mathrm{m}$, resulting in a 40 dB increase in the mechanical signal level. Black circles indicate the values of γ_i obtained by taking the average of the detuned data. The inset shows γ_OM determined by subtracting γ_i from the red-detuned γ (circles), and from cooperativity C using the calibrated $\langle n \rangle$ (squares). A linear fit (red line) yields $g_0/2\pi = 840$ kHz. **b**, Calibrated mechanical mode occupancy $\langle n \rangle$ versus intracavity photon number n_c. Blue and red circles are measured with drive laser detunings $\Delta = \pm\omega_\mathrm{m}$, respectively. The mode occupancy n_f corresponding to $T_\mathrm{f} = 4$ K is indicated by a black solid line. **c**, Series of red-detuned NPSD for range of n_c. Here the NPSD is plotted as $S_\mathrm{xx} = x_\mathrm{ZPF}^2 \bar{S}_\mathrm{bb}$, where $x_\mathrm{ZPF} = 4.1$ fm is the zero-point amplitude of the "breathing" mode.

saturates to a constant value in good agreement with the expected mechanical bath occupancy at a fridge base temperature of $T_\mathrm{f} = 4$ K (solid black line). Initially the cooling trend follows the expected result derived in section 1.1.3, where $\langle n \rangle = n_\mathrm{f}(1+C)^{-1}$ ($C = \gamma_\mathrm{OM}/\gamma_\mathrm{i}$ is the cooperativity), while at high n_c the cooling deviates from this ideal behavior and begins to taper off as a function of n_c. This behavior has been observed in previous attempts to laser cool nanobeams into the ground state [45], and is attributed to a combination of n_c-dependent heating of the mechanical bath and an increase in γ_i with n_c. The latter effect is observable, though very weak, in the low power γ_i data shown in Fig. 3.2a. The n_c-dependence of both the bath temperature and γ_i will become much more significant as the fridge base temperature is lowered, and will be discussed in greater detail in the following sections.

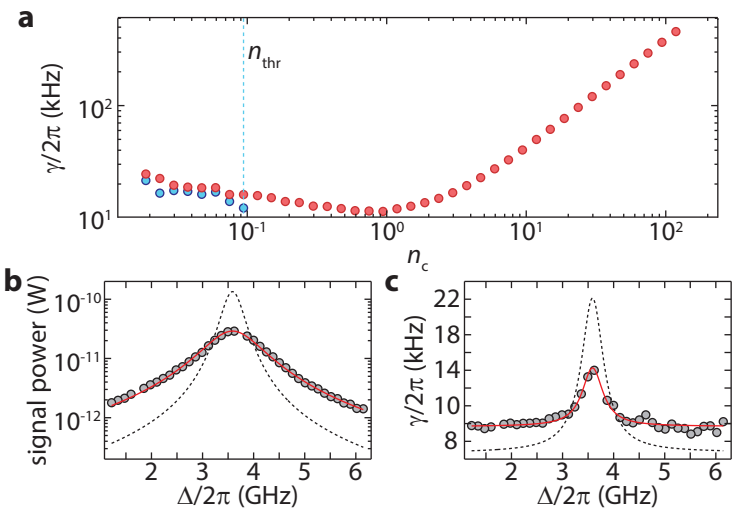

Figure 3.3: **Mechanical frequency jitter at $T_f = 185$ mK. a**, Mechanical linewidth versus n_c for red- (red circles, $\Delta = \omega_m$) and blue-detuned (blue circles, $\Delta = -\omega_m$) drives at $T_f = 185$ mK. The blue dashed line indicates the self-oscillation threshold. **b**, Measured resonant signal power (circles) versus optical detuning, Δ, for $n_c \sim 1.5$. The red curve shows a best fit to the data, yielding $C = \gamma_{OM}/\gamma_i = 3.9$. The black dashed curve shows the best fit in the absence of backaction cooling ($C \to 0$), consistent with assuming γ_i is equal to the measured time-averaged linewidth. **c**, Corresponding measured (circles) mechanical linewidth versus Δ. The red curve is a fit with C constrained ($C = 3.9$), but assuming a Voigt lineshape with additional Gaussian frequency jitter term. The dashed black curve is the best fit for C constrained but with no additional frequency jitter term.

3.4 Mechanical Frequency Jitter

As the fridge temperature is lowered into the sub-kelvin range, the behavior of the observed mechanical signal begins to deviate sharply from the near ideal behavior observed at 4 K and above. The first thing we notice when performing mechanical spectroscopy is a very different dependence of measured linewidth on optical drive power. In particular, the measured mechanical linewidth versus n_c, shown in Fig. 3.3a for $T_f = 185$ mK, increases with decreasing drive power below an apparent minimum at $n_c \sim 1$, rather than reaching a fixed intrinsic value. However, there are several clues that the observed linewidth in this range is not the true mechanical damping rate. For one, the observed self-oscillation threshold (which occurs when $\gamma_{OM} \approx -\gamma_i$) occurs at $n_{thr} \approx 0.1$, an order of magnitude lower than the observed threshold at $T_f = 4$ K. As κ and g_0 should not change with decreasing temperature, γ_{OM} as a function of n_c should be the same as in Fig. 3.2a, implying that γ_i is indeed much smaller. Furthermore, the observed linewidths for red- and blue-detuned drives are identical at n_{thr}, and are in fact larger than the minimum observed linewidth. These inconsistencies indicate that the linewidth associated with the true energy decay rate of the mechanics is likely obscured in the time-averaged spectrum (from which the measured values in Fig. 3.3a are extracted) due to frequency jitter [141, 142]. As a result of the long averaging times (minutes to

hours) required at low optical drive power, direct observation of the mechanical frequency jitter is not possible. However, indirect confirmation of the frequency jitter can be ascertained by studying the detuning (Δ) dependence of the transduced signal power and linewidth for a fixed n_{c}, as shown in Fig. 3.3b and c.

By adjusting the input power to keep n_{c} constant (in this case, $n_{\mathrm{c}} \approx 1.5$), such a measurement keeps constant any effects, such as optical heating or frequency jitter, that might depend on the intracavity photon number. To good approximation (see Eq. 1.18), when $\omega_{\mathrm{m}}/\kappa \gg 1$ and $(\Delta - \omega_{\mathrm{m}})^2 \gg (\Delta + \omega_{\mathrm{m}})^2$, the detuning dependent optomechanical contribution to the mechanical damping is given by

$$\delta\gamma(\Delta) \approx \gamma_{\mathrm{OM}} \frac{1}{1 + 4(\Delta - \omega_{\mathrm{m}})^2/\kappa^2}, \tag{3.5}$$

while the transduced heterodyne signal power (i.e., the integrated power in the sideband noise peak) will be proportional to $\delta\gamma\langle n\rangle \propto \delta\gamma(\Delta)[1 + \delta\gamma(\Delta)/\gamma_{\mathrm{i}}]^{-1}$. When there is negligible backaction (i.e., $C = \gamma_{\mathrm{OM}}/\gamma_{\mathrm{i}} \to 0$), the dependence of the total signal power on Δ is simply a Lorentzian with a width equal to the cavity linewidth κ, while the response broadens for increasing cooperativity C. As seen in Fig. 3.3b, a fit to the observed signal power assuming negligible backaction (black dashed line) provides an extremely poor fit to the data, while allowing C to vary freely leads to a best-fit result (red line) of $C = 3.9$. As n_{c} and κ are known exactly, this constrains the quantity $g_0^2/\gamma_{\mathrm{i}}$. In principle we do not expect g_0 to vary appreciably at temperatures below 4 K, and a linear fit to the observed γ in Fig. 3.3a for the large n_{c} yields the same value of g_0, so one may infer γ_{i} from the constrained value of C alone. However, while we can directly measure κ at low n_{c} via a sufficiently low power wavelength scan of the cavity, g_0 can only be measured directly once n_{c} is large enough that $\gamma_{\mathrm{OM}} \gtrsim \gamma_{\mathrm{i}}$. As large n_{c} can (and, as we will see in the next section, does) heat the cavity via optical absorption, we cannot rule out the possibility that any temperature dependence of g_0 is merely masked by heating effects at large n_{c}. Consequently, in order to separately determine both g_0 and γ_{i} we also fit the linewidth dependence on Δ.

The total linewidth is ideally given as a function of detuning by $\gamma(\Delta) = \gamma_{\mathrm{i}} + \delta\gamma(\Delta)$, where $\delta\gamma(\Delta)$ is approximately given by Eq. 3.5. Given the known value of C, we can fit this function to the observed linewidth data shown in Fig. 3.3c by using γ_{i} as the fitting parameter while using the constraint $C = 3.9$ to determine g_0, which results in a poor fit to the data (black dashed line). On the other hand, in the presence frequency jitter corresponding to a random Gaussian process the mechanical noise peak will technically be a Voigt lineshape (the convolution of a Gaussian with a Lorentzian) rather than a pure Lorentzian, with a linewidth approximately given by [143]

$$\gamma(\Delta) \approx 0.5246(\gamma_{\mathrm{i}} + \delta\gamma(\Delta)) + \sqrt{0.2166(\gamma_{\mathrm{i}} + \delta\gamma(\Delta))^2 + \gamma_{\mathrm{G}}^2}, \tag{3.6}$$

where γ_{G} is the full width at half maximum of the Gaussian process describing the frequency jitter.

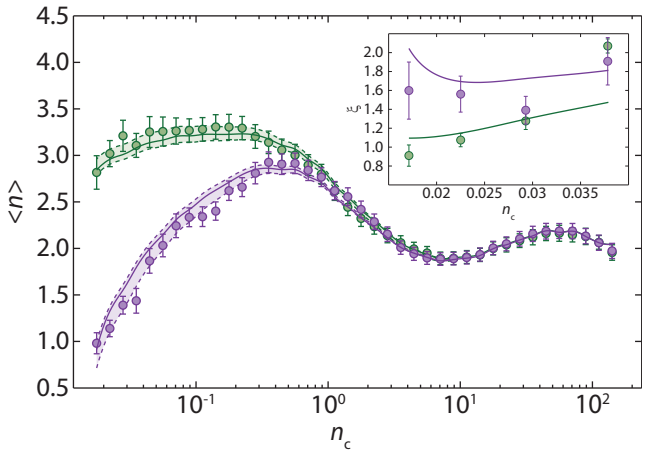

Figure 3.4: **Mechanical occupancy at $T_\mathrm{f} = 10$ mK and 635 mK.** Calibrated mode occupancy $\langle n \rangle$ versus n_c for $T_\mathrm{f} = 10$ mK and 635 mK. A best fit to both temperature data sets using the proposed phenomenological heating model are shown as solid curves, with shaded region representing the variation in the fit for $\gamma_0/2\pi = 306 \pm 28$ Hz. The inset shows the measured asymmetry parameter ξ as a function of n_c, with the prediction of the best fit model shown as solid curves. $T_\mathrm{f} = 10$ mK data/fits are shown as purple circles/solid curves. $T_\mathrm{f} = 635$ mK data/fits are shown as green circles/solid curves.

Fitting this functional form to the observed linewidth data, using γ_i and γ_G as fitting parameters while keeping C fixed by adjusting g_0 as before, we obtain an excellent fit (red line) that yields $\gamma_\mathrm{i}/2\pi = 2.3$ kHz for the intrinsic mechanical energy decay rate, $\gamma_\mathrm{G}/2\pi = 6.1$ kHz for the Gaussian frequency jitter, and $g_0/2\pi = 805$ kHz, consistent with the $T_\mathrm{f} = 4$ K value as expected. These measurements confirm the presence of n_c dependent frequency jitter in the mechanics, as well as the fact that γ_i is in fact much lower at lower base temperatures while g_0 does not vary appreciably. While the method applied here works well for determining γ_i at moderate n_c, for very small n_c it is ineffective as the amount of optomechanical backaction is too small to accurately determine C from measurements of the total signal power versus Δ. Without this constraint, one cannot obtain reliable estimates for γ_i, as there is a much wider range of γ_i and γ_G values which can provide good fits to the linewidth curve. In what follows, we will instead determine γ_i by fitting the heating behavior of the device over the full range of n_c.

3.5 Optically Induced Heating

An even more striking deviation from the behavior at $T_\mathrm{f} = 4$ K is observed in the calibrated occupancy of the mechanical resonator. The calibrated mechanical mode occupation $\langle n \rangle$ for a red-detuned ($\Delta = \omega_\mathrm{m}$) drive is plotted against n_c in Fig. 3.4 for $T_\mathrm{f} = 10$ mK (purple) and 635 mK (green). Both curves exhibit a series of dramatic heating and cooling trends, and in fact coincide for $n_\mathrm{c} \gtrsim 1$.

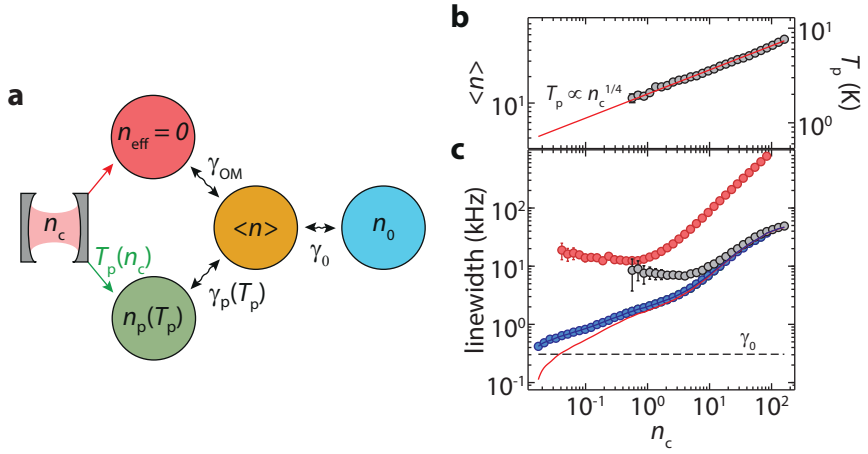

Figure 3.5: **Heating model.** **a**, Schematic showing the phenomenological three-bath model for the mechanical occupancy. See text for details. **b**, Measured phonon occupancy versus intracavity photon number for a resonant drive at $T_f = 10$ mK. The red line shows a power law fit to the occupancy. The right axis shows the equivalent bath temperature, T_p. **c**, Measured mechanical linewidth versus n_c for resonant (gray circles) and red-detuned (red circles) drives. The blue circles show the best fit values of $\gamma_i = \gamma_0 + \gamma_p$ to both $T_f = 10$ and 635 mK data sets. The best fit value for γ_0 and a smooth curve fit to the inferred γ_p are shown as a black dashed line and a solid red line, respectively.

At the lowest optical drive powers ($n_c = 0.016$) and lowest fridge temperature ($T_f = 10$ mK), the calibrated phonon occupancy reaches a minimum $\langle n \rangle = 0.98 \pm 0.11$, corresponding to $T \approx 270$ mK. The complex behavior of these two cooling curves can be understood by incorporating both n_c dependent heating of the thermal bath as well as temperature (and thus n_c) dependence of the intrinsic damping rate γ_i into the equations determining $\langle n \rangle$. Indeed, such effects have already been seen to play a small role at large n_c in previous experiments at $T_f = 4 - 15$ K (see Ref. [45] as well as comments in section 3.3). Due to the sharp drop in thermal conductance with decreasing temperature [144], heating of the mechanical mode by optical absorption is expected to be much more significant at sub-kelvin base temperatures.

3.5.1 Phenomenological heating model

The simplest approach to modeling heating effects is to use a single thermal bath, which couples to the mechanics at rate γ_i and has average occupancy n_b, both of which are now assumed to have a dependence on n_c. However, more realistically one expects multiple effective thermal baths for the mechanics, corresponding to different physical regions on the chip and coupling via different physical processes. In particular, there exist two fundamental ways in which the localized mechanical mode of interest can exchange energy with its environment. The first is coupling to local phonons of different frequencies, typically via anharmonic phonon-phonon scattering [145], the rate of which strongly depends upon the local phonon temperature. The second is direct coupling to the environment

external to the cavity region via phonon radiation out the ends of the beam. In this case, the coupling rate should depend only on the nature of clamping at the end of the beam (i.e., the engineered phononic shield surrounding the cavity), and not on temperature. Moreover, while one expects that optical absorption will create some local population of hot phonons in the cavity region, one expects negligible heating of the environment external to the cavity region, due to the small amount of absorbed optical power, the weak thermal link between the cavity region and the external chip, and the large thermal link between the external chip and the fridge. For these reasons we explicitly separate these two physical processes when modeling the system, leading to the phenomenological three-bath model illustrated in Fig. 3.5a. Here the local breathing mode, with average occupancy $\langle n \rangle$, is coupled to the exterior environmental bath, which has a fixed occupancy n_0, at a constant rate γ_0 which depends only on the properties of the acoustic radiation shield. It is also coupled to a localized hot-phonon bath driven by optical absorption, with effective temperature $T_\mathrm{p}(n_\mathrm{c})$ and corresponding occupancy $n_\mathrm{p}(T_\mathrm{p})$, at a temperature-dependent rate $\gamma_\mathrm{p}(T_\mathrm{p})$. Finally, the optomechanical interaction couples the mechanical mode to an effectively zero temperature bath [1] at rate γ_OM. One should note that, due to finite thermal resistance between the sample and the mixing chamber plate, the exterior bath occupancy n_0 will not necessarily be equal to the fridge occupancy n_f even if excess heating is negligible. For the measurements made in this chapter one finds that the results are not strongly dependent on the specific value of n_0 due to the large value of n_p, so for simplicity we will take $n_0 \approx n_f$.

Under this model, the average mode occupancy is given by [2]

$$\langle n \rangle = \frac{\gamma_0 n_0 + \gamma_\mathrm{p}(T_\mathrm{p})\, n_\mathrm{p}(T_\mathrm{p})}{\gamma_0 + \gamma_\mathrm{p}(T_\mathrm{p}) + \gamma_\mathrm{OM}(n_\mathrm{c})}, \tag{3.7}$$

where we have explicitly noted the n_c dependence of γ_OM and where γ_p and n_p are implicitly dependent on n_c via their dependence on the local hot-phonon effective temperature $T_\mathrm{p}(n_\mathrm{c})$. Initially there exist a number of unknowns in this model. In particular, while n_0 can be approximated as n_f and a guess at γ_0 can be made based on FEM simulation of the phononic shield (see section 1.3.5), we have no a priori model for either γ_p or n_p as a function of n_c, making a proper fit of the complex heating and cooling behavior observed in Fig. 3.4 highly error-prone. However, when the optical drive laser is placed on resonance with the cavity ($\Delta = 0$), the optomechanical damping vanishes ($\delta\gamma = 0$, as seen in Eq. 1.18). In this case, for sufficiently large n_c we expect $n_\mathrm{p} \gg n_0$ and therefore $\langle n \rangle \approx n_\mathrm{p}$. From such on-resonance measurements of $\langle n \rangle$, shown in Fig. 3.5b, we find that n_p follows

[1] As discussed in section 1.1.3 and Appendix A.1, the effective bath provided by the optomechanical interaction is not really zero-temperature, but rather has an effective occupancy $n_\mathrm{eff} = (\kappa/4\omega_\mathrm{m})^2$ due to the presence of non-resonant phonon emission. As this is extremely small due to the high degree of sideband-resolution in our system ($n_\mathrm{eff} \approx 0.001$) we simply approximate $n_\mathrm{eff} = 0$.

[2] The derivation of this expression is almost identical to that of a single bath, only now we assume multiple *uncorrelated* thermal baths driving the mechanics in Eq. 1.8. This is similar to the separation of the optical dissipation into extrinsic and intrinsic channels discussed in section 1.1.1.

a simple power law $n_{\rm p} \propto n_{\rm c}^{0.25}$. Unfortunately, due to lack of resonant enhancement by the optical cavity, the sideband scattering amplitudes are reduced by a factor of $\frac{\kappa}{2\omega_{\rm m}}$ when driving on-resonance, preventing us from measuring $n_{\rm p}$ directly down to the lowest measured $n_{\rm c}$. However, as discussed below in section 3.6, it is reasonable based on a proposed physical model of the heating process to expect this power law will hold down to our lowest measurement powers, and we extrapolate to obtain $n_{\rm p}$ as a function of $n_{\rm c}$ across the whole measurement range.

With n_0, $n_{\rm p}$ and $\gamma_{\rm OM}$ thus determined, we need only determine $\gamma_{\rm p}$ and γ_0. To do so, we guess an initial value of γ_0 based on FEM simulations of the phononic shield and find a single spline fit to $\gamma_{\rm p}\left(n_{\rm c}\right)$ which best matches the two sets of occupancy curves shown in Fig. 3.4. Iterating this procedure as a function of γ_0 yields the best fit to our phenomenological heating model. The resulting $\gamma_{\rm i}(n_{\rm c}) = \gamma_{\rm p}(n_{\rm c}) + \gamma_0$ curve that best fits the measured $\langle n \rangle$ data for both $T_{\rm f} = 10$ mK and 635 mK is plotted in Fig. 3.5c (blue circles). Also shown in Fig. 3.5c are the best fit value of the coupling to the fridge bath $\gamma_0/2\pi = 306 \pm 28$ Hz (dashed black horizontal curve) and a smooth spline curve fit to the inferred values of $\gamma_{\rm p}(n_{\rm c})$ (red solid curve). A plot of the best fit model is shown alongside the measured $\langle n \rangle$ cooling curves in Fig. 3.4, with the associated shading regions showing the variation in the model fit within the uncertainty in γ_0. In addition to the good agreement of the model for both fridge temperatures, the inferred intrinsic energy damping rate is also consistent with the measured self-oscillation threshold (Fig. 3.3a). At the lowest drive powers ($n_{\rm c} = 0.016$), the energy damping mechanical Q-factor reaches an impressively high value of $Q_{\rm m} = 9 \times 10^6$.

Alongside the calibrated mode occupancy $\langle n \rangle$, we have also measured the sideband asymmetry, ξ, defined in Eq. 1.35, shown in the inset to Fig. 3.4. The asymmetry is sensitive to both the absolute mode occupancy and to the sum of γ_0 and $\gamma_{\rm p}$ through the cooperativity $C = \gamma_{\rm OM}/(\gamma_0 + \gamma_{\rm p})$. Good correspondence can also be seen between the best fit model (solid curves) and the measured ξ (circles).

3.6 Discussion

While the phenomenological model presented in the previous section does not attempt to explain the microscopic physics of the optically induced heating, some details of the underlying physical processes can be elucidated by studying how the frequency jitter and hot phonon coupling rate scale with the local hot phonon temperature.

3.6.1 Temperature Dependence of Frequency Noise

As detailed in section 3.4, at sub-kelvin fridge temperatures frequency noise of the mechanical resonance is seen to dominate the time-averaged measured linewidth for low optical drive powers. In the case of the red-detuned drive, frequency noise is observed for $n_{\rm c} < 1$ before optomechanical

damping kicks in, while the on-resonance measurement shows frequency noise for $n_c < 10$ before the intrinsic energy damping (γ_p) begins to dominate. By comparing to the on-resonance measured mode occupancy (Fig. 3.5b), the frequency noise can be plotted with respect to the optical absorption driven temperature T_p, as shown in Fig. 3.6a. A power law fit of the time-averaged linewidth in the frequency jitter dominated regime shows a frequency jitter scaling proportional to $T_p^{-0.9}$.

Similar inverse temperature scaling of frequency noise has been observed in superconducting microwave resonators [142], with substantial evidence indicating the source to be fluctuations in two-level tunneling states (TLS) of near-field amorphous materials [146,147]. These TLS also couple to phonons, contributing to not only the dielectric properties of the material, but also the elastic properties. As such, the mechanical modes of the devices studied here can be expected to couple to TLS in a similar fashion to microwave electromagnetic resonators. The decrease in frequency noise with temperature can then be attributed to the thermal excitation and saturation of the TLS. The presence of TLS in the Si devices of this work likely stems from the formation of a native oxide on the Si surfaces of the patterned nanobeam. Careful measures, involving a final wet etch in HF acid [148], are made to passivate the Si surfaces of the nanobeam in order to reduce the optical absorption from mid-bandgap surface electronic states. While this cleaning should have been effective in removing the surface oxide as well, a one hour procedure was required to load the sample into the dilution refrigerator, during which time a native oxide would at least partially reform on the Si surfaces of the device. This native oxide is likely to be the source of the TLS. Future work will look to substantially reduce the device loading time after HF surface cleaning, hopefully reducing both the optical absorption and TLS density.

3.6.2 Heating and Damping via Three-Phonon Scattering Processes

Though a detailed microscopic calculation of the additional heating and damping due to optical absorption is beyond the scope of this work, some qualitative insight into the nature of the locally heated mechanical bath and its coupling to the mechanical mode of interest can be obtained from consideration of a simplified model of the phonon-phonon interactions. At the temperatures considered in this work ($T < 10$ K) the mean free path of the thermal phonons is expected to be much larger than the wavelength of the mechanical mode of interest. Consequently, the damping should be described by the Landau-Rumer theory, where losses occur primarily due to three-phonon mixing with the local thermal environment due to anharmonicity in the Si lattice [145, 149]. As shown below, a simple toy model of three-phonon scattering assuming a hot phonon bath consisting of high frequency phonons above some frequency cutoff $\omega_{co} \gg \omega_m$ predicts a universal low temperature scaling $\gamma_p \propto T_p e^{-T_{co}/T_p}$, where $T_{co} = \hbar\omega_{co}/k_B$.

Consider a toy model where two modes, with frequencies ω_1 and ω_2, respectively, are coupled with the mode of interest at frequency ω_m. For now we assume that both auxiliary modes are high-

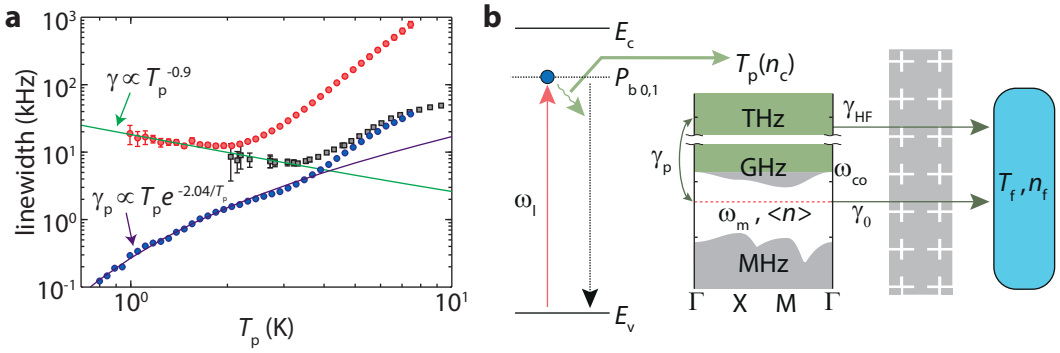

Figure 3.6: **Linewidth scaling and physical heating model.** **a**, Mechanical linewidth vs optical-absorption-driven hot phonon bath temperature T_p as determined from the resonant heating measurement of Fig. 3.5b. The time-averaged measured total linewidth is shown for red-detuned ($\Delta = \omega_\mathrm{m}$, red circles) and resonant ($\Delta = 0$, gray squares) drives, as well as the hot phonon coupling rate γ_p, extracted from the heating model (blue circles). A power law fit of the total measured linewidth corresponding to $T_\mathrm{p}^{-0.9}$ is shown as a solid green curve, and a fit to γ_p at the low temperature end of the curve, corresponding to the three-phonon scattering model discussed in the main text, is shown as a purple curve. **b**, Diagram illustrating the proposed model of optical-absorption-driven heating of the mechanics. See text for details.

frequency, so that ω_1, $\omega_2 > \omega_\mathrm{m}$, and without loss of generality $\omega_1 - \omega_2 = \omega_\mathrm{m}$. As will become clear, the reason for considering this special case, is that it leads to a unique, universal temperature scaling for the three-phonon coupling rate at low temperatures. To first order in perturbation theory, the scattering rates into and out of the mechanical cavity mode due to the lowest order anharmonic interaction can be given by [149] $\Gamma_+ = A(n_\mathrm{m}+1)(n_2+1)n_1$ and $\Gamma_- = A(n_1+1)n_\mathrm{m}n_2$, where A is a constant that depends on the matrix element of the anharmonic potential, and n_1, n_2, and n_m are the number of quanta in each of the three mechanical modes. In the absence of other dissipative processes a simple rate equation for the population of the mechanical cavity mode is given by

$$\dot{n}_\mathrm{m} = \Gamma_+ - \Gamma_- = -A(n_2 - n_1)n_\mathrm{m} + A(n_2+1)n_1, \tag{3.8}$$

which has the same form as the equation for a harmonic oscillator interacting with a bath of occupation n_p with a coupling rate γ_p, where

$$n_\mathrm{p} = \frac{n_2(n_1+1)}{n_2 - n_1}, \quad \gamma_\mathrm{p} = A(n_2 - n_1). \tag{3.9}$$

If the two high-frequency modes are both in equilibrium with each other at some elevated temperature T_p, such that $n_\mathrm{i} = (\exp\left(\frac{\hbar\omega_\mathrm{i}}{k_\mathrm{B}T_\mathrm{p}}\right) - 1)^{-1}$, it is easy to show that n_p is simply given by the Bose-Einstein occupation factor for the mechanical mode at temperature T_p ($n_\mathrm{p} = (\exp\left(\frac{\hbar\omega_\mathrm{m}}{k_\mathrm{B}T_\mathrm{p}}\right) - 1)^{-1}$). We can also see that the scattering rate γ_p will depend on T_p through the temperature dependence of the population difference $n_2 - n_1$. This dependence can be approximately linear or exponential

in $T_{\rm p}$, depending on the value of $\hbar\omega_2/k_{\rm B}T_{\rm p}$.

More realistically, the optical absorption process will populate a range of high-frequency phonon modes above some cutoff frequency ($\omega_{\rm co}$), which can contribute to the heating. On the assumption that these modes come into quasi-equilibrium with each other at some elevated temperature $T_{\rm p}$, we can easily show that that the expression for $n_{\rm p}$ is unchanged, and the effective bath occupancy $n_{\rm p}$ will be given by the above expression. The scattering rate will now be given by

$$\gamma_{\rm p} = \int_{\omega_{\rm co}}^{\infty} d\omega \rho(\omega) A(\omega, \omega_{\rm m}) \left(n(\omega, T_{\rm p}) - n(\omega + \omega_{\rm m}, T_{\rm p}) \right), \tag{3.10}$$

where $\rho(\omega)$ is the density of modes at frequency ω, the matrix element A now can depend explicitly on frequency, and $n(\omega, T)$ is just the Bose-Einstein occupation at frequency ω and temperature T. In general, we can obtain limited analytical results if we assume that the product of the density of states and the matrix element obeys some power law as a function of frequency, $\rho(\omega)A(\omega, \omega_{\rm m}) \propto \omega^a$. For example, under a simple continuum elastic model the product would be given by $\rho(\omega)A(\omega, \omega_{\rm m}) \propto \omega_{\rm m}(\omega_{\rm m} + \omega)\omega^3$ [149]. Making a simple change of variables $x = \frac{\hbar\omega}{k_{\rm B}T_{\rm p}}$, and assuming $\omega_{\rm m} \ll \omega_{\rm co}$, we then arrive at the approximate relation:

$$\gamma_{\rm p} \propto \omega_{\rm m} T_{\rm p}^{a+1} \int_{x_{\rm co}}^{\infty} dx \frac{x^a e^x}{(e^x - 1)^2}; \quad x_{\rm co} = \frac{\hbar\omega_{\rm co}}{k_{\rm B}T_{\rm p}}. \tag{3.11}$$

Considering first the limiting case $x_{\rm co} \gg 1$ (low temperature), we note that the integral can be approximated by

$$\int_{x_{\rm co}}^{\infty} dx \frac{x^a e^x}{(e^x - 1)^2} \approx \int_{x_{\rm co}}^{\infty} dx \frac{x^a}{e^x} = \Gamma(a + 1, x_{\rm co}), \tag{3.12}$$

where $\Gamma(\alpha, z)$ is the upper incomplete Gamma function. For real values of z, this function has the asymptotic behavior $\Gamma(\alpha, z) \to z^{\alpha-1}e^{-z}$ as $|z| \to \infty$. This leads to the approximate scaling law

$$\gamma_{\rm p} \propto T_{\rm p}^{a+1} x_{\rm co}^a e^{-x_{\rm co}} = T_{\rm p} e^{-(T_{\rm co}/T_{\rm p})}, \tag{3.13}$$

where $T_{\rm co} = \frac{\hbar\omega_{\rm co}}{k_{\rm B}}$ and $T_{\rm p} \ll T_{\rm co}$. It is important to note that this limiting form has no dependence on a. As we can always expand the phonon density of states and scattering matrix elements as power series in ω, finding that each term reduces to the same functional dependence on $T_{\rm p}$, this scaling with $T_{\rm p}$ can be taken as universal behavior at sufficiently low temperatures for the coupling to high-frequency phonons.

In the opposite limiting case $x_{\rm co} \ll 1$ (high temperature), the lower limit of the integral can be extended to 0, and a simple integration by parts shows that, for $a > 1$,

$$\int_0^{\infty} dx \frac{x^a e^x}{(e^x - 1)^2} = a \int_0^{\infty} dx \frac{x^{a-1}}{e^x - 1} = a\Gamma(a)\zeta(a), \tag{3.14}$$

where $\Gamma(a)$ is the gamma function and $\zeta(a)$ is the Riemann zeta function. As this is a simple constant, we get the scaling law $\gamma_{\rm p} \propto T_{\rm p}^{a+1}$, when $T_{\rm p} \gg \frac{\hbar\omega_{\rm co}}{k_{\rm B}}$ and $a > 1$.

Though we have considered explicitly the case where both bath phonons participating in the three-phonon scattering are much higher in frequency than $\omega_{\rm m}$, it is straightforward to repeat this analysis for the cases when one or both of the phonons are lower frequency. In this case, one does not find the particular universal low-temperature behavior noted above, but rather a power law which depends on the particulars of the density of states and scattering matrix elements. Thus, the appearance of a bath coupling rate of the particular form in Eq. 3.13 at low temperatures can be taken as evidence (though not definitive proof) of a predominantly high frequency phonon bath.

As displayed in Fig. 3.6a the low temperature end of the inferred $\gamma_{\rm p}(T_{\rm p})$ data fits well to this functional form, suggesting that the local hot phonon bath in our device is indeed composed primarily of a quasi-equilibrium population of high-frequency phonons. The effective cutoff temperature is $T_{\rm co} \approx 2$ K, corresponding to a cutoff frequency of $\omega_{\rm co}/2\pi \approx 35$ GHz. The phonon wavelength in Si at this cutoff frequency is $\lambda_{\rm co} \sim 200$ nm, which is comparable to the cross-sectional dimensions of the nanobeam ($\sim 500 \times 220$ nm). This fact suggests a possible physical explanation for the formation of such a hot phonon population.

Our proposed microscopic picture of the heating process is illustrated schematically in Fig. 3.6b. The source of optical absorption in our structures is most likely due to the well known $P_{\rm b0,1}$ electronic defect states at the surface of Si [148, 150], which lie in the middle of the bandgap of silicon (represented schematically by valence and conduction band energy levels $E_{\rm v}$ and $E_{\rm c}$) in an energy range which encompasses the frequency of our drive laser ($\hbar\omega_{\rm l} \approx 0.8$ eV). Though phonon emission processes from such excited electronic states have not been extensively studied, it is reasonable to assume that phonon assisted emission can occur primarily involving relatively high-frequency phonons in the 100s of GHz to THz range, rather than low-frequency acoustic phonons near the breathing mode frequency. In addition to radiating into the external environment at some fast rate $\gamma_{\rm HF}$ (as these high-frequency phonons are not affected by the phononic shield), the hot phonon population will approach local thermal equilibrium by scattering energy to lower frequencies. However, below the cutoff frequency the nanobeam will appear effectively one-dimensional due to the small dimensions relative to phonon wavelength, and the reduced dimensionality will result in a reduced number of decay channels compared to bulk Si. The corresponding reduced down-conversion rate, in conjunction with the relatively fast pumping and radiative escape of the high-frequency phonons, will result in a population above $\omega_{\rm co}$ (green shaded region) which is in quasi-equilibrium at an elevated effective temperature $T_{\rm p}$ much larger than the effective temperature of the rest of the phonon spectrum (gray shaded region). This type of "1D bottleneck" effect and the corresponding generation of a high frequency hot phonon population under optical pumping conditions has been observed in low-dimensional systems such as carbon nanotubes [151, 152] and graphene [153].

The existence of a predominantly high-frequency hot phonon bath helps to explain the power law behavior of $n_{\rm p}$ versus $n_{\rm c}$ observed in the on-resonance measurements, as phonons with wavelengths small relative to the dimensions of the structure are expected to possess a thermal conductance scaling as $G_{\rm th} \sim T^3$ [154], which leads to a relation $P_{\rm abs} = \int_{T_0}^{T_{\rm p}} dT G_{\rm th}(T) \propto T_{\rm p}^4 - T_0^4$ between the absorbed power $P_{\rm abs}$, the local hot-phonon temperature, and the external temperature T_0. Assuming $P_{\rm abs} \propto n_{\rm c}$ and $T_{\rm p} \gg T_0$, this results in the measured power law $T_{\rm p} \propto n_{\rm c}^{0.25}$. This also helps to justify the extrapolation of the observed power law to lower $T_{\rm p}$. One might expect a transition to a different power law for $G_{\rm th}$ as the bath temperature lowers (e.g., to one-dimensional conduction where $G_{\rm th} \sim T$, as has been observed in dielectric nanowires [155]). However, such transitions take place when the temperature becomes low enough that most of the thermal energy is transported by low frequency phonons with wavelengths comparable to the nanostructure's cross-sectional dimensions. In the situation proposed here, where the excess heat is carried primarily by high-frequency, short wavelength phonons, such a transition would not be expected to occur. Of course, this simple toy model does not capture all of the features of the measured $\gamma_{\rm p}(T_{\rm p})$ curve (in particular, the kink at $T_{\rm p} \approx 3$ K). In addition, the estimated escape time of ballistic high-frequency phonons from the nanobeam ($\gamma_{\rm HF}^{-1} \sim 1 - 10$ ns) suggests that they must quickly come into quasi-equilibrium. Given the slow rates of most bulk relaxation processes [156], this fast thermalization may indicate a unexpectedly large degree of diffusive and inelastic scattering at the surfaces of the patterned nanobeam [157].

Given that we have tentatively attributed the mechanical frequency jitter in our system to dispersive coupling to TLS, one wonders whether the observed temperature dependent damping can be adequately explained by dissipative coupling to TLS rather than by three-phonon scattering as proposed above. In general, the TLS contribution to the damping of an acoustic mode with frequency ω is given by [158]

$$\gamma_{\rm TLS}(\omega) = \frac{D^2}{\rho v^2}\frac{1}{k_{\rm B}T}\int_0^\infty d\Delta \int_0^\infty d\Delta_0 \; f(\Delta,\Delta_0)\left(\frac{\Delta}{E}\right)^2 \mathrm{sech}^2\left(\frac{E}{2k_{\rm B}T}\right)\frac{\omega^2\tau}{1+\omega^2\tau^2}, \tag{3.15}$$

where D is the deformation potential coupling between the TLS and the acoustic mode, ρ and v are the material density and speed of sound, Δ is the TLS energy asymmetry, Δ_0 the TLS tunnel splitting, and $f(\Delta,\Delta_0)$ is the TLS density of states. The total energy E of a given TLS is defined by $E = \sqrt{\Delta^2 + \Delta_0^2}$, while its relaxation time is $\tau \propto \tanh(E/2k_{\rm B}T)/E\Delta_0^2$ with a constant of proportionality dependent on material properties. The specific temperature dependence of this damping rate is highly dependent on both the thermal energy scale $k_{\rm B}T$ relative to the average TLS energy, as well as the specific density of states.

In the case when all TLS have nearly the same energy and relaxation time, so that we may take $f(\Delta,\Delta_0)$ to be approximately a delta function in both Δ and Δ_0, we find that for $T \ll E/k_{\rm B}$ we

obtain $\gamma_{\rm TLS} \propto T^{-1}\ e^{-E/k_{\rm B}T}$. This functional form fits reasonably well to the $\gamma_{\rm p}$ data at low $T_{\rm p}$, with $E \approx 0.34$ meV (~ 80 GHz). While this shows that a TLS model could in principle lead to the type of low temperature damping behavior observed in our devices, it would require a distribution of TLS very different from that typically observed in real materials. In structural glasses (such as silicon oxide) one expects a density of states of the form $f(\Delta, \Delta_0) \approx P/\Delta_0$, where P is a constant, resulting in a broad range of energies and relaxation times which lead to damping rates scaling as T^3 [158]. On the other hand, it has been suggested that TLS in the bulk of crystalline solids such as silicon have a well defined tunnel splitting Δ_0 owing to the ordered nature of the solid, with the density of states $f(\Delta)$ having an approximate Gaussian form with standard deviation Δ_1 dependent upon the variation of strain within the material [159]. Provided $\Delta_1 \ll \Delta_0$, this would be consistent with the observed damping in this work. However, to date experiments studying the effects of TLS on microscale mechanical resonators in crystalline Si and GaAs have observed $\gamma_{\rm TLS} \propto T$ (albeit at lower mechanical frequencies, on the kHz-MHz scale), suggesting a reasonably broad distribution of TLS energy (in particular $\Delta_1/\Delta_0 \sim 100$) [160, 161]. Overall, TLS coupling seems a less plausible explanation for the observed dependence, though it cannot be completely ruled out at this time.

Although significant work remains to determine the exact microscopic details of the optical absorption heating and frequency jitter observed in the measurements of the quasi-1D OMC cavity studied here, there are nonetheless several interesting points to note. Firstly, the phononic shield provides excellent mechanical isolation of the 'breathing' mode, while at the same time providing good mechanical coupling to the fridge bath for heat carrying phonons above the acoustic bandgap. Secondly, although lower phonon occupancies could have been measured using thinner phononic shields, effectively increasing the coupling rate γ_0 to the fridge bath at $T_{\rm f}$, this would come with a commensurate reduction in cooperativity $C = \gamma_{\rm OM}/\gamma_{\rm i}$. Coherent quantum interactions between the optical cavity field and the mechanical resonator require $C > 1$ and $\langle n \rangle < 1$, and although the devices of this work closely approach this limit, a move to quasi-2D Si OMC devices [162] with larger thermal conductance should enable future work in the quantum regime as envisioned in recent proposals [163–166]. Further discussion of this will be given in Chapter 5.

Chapter 4

Pulsed Optomechanical Measurements at Millikelvin Temperatures

The initial cryogenic measurements performed in Chapter 3 have demonstrated that operation deep in the mechanical ground-state is not possible in the steady state in our Si nanobeam OMCs due to excessive amounts of heating from even modest parasitic optical absorption. However, the possibility is still left open for performing the kind of quantum engineering applications discussed in section 1.2 using sufficiently short drive pulses, provided the mechanical heating is slow enough. Indeed, certain applications (in particular, the heralded generation of mechanical Fock or entangled states) *require* the use of short pulses. In this chapter we will use pulsed measurement in conjunction with phonon counting techniques [56] to study the heating dynamics of a Si nanobeam OMC at millikelvin temperatures.

The measurement is shown schematically in Fig. 4.1a. Electro-optic modulation of the laser drive generates optical pulses with frequency and duty-cycle controlled by a variable delay electrical pulse generator. The red arrows indicate a red-detuned coherent drive ($\Delta = \omega_{\mathrm{m}}$). In this case absorption of a single phonon from the nanomechanical resonator results in upconversion of a drive photon to the anti-Stokes sideband at ω_{c}, represented by the black arrows. The cavity reflection is filtered to reject the drive frequency and subsequently directed to a high-efficiency single-photon detector (SPD) and a time-correlated single photon counting system synced to the pulse generator. This measurement repeats each pulse period, building up a histogram with respect to photon arrival time relative to a synchronization signal sent by the pulse generator. As the number of anti-Stokes photons is directly proportional to the phonon occupancy of the mechanical resonator $\langle n \rangle$, the photon count rate in each time bin then portrays the time-evolution of $\langle n \rangle$ during the pulse on-state, as described theoretically in section 1.1.5. For a blue-detuned drive ($\Delta = -\omega_{\mathrm{m}}$) the observed Stokes count rate is proportional to $\langle n \rangle + 1$. As described in Chapter 3 and illustrated in Fig. 4.1c, the observed

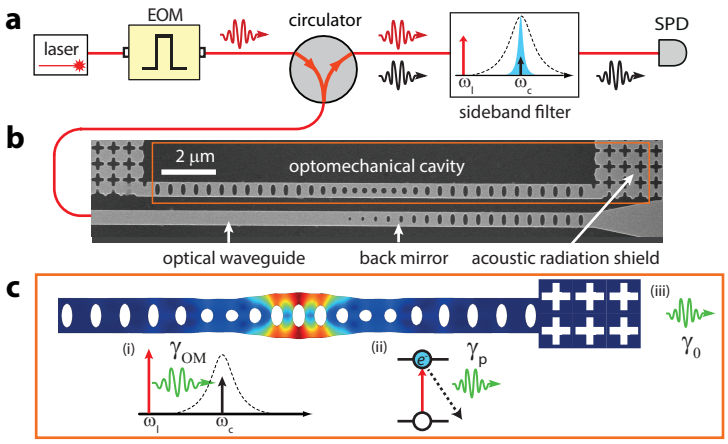

Figure 4.1: **Pulsed measurement schematic. a**, Schematic of the pulsed phonon counting measurement. Pump light at frequency ω_l (red arrows) is pulsed using an electro-optic modulator (EOM), and subsequently directed to the OMC inside a dilution fridge via an optical circulator. The cavity reflection is then filtered at the cavity frequency ω_c (black arrows) and directed to a single photon detector (SPD). **b**, SEM image of the nanobeam OMC cavity used for this measurement. **c**, The phonon occupancy of the OMC acoustic resonance is affected by three processes. Detuned drive light at ω_l (red) exchanges energy with the acoustic mode at rate γ_{OM}, generating scattered photons at ω_c (black) in the process. Additionally, the laser drive causes excitation of electronic defect states in the silicon device layer, subsequently exciting a hot phonon bath that heats the mechanical resonator at rate γ_p, and phonons can escape the cavity volume via the acoustic shield at intrinsic decay rate γ_0.

time-resolved heating curves will depend on a combination of the optomechanical damping/heating due to sideband scattering at rate γ_{OM}, local coupling, at rate γ_p, to hot phonons generated via optical absorption, and coupling to the external environment at a rate γ_0 determined by the acoustic radiation shield. In particular, these time-resolved pulsed measurements allow us to gauge how quickly the hot phonon bath "turns on", which directly impacts the feasibility of performing pulsed quantum optomechanical experiments in the ground state. All measurements presented herein were performed at a fridge base temperature of $T_f = 10$ mK.

4.1 Experimental Setup

The full measurement setup is shown in Fig. 4.2. A fiber-coupled, wavelength tunable external cavity diode laser is used as the light source, with a small portion ($\sim 1\%$) of the laser output sent to a wavemeter (λ-meter) for frequency stabilization. The remaining laser power is sent through both a C-band optical demultiplexer (DeMux) to reject broadband spontaneous emission from the laser, and a high-finesse tunable filter to remove laser phase noise at the mechanical frequency. As discussed in detail below in section 4.1.3, this avoids excess technical noise on the single photon detector [56], as well as issues of noise squashing/anti-squashing, as discussed in section 1.1.6.2.

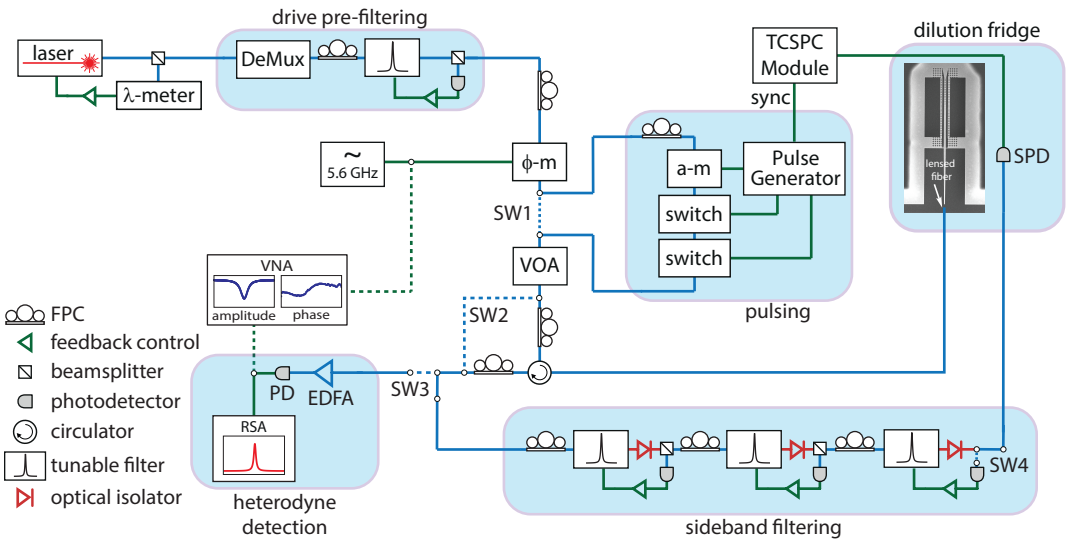

Figure 4.2: **Experimental setup.** Full setup for performing pulsed phonon counting. λ-meter: wavemeter, DeMux: C-band optical demultiplexer, ϕ-m: electro-optic phase modulator, SW: optical switch, a-m: high-extinction electro-optic amplitude modulator, switch: high-speed electro-optic switch, VOA: variable optical attenuator, EDFA: erbium-doped fiber amplifier, PD: high-speed photodetector, RSA: real-time spectrum analyzer, VNA: vector network analyzer, SPD: single photon detector, TCSPC: time-correlated single photon counting, FPC: fiber polarization controller.

After pre-filtering the laser drive is sent through an electro-optic phase modulator (ϕ-m), which may be driven by either an RF signal generator, to generate optical sidebands for locking the filter cavities, or a vector network analyzer (VNA) for obtaining the amplitude and phase response of the optical cavity.

To generate pulses, an optical switch (SW1) allows the drive tone to be sent through a series of electro-optic amplitude modulators for the generation of optical pulses. The first is a fast high-extinction modulator (a-m) with a rise and fall time of $\sim$ 25 ns. Though in principle the high-extinction modulator can provide $50-60$ dB extinction on its own, its transmission level is much less stable and difficult to lock at the maximum extinction point. Thus, we lock the modulator at $\sim$ 30 dB of extinction and use two electro-optic switches, each providing $\sim 18-20$ dB extinction, to fully extinguish the drive. While these switches are much slower ($\sim$ 200 ns rise time, $\sim$ 30 μs fall time), the extra switching time should have a negligible impact, as it is much smaller than the pulse periods used here ($\sim$ 1 ms or more). The modulators are collectively driven by a variable delay electrical pulse generator, which also provides a synchronization pulse to the single photon counting electronics.

A variable optical attenuator (VOA) controls the power input to the cavity, after which an optical circulator routes the laser light to a lensed fiber tip inside the dilution refrigerator for end-fire coupling to the device. Subsequently, the cavity reflection can be switched (SW3) into one of two

detection setups. The first setup sends the signal through an erbium-doped fiber amplifier (EDFA) followed by a high-speed photodetector (PD). The resulting amplified photocurrent may be directed to a real-time spectrum analyzer (RSA) in order to measure the optical noise power spectral density (NPSD) for mechanical characterization [1], or to the VNA, which is used in conjunction with the phase modulator to measure the full complex response of the optical cavity for purposes of optical characterization, as discussed in section 2.5.2. The second detection path sends the reflected signal through three additional tunable filters in order to reject the drive frequency, and then back into the dilution refrigerator where it is detected by a superconducting single photon detector (SPD). The output of this detector is sent to a time-correlated single photon counting (TCSPC) module to build up a histogram of photon count events as a function of time relative to the sync pulse received from the pulse generator. Further details of the single photon detectors used in this work are given in Appendix C.

The tunable filters used for both pre-filtering the drive and filtering the cavity reflection are commercially available, piezo-tunable Fabry-Perot filters (Micron Optics, FFP-TF2), all with a $\sim$ 50 MHz bandwidth, a free-spectral range of $\sim$ 20 GHz, and a tuning voltage of $\lesssim$ 18 V per free-spectral range. The filters each offer roughly 40 dB of drive suppression (relative to the peak transmission) measured at a filter-drive detuning of $\Delta_{\text{filter}} \approx \omega_{\text{m}}$, though the exact amount varies by $1-2$ dB from filter to filter. When locking the post-cavity filters a switch is used to bypass the cavity (SW2), as a relatively large amount of CW power is used during the locking procedure and we would like to avoid sending large amounts of power into the cavity unless necessary. More importantly, the SPD is also bypassed (SW4), as allowing too much power to reach the SPD will saturate it, resulting in significantly elevated dark counts ($\Gamma_{\text{dark}} \approx 500 - 1000$ c.p.s.) for $1-2$ minutes after the signal is turned off. Once the switches are set an RF signal generator is used to drive the phase modulator at ω_{m}, producing an optical sideband that is aligned with the cavity (and motional sideband) frequency. To lock the filter chain, a sinusoidal voltage (with an initial range of ± 10 V) is used to dither each filter while monitoring its transmission. The offsets of the sinusoids are then adjusted, and their amplitudes reduced, to maximize transmission of the desired sideband, fixing the voltages once all three filters are well aligned. Over time the filters will drift and the locking procedure will need to be repeated, though in subsequent re-locking attempts a much smaller dithering amplitude of ~ 1 V is sufficient. As the piezo elements controlling each filter have a finite settling time, the filters will drift much more rapidly following the initial locking attempt (during which large voltage shifts are applied) than following subsequent re-locks. After several re-locks, the filters will typically become stable enough that the total transmission at the sideband frequency changes by $\lesssim 5\%$ over several

[1] Here we use direct detection, where the classical cavity reflection is used as the local oscillator, rather than the more complex heterodyning scheme utilized in Chapter 3. While this results in a much lower signal-to-noise ratio (SNR), as discussed in section 3.1, it is technically much simpler. As heterodyne spectroscopy is used here only for the initial characterization of the device parameters at $T_{\text{f}} = 4$ K, the reduced SNR is of no real consequence.

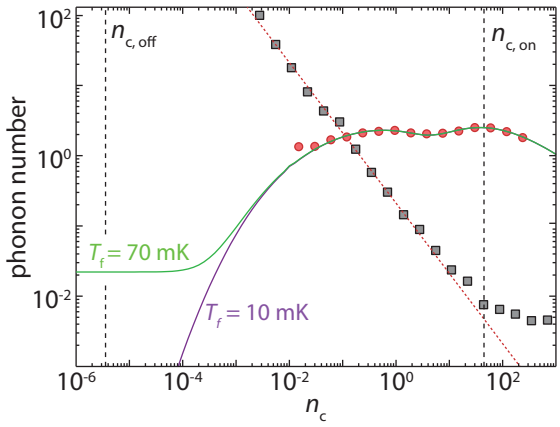

Figure 4.3: **Pulsed measurement sensitivity.** Noise-equivalent phonon number n_{NEP} (gray squares) and phonon occupancy $\langle n \rangle$ (red circles) for red-detuned ($\Delta = \omega_{\mathrm{m}}$) continuous-wave (CW) driving versus intracavity photon number n_{c}. The red dashed line indicates the expected n_{NEP} contribution from SPD dark counts. The vertical dashed line at $n_{\mathrm{c}} \approx 45$ indicates the photon number during the on-state of the pulse ($n_{\mathrm{c,on}}$), while the dashed line at $n_{\mathrm{c}} \approx 4 \times 10^6$ indicates the levels reached during the off-state ($n_{\mathrm{c,off}}$). Solid green and purple lines show fits to the CW heating model for fridge bath temperatures of $T_{\mathrm{f}} = 70$ mK and 10 mK, respectively, extrapolated to the range of $n_{\mathrm{c,off}}$.

minutes. At this point the phase modulator is turned off, the pulse generator, cavity, and SPD are switched back into the optical train, and we begin accumulating pulsed data for 2 minutes before re-locking the filters. The total filter transmission is recorded at the end of a locking procedure and again prior to re-locking, and if the transmission has shifted by more than a few percent the previous pulsed dataset is discarded. All the time-resolved data shown in this chapter are taken in this manner, with all 2 minute datasets averaged together to produce the final pulsed histogram.

4.1.1 Device Characterization

The basic device parameters are characterized at a fridge base temperature of $T_{\mathrm{f}} = 4$ K in the same manner described in sections 2.5.2 and 3.3. For the device measured in this chapter the relevant parameters are $\eta_{\mathrm{cpl}} = 68\%$, $\kappa/2\pi = 443$ MHz, $\kappa_{\mathrm{e}}/\kappa = 0.5$, $g_0/2\pi = 710$ kHz, and $\omega_{\mathrm{m}}/2\pi = 5.6$ GHz.

The measurements of this chapter utilize sideband asymmetry to determine $\langle n \rangle$ (as discussed in section 1.1.5), and as such do not strictly require us to calibrate the total detection efficiency η, as it is common to both red and blue detunings. However, we may calibrate the total detection efficiency in a similar fashion to Chapter 3, yielding $\eta_{\mathrm{2det}} \approx 2\%$ and $\eta_{\mathrm{SPD}} \approx 68\%$, for a total detection efficiency of $\eta = \eta_{\mathrm{cpl}}\eta_{\mathrm{2det}}\eta_{\mathrm{SPD}} \approx 1\%$, consistent with the observed count rates and occupancy $\langle n \rangle$ estimated via sideband asymmetry.

4.1.2 Phonon Counting Sensitivity

As discussed in section 1.1.5, the signal-to-noise ratio (SNR) of the phonon counting measurement is determined by the sideband scattering rate γ_{OM}, the total system detection efficiency η, the dark count rate of the SPD Γ_{dark}, and the residual transmission A, relative to the peak transmission, of the filters at the drive frequency. A useful parameterization of the sensitivity to low $\langle n \rangle$ is the noise-equivalent phonon occupancy n_{NEP}, given by

$$n_{\mathrm{NEP}} = \frac{\kappa \Gamma_{\mathrm{dark}}}{\eta \kappa_{\mathrm{e}} \gamma_{\mathrm{OM}}} + A \left(\frac{\kappa \omega_{\mathrm{m}}}{2 \kappa_{\mathrm{e}} g_0} \right)^2 . \tag{4.1}$$

A typical n_{NEP} curve for the setup illustrated in Fig. 4.2, measured using a comparable device at room temperature, is shown in Fig. 4.3 as gray squares. To compute n_{NEP}, the noise count rate $\Gamma_{\mathrm{dark}} + \Gamma_{\mathrm{drive}}$ is measured with the drive laser detuned far off-resonance from the optical cavity ($\Delta \gg \omega_{\mathrm{m}}$) while the drive-filter detuning is kept constant. While this eliminates the signal due to motional sideband photons, because of the flat cavity reflection for $\Delta \geq \omega_{\mathrm{m}}$ in the sideband resolved regime (see Fig. 2.10b) it does not change the reflected drive amplitude or the corresponding noise count rate Γ_{drive}. To determine the per-phonon sideband count rate $\Gamma_{\mathrm{SB,0}}$, the sideband count rate (i.e., total count rate minus the measured noise count rate) is measured at low n_{c} where back-action is negligible and $\langle n \rangle$ is equal to the known room temperature thermal mode occupancy of $n_{\mathrm{b}} \approx 1100$. Since $\Gamma_{\mathrm{SB,0}}$ scales linearly with n_{c}, we can then determine $\Gamma_{\mathrm{SB,0}}$ for all n_{c} from this single measurement without relying on calibration of the optomechanical backaction. We may then compute n_{NEP} by dividing the measured noise count rate by the extrapolated $\Gamma_{\mathrm{SB,0}}$ at each value of n_{c}. The measured sensitivity follows the expected curve at low power due to detector dark counts (dashed red curve), and at the highest n_{c} begins to saturate to a limiting value of $n_{\mathrm{NEP}} \approx 4 \times 10^{-3}$, corresponding in this case to ~ 102 dB of drive suppression.

For sufficiently high probe power ($n_{\mathrm{c}} > 40$), n_{NEP} falls below 0.01, enabling sensitive detection of the mechanical resonator deep in its quantum ground state. However, as we know, at sub-kelvin temperatures optical absorption heating produces a steady-state $\langle n \rangle > 1$ for $n_{\mathrm{c}} > 0.01$ (red circles in Fig. 4.3) during continuous-wave (CW) detection. In order to maintain the OMC in the mechanical ground state, the duty cycle of the pulse train must be kept sufficiently low, and the modulation depth sufficiently high, such that the mechanical mode thermalizes to the ambient bath temperature provided by the dilution fridge (with occupancy n_0) between successive pulses.

As in Chapter 3, the CW behavior fits well to the phenomenological heating model detailed in section 3.5.1. When extrapolated to low n_{c}, fits of the CW $\langle n \rangle$ assuming n_0 corresponding to fridge base temperatures $T_{\mathrm{f}} = 10$ mK (purple solid line) and 70 mK (green solid line) reveal a relevant pulse-off-state regime of $n_{\mathrm{c}} < 10^{-4}$ in which absorption heating effects are negligible. We are able to reliably achieve pulse extinction ratios in the range $65 - 75$ dB, allowing for high-sensitivity

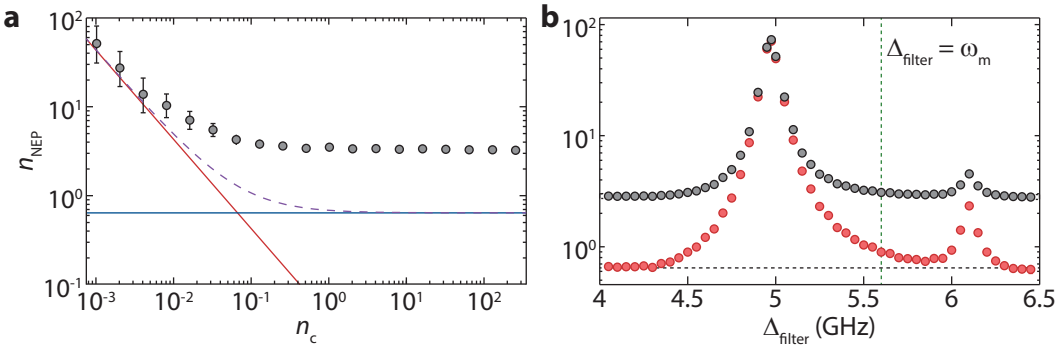

Figure 4.4: **Technical laser noise contributions to n_{NEP}. a**, Noise equivalent phonon number n_{NEP} versus intracavity photon number n_{c} calculated using the measured signal and noise count rates for a two-filter setup (gray circles). Solid lines indicate the theoretically expected contributions due to dark counts (red) and drive bleed-through (blue), based on the measured system efficiency and drive suppression, with the sum of the two contributions displayed as a purple dashed line. **b**, n_{NEP} versus filter-drive detuning Δ_{filter} for $n_{\mathrm{c}} \approx 65$, with (red) and without (gray) an additional C-band band-pass filter inserted. The vertical green line indicates the detuning corresponding to the data from **a**, and the horizontal black line indicates the expected limiting sensitivity.

detection during the on-state ($n_{\mathrm{c,on}} = 45$, corresponding to an optomechanical damping rate of $\gamma_{\mathrm{OM}}/2\pi = 200$ kHz) and fridge thermalization during the off-state ($n_{\mathrm{c,off}} < 1.2 \times 10^{-5}$).

4.1.3 Effects of Technical Laser Noise

For this experiment we use a higher frequency nanobeam OMC design than that studied in Chapter 3, with a mechanical frequency of $\omega_{\mathrm{m}}/2\pi = 5.6$ GHz. While the higher frequency is desirable for achieving lower phonon occupancies at a given temperature, it also raises concerns about the effects of laser phase noise on the measurements, as the laser used in this experiment has previously been observed to possess a prominent phase noise peak at 5 GHz [23]. In addition to phase noise, most diode lasers typically have a small amount of broadband spontaneous emission. While this additional noise is orders of magnitude weaker than the laser tone itself, it exists outside the wavelength region ($\lambda \sim 1520-1570$ nm) where the filters are guaranteed to be high finesse, and thus can be transmitted with high efficiency to the SPDs.

These noise contributions to the phonon counting noise floor were previously examined using a similar OMC device in an earlier version of the setup incorporating only two tunable filters [56]. In Fig. 4.4a we show the measured n_{NEP} of the two-filter setup (solid gray circles), as well as the expected theoretical curve given by Eq. 1.41 (purple dashed line) and the individual contributions of dark counts (red solid line) and drive bleed-through (blue solid line). At high n_{c} we observe n_{NEP} saturate to a value several times larger than expected given the ~ 80 dB of drive attenuation provided by the two filters. To characterize this excess noise we measure n_{NEP} as a function of filter-drive detuning Δ_{filter} at a high power, where the drive bleed-through dominates the noise

($n_c \approx 65$), as shown in Fig. 4.4b. A strong dependence on $\Delta_{\rm filter}$ is observed, with a peak in the noise at 5 GHz and a secondary peak at 6.1 GHz, consistent with measured phase-noise of the laser [23]. With the addition of a C-band bandpass filter prior to the SPD to remove broadband spontaneous emission from the drive laser, and at frequencies far from the laser phase-noise peaks, the measured $n_{\rm NEP}$ agrees well with the theoretical predictions based on the expected drive suppression (horizontal dashed curve). Recall that, as the measurement of noise counts is taken when $\Delta \gg \omega_{\rm m}$, the contribution of the phase noise to $n_{\rm NEP}$ will be given by Eq. 1.58 rather than Eq. 1.56. Using the maximum value of $n_{\rm NEP} \approx 70$ observed in Fig. 4.4b, which is due almost entirely to phase noise, we extract a peak phase noise spectral density of $S_{\phi\phi}(\omega) \approx 1.3 \times 10^{-14}$ Hz^{-1} at $\omega \approx 5$ GHz, consistent with the value previously measured using Mach-Zender interferometry in Ref. [23].

To get rid of this excess noise, we insert both the bandpass filter (for filtering spontaneous emission) and a tunable high-finesse filter (for filtering phase noise) immediately after the drive laser output as shown in Fig. 4.2, enabling us to reach $n_{\rm NEP} \ll 1$ using a three-filter phonon counting setup. As discussed in section 4.1.2 the effective drive attenuation is only ~ 102 dB rather than the expected value of ~ 120 dB, likely due to a small amount of phase noise not removed by pre-filtering the drive laser [2]. A conservative estimate of the residual phase noise can be made by assuming that the limiting value of $n_{\rm NEP} \approx 4 \times 10^{-3}$ observed in Fig. 4.3 is entirely due to phase noise (i.e., perfect filtering of the drive tone). Using Eq. 1.58 we find $S_{\phi\phi}(\omega_{\rm m}) \lesssim 8 \times 10^{-19}$ Hz^{-1}. For the drive power during the on-state of the pulse ($n_{\rm c,on} \approx 45$), the corresponding effective phase noise occupancy is $n_\phi \approx 3.2 \times 10^{-5}$, which has a negligible effect on the measurements here.

4.2 Experimental Results

The sideband photon count rate versus time is shown in Fig. 4.5a, for both red- ($\Delta = \omega_{\rm m}$) and blue-detuned ($\Delta = -\omega_{\rm m}$) driving. Henceforth, the variable t refers to time relative to the synchronization signal generated by the pulse generator, while $t_{\rm pulse}$ refers to time relative to the start of the optical pulse occurring around $t \approx 1$ μs. For the measurements shown here the pulse period is $T_{\rm per} = 5$ ms. Vertical dashed lines indicate the time bins corresponding to the start and stop of the pulse, determined from observing the rising and falling edges of the pulse when bypassing the cavity.

Throughout the pulse a pronounced asymmetry is observed between count rates for red-detuned versus blue-detuned driving, which can be quantified by the asymmetry parameter $\xi = \Gamma_-/\Gamma_+ - 1$, where $\Gamma_\pm$ is the sideband photon count rate for a drive detuning $\Delta = \pm\omega_{\rm m}$. This asymmetry, shown versus $t_{\rm pulse}$ in Fig. 4.5b, initially decreases with time before leveling off and beginning to increase for sufficiently long pulse times. The increase at later times can be ascribed to the effect of

[2] It is also possible that some of the residual noise is spontaneous emission. However, this is of much less concern as, unlike phase noise, it is far off-resonant with the optical cavity and cannot contribute to real heating of the mechanical mode or noise squashing/anti-squashing effects.

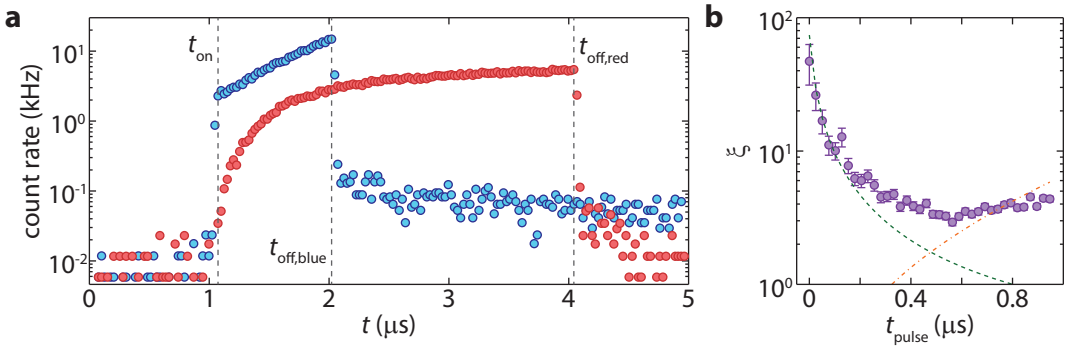

Figure 4.5: **Time-resolved heating and asymmetry.** **a**, Total photon count rate versus time t for a red-detuned (red circles, $\Delta = \omega_{\rm m}$) and blue-detuned (blue circles, $\Delta = -\omega_{\rm m}$). For both data sets $n_{\rm c,on} \approx 45$ and $T_{\rm per} = 5$ ms. Vertical dashed lines indicate the start and stop times of the pulses. **b**, Sideband asymmetry ξ versus time within the pulse $t_{\rm pulse}$. The green dashed and orange dash-dotted lines show the theoretically expected $\xi(t)$ in the absence of backaction and heating effects, respectively.

optomechanical backaction, which results in cooling or heating of the mechanical resonator for red- or blue-detuned drives, respectively. However, for $t_{\rm pulse} \ll \gamma_{\rm OM}^{-1} \approx 750$ ns the effects of backaction can be neglected and the phonon occupancy $\langle n \rangle$ may be assumed equal for both drive detunings. In this case, the asymmetry is simply related to the occupancy by $\xi = \langle n \rangle^{-1}$, and arises from the fundamental asymmetry between phonon absorption ($\Gamma_+ \propto \langle n \rangle$) and emission ($\Gamma_- \propto \langle n \rangle + 1$) processes [13, 167]. Theoretical plots of the two contributions to $\xi(t)$ are shown in Fig. 4.5b. The green dashed line shows the expected $\xi(t)$ assuming that $\langle n \rangle|_{\Delta=\omega_{\rm m}} = \langle n \rangle|_{\Delta=-\omega_{\rm m}}$ (no backaction), while the orange dash-dotted line shows $\xi(t)$ in the absence of optical heating and in the case when $n_{\rm b} \gg 1$ such that $\xi(0) = 0$ and asymmetry arises solely from backaction effects. From the asymmetry measured in the first 25 ns time bin, we extract a minimum phonon occupancy of $\langle n \rangle_{\rm min} = 0.021 \pm 0.007$ ($T_{\rm min} \approx 70$ mK). This measured occupancy is lower than previous results in both cavity optomechanical [45] and electromechanical [46] systems by more than an order of magnitude, and is comparable to the occupancy measured in other GHz frequency mechanical oscillators at mK temperatures [37]. Moreover, while motional sideband asymmetry has previously been measured in a variety of optomechanical and electromechanical systems using linear detection schemes [13–16] , the use of phonon counting techniques allows the observed asymmetry to be directly and unambiguously attributed to quantum fluctuations of the mechanical oscillator [14, 168].

During the off-state of the pulse, optical heating of the mechanical resonator should be negligible and the phonon occupancy should cool at the intrinsic damping rate γ_0. Thus, the initial and final occupancies during the pulse ($\langle n \rangle_{\rm i}$ and $\langle n \rangle_{\rm f}$, respectively) should obey the relation $\langle n \rangle_{\rm i} = e^{-\gamma_0 T_{\rm per}} \langle n \rangle_{\rm f}$, assuming the pulse period $T_{\rm per}$ is much larger than the pulse width $T_{\rm pulse}$. The ratio $\langle n \rangle_{\rm i} / \langle n \rangle_{\rm f}$, shown in Fig. 4.6a versus $T_{\rm per}$, displays the expected exponential decay with $\gamma_0/2\pi =$

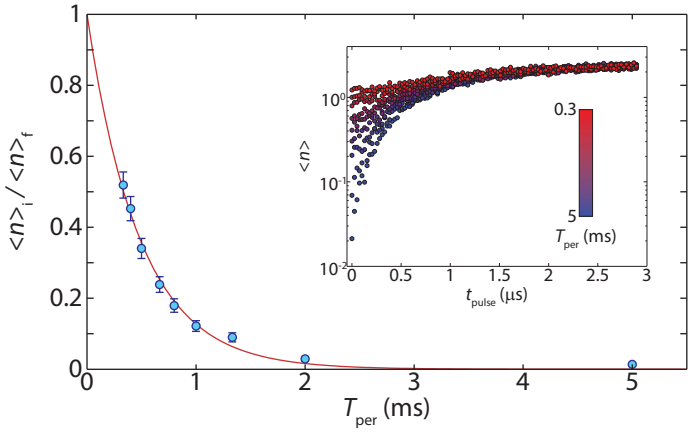

Figure 4.6: **Ringdown measurement of γ_0.** Ratio of initial/final phonon number $\langle n\rangle_{\rm i} / \langle n\rangle_{\rm f}$ versus pulse period $T_{\rm per}$. The solid red curve shows a fit to exponential decay. The inset shows $\langle n\rangle$ versus $t_{\rm pulse}$ for the different values of $T_{\rm per}$.

328±14 Hz and a corresponding intrinsic mechanical quality factor of $Q_{\rm m} \approx 1.7{\times}10^7$. This decay rate corresponds well with the value inferred from fitting the proposed heating model to CW occupancy data, as in Chapter 3. From our measurement of γ_0 and $\langle n\rangle_{\rm min}$ we can estimate a thermal coherence time of $\tau_{\rm th} = (\gamma_0(1+\langle n\rangle_{\rm min}))^{-1} = 475 \pm 21$ μs. This value of $Q_{\rm m}$ also shows excellent agreement with that expected from simulation for 4 cross periods and roughly 2% disorder in the hole sizes and periods (see section 1.3.5), suggesting that we can reach the regime where the mechanical damping is limited by the radiation shield.

Once $\langle n\rangle$ is determined in the initial time bin it may be used to convert measured count rates to phonon occupancies throughout the pulse for either drive detuning. The resulting calibrated occupancy versus $t_{\rm pulse}$ is shown for $\Delta = \omega_{\rm m}$ in Fig. 4.7a, and for $\Delta = -\omega_{\rm m}$ in Fig. 4.7b. The simplest model of the heating dynamics assumes the creation of the hot phonon bath to be effectively instantaneous. In this case $\langle n\rangle$ will obey the following rate equation during the pulse on-state:

$$\dot{\langle n\rangle}|_{\Delta=\pm\omega_{\rm m}} = -\gamma\,\langle n\rangle + \gamma_{\rm p} n_{\rm p} + \gamma_0 n_0 + \frac{1}{2}\,(1\mp 1)\,\gamma_{\rm OM}, \tag{4.2}$$

where $\gamma = \gamma_0 + \gamma_{\rm p} \pm \gamma_{\rm OM}$, and the extra factor of $\gamma_{\rm OM}$ for a blue-detuned drive accounts for the possibility of spontaneous emission into the mechanical resonator due to the optomechanical interaction. This has the simple solution

$$\langle n\rangle\,(t)|_{\Delta=\pm\omega_{\rm m}} = \langle n\rangle\,(t_0) e^{-\gamma t} + n_{\rm H}\left(1 - e^{-\gamma t}\right), \tag{4.3}$$

$$n_{\rm H} = \gamma^{-1}\left(\gamma_{\rm p} n_{\rm p} + \gamma_0 n_0 + \frac{1}{2}\,(1\mp 1)\,\gamma_{\rm OM}\right), \tag{4.4}$$

where t_0 is the start time of the pulse, and $t_0 \le t \le t_0 + T_{\rm pulse}$.

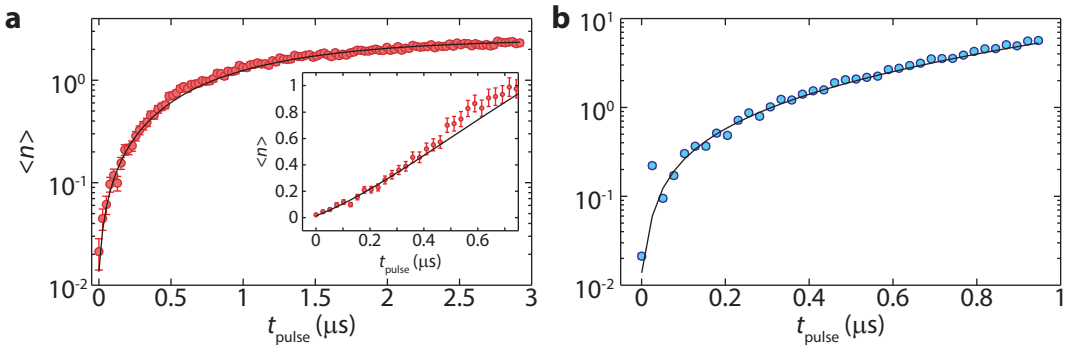

Figure 4.7: **Calibrated occupancy.** Calibrated phonon occupancy $\langle n \rangle$ versus t_{pulse} for **a**, $\Delta = \omega_{\mathrm{m}}$ and **b**, $\Delta = -\omega_{\mathrm{m}}$. The solid black lines show fits to a model including a slow exponential turn-on of the hot phonon bath. The inset of **a** shows detail for $t_{\mathrm{pulse}} < 750$ ns on a linear scale.

In principle, γ_{OM} can be determined independently by calibrating the device parameters, γ_0 is found by fitting the occupancy decay during the pulse off-state (see Fig. 4.6 and associated discussion above), and γ_{p} and n_{p} can be subsequently determined by fitting the CW occupancy curve shown in Fig. 4.3. However, using these independently determined values to fit Eq. 4.4 to the red- and blue-detuned data shown in Fig. 4.7a and b results in a poor fit and inconsistent results. In particular, the apparent exponential heating rate γ is much smaller than expected for the red-detuned pulse and larger than expected for the blue-detuned pulse. Moreover, as can be seen in the inset to Fig. 4.7a, the heating rate abruptly changes around $t_{\mathrm{pulse}} \approx 600$ ns. These inconsistencies indicate additional complexity in the heating dynamics, likely due to a finite equilibration time of the hot phonon bath.

Though it does not perfectly model the abruptness of the change in heating rate, the simplest addition to the heating model which yields good agreement with the experimental data is to approximate the finite hot phonon equilibration time by allowing a fraction of the bath occupancy n_{p} to turn on exponentially over time. Thus, the phenomenological rate equation becomes

$$\dot{\langle n \rangle}|_{\Delta=\pm\omega_{\mathrm{m}}} = -\gamma \langle n \rangle + \gamma_{\mathrm{p}} n_{\mathrm{p}} \left(1 - \delta_{\mathrm{b}} e^{-\gamma_{\mathrm{S}} t}\right) + \gamma_0 n_0 + \frac{1}{2}\left(1 \mp 1\right)\gamma_{\mathrm{OM}}, \tag{4.5}$$

where δ_{b} is the slow growing fraction of n_{p} and γ_{S} is the turn-on rate. Strictly speaking γ_{p} should depend on the phonon distribution of the hot phonon bath, and thus would be expected to be time-dependent in this model as well. However, as the resulting rate equation becomes intractable, and the effects should be negligible in the regime $\gamma_{\mathrm{OM}} \gg \gamma_{\mathrm{p}}$, we approximate γ_{p} equal to its steady-state value. This modified rate equation has the solution

$$\langle n \rangle (t)|_{\Delta=\pm\omega_{\mathrm{m}}} = \langle n \rangle (t_0) e^{-\gamma t} + n_{\mathrm{H}} \left(1 - e^{-\gamma t}\right) + n_\delta \left(e^{-\gamma_{\mathrm{S}} t} - e^{-\gamma t}\right), \tag{4.6}$$

$$n_\delta = \frac{\gamma_{\mathrm{p}} n_{\mathrm{p}} \delta_{\mathrm{b}}}{\gamma_{\mathrm{S}} - \gamma}, \tag{4.7}$$

which is used to obtain the fits shown in Fig. 4.7a and b, using $\langle n \rangle (t_0)$, γ_S, and δ_b as free parameters, yielding $\delta_b = 0.79 \pm 0.08$, $\gamma_S/2\pi = 215 \pm 29$ kHz, and $\langle n \rangle (t_0)$ in good agreement with the value $\langle n \rangle_{min}$ obtained via asymmetry measurements.

During the off-state of the pulse ($t_0+T_{pulse} \leq t \leq t_0+T_{per}$), the resonator will simply cool towards the ambient fridge occupancy n_0 at the intrinsic damping rate γ_0. Using the initial condition that $\langle n \rangle (t_0) = n_0$ for the first pulse ($t_0 = 0$), and iterating many pulses, we find that in the steady-state the initial phonon occupancy during a pulse (assuming $T_{per} \gg T_{pulse}$) is

$$\langle n \rangle (t_0) = \frac{n_0 \left(1 - e^{-\gamma_0 T_{per}}\right) + n_H \left(1 - e^{-\gamma T_{pulse}}\right) e^{-\gamma_0 T_{per}} + n_\delta \left(e^{-\gamma_S T_{pulse}} - e^{-\gamma_S T_{per}}\right) e^{-\gamma_0 T_{per}}}{1 - e^{-\gamma T_{pulse} - \gamma_0 T_{per}}}. \tag{4.8}$$

Thus, once n_0, γ_S and δ_b are determined by fitting the occupancy curves, we may use Eqs. 4.7 and 4.8 to determine the occupancy throughout the pulse for arbitrary T_{per} and T_{pulse}. This allows us to determine the maximum attained phonon occupancy as a function of pulse parameters in order to assess the feasibility of various quantum protocols that require ground state operation, as discussed further in Chapter 5.

Chapter 5

Future Work and Outlook

At this point we can envision two main routes to progress beyond the initial cryogenic measurements performed in this work. One is to attempt to perform initial wavelength conversion and state heralding experiments using suitably short pulses in the current nanobeam structures. Using our measurements of the steady-state and time-resolved heating dynamics, we are now in a position to quantitatively assess the feasibility of such experiments.

The second option is to attempt to mitigate the effects of heating in the steady-state by moving to an alternative structure. For this task we may turn to the snowflake OMCs discussed briefly in section 1.3.6, as the two-dimensional nature of such cavities should offer a thermal conductivity that is superior to that of the one-dimensional nanobeams.

5.1 Prospects for 1D Nanobeams

5.1.1 Wavelength Conversion

While the low initial thermal occupancies and long mechanical coherence times measured in this work are promising, the utility of cavity optomechanical systems for performing coherent quantum operations between the optical and mechanical degrees of freedom is ultimately predicated upon the ability to simultaneously achieve $\langle n \rangle \ll 1$ and large cooperativity $C = \gamma_{\mathrm{OM}}/\gamma_{\mathrm{i}}$, where $\gamma_{\mathrm{i}} = \gamma_0 + \gamma_{\mathrm{p}}$ is the total coupling rate between the mechanical resonator and its thermal bath. In the specific example of optomechanically mediated coherent transfer of photons between optical and superconducting microwave resonators discussed in section 1.2.1, the relevant figure of merit is the effective cooperativity $C_{\mathrm{eff}} = C/n_{\mathrm{b}}$, where n_{b} is the effective bath occupancy defined such that $\gamma_{\mathrm{i}} n_{\mathrm{b}} = \gamma_0 n_0 + \gamma_{\mathrm{p}} n_{\mathrm{p}}$. A necessary condition for achieving low noise conversion at the single photon level is $C_{\mathrm{eff}} \gg 1$. Using the values of γ_0 and γ_{S} measured in Chapter 4, we may use Eqs. 4.7 and 4.8 to calculate the maximum phonon occupancy achieved during the pulse, $\langle n \rangle_{\mathrm{max}}$, for a given T_{pulse} and T_{per} and $\Delta = \omega_{\mathrm{m}}$ (Fig. 5.1a), as well as C_{eff} as a function of t_{pulse} (Fig. 5.1b).

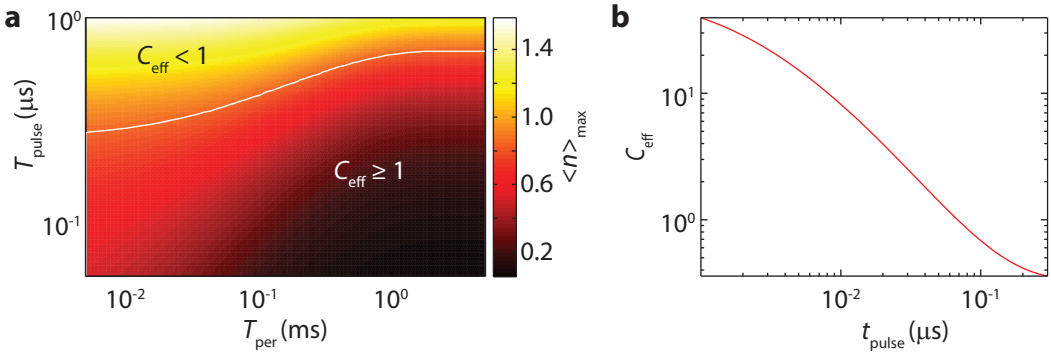

Figure 5.1: **Pulsed $C_{\mathbf{eff}}$. a**, Maximum phonon occupancy $\langle n \rangle_{\max}$ versus pulse period T_{per} and pulse width T_{pulse}. The white contour delineates the region where the effective cooperativity $C_{\mathrm{eff}} \geq 1$ throughout the pulse. **b**, C_{eff} versus time within the pulse t_{pulse} for $T_{\mathrm{per}} = 5$ ms, $T_{\mathrm{pulse}} = 3$ μs.

Though optical heating of our devices precludes the use of a CW drive, due to the slow turn-on of the hot phonon bath observed in this work, we find that $\langle n \rangle < 1$ and $C_{\mathrm{eff}} > 1$ can be maintained throughout the entire pulse period for $T_{\mathrm{pulse}} \lesssim 300$ ns at pulse rates approaching 1 MHz. Furthermore, we find $C_{\mathrm{eff}} \gg 1$ during the initial 100 ns of the pulse and reaches values as large as $C_{\mathrm{eff}} \approx 40$. This demonstrates that coherent optical to microwave conversion should in principle be possible at the single photon level using current systems, provided we are willing to operate in the pulsed regime.

Of course, in this discussion we have neglected any mention of the coupling of superconducting microwave resonators to the high-frequency modes of a silicon nanobeam, which is another necessary component for optical-to-microwave state conversion. Although a detailed discussion of such matters is beyond the scope of this thesis, it has been shown theoretically that it is possible to fabricate electromechanical systems which are compatible with the type of nanobeam OMCs studied in this work, with sufficiently strong coupling between the microwave resonator and the localized "breathing" mechanical mode to enable the type of state transfer protocols considered here [169]. Efforts to experimentally realize such devices are currently ongoing in the Painter group [170, 171].

5.1.2 Heralded Fock State Generation

We may also evaluate the fidelity of the creation of a single phonon Fock state [34] in our system as a function of pulse width and period. To properly account for the effects of heating during the pulse, we must extend the simple picture described in section 1.2.2 in order to include to lowest order the effects of interaction with mechanical bath.

The fidelity of heralded phonon addition/subtraction, as discussed in section 1.2.2, is in principle degraded by a number of non-ideal factors, including counter-rotating terms in the interaction Hamiltonian, a small probability of multi-phonon events, the initial thermal occupancy of the mode, and

finite amounts of heating or mechanical bath interaction during the pulse. In the specific optomechanical systems considered in this work, the predominant source of non-ideal behavior is optically induced heating during the pulse which, as discussed in Chapters 3 and 4, can be substantial at the cryogenic temperatures where we wish to operate. Consequently, in order to assess the feasibility of heralded state preparation in our devices, it is necessary to develop a theoretical expression for the fidelity of the generated state, including the effects of mechanical heating to lowest order.

For concreteness we will consider the case $\Delta = -\omega_{\mathrm{m}}$ (blue-detuned drive). We start from the linearized Heisenberg-Langevin equations given in Eq. 1.12 and 1.13. Moving into a frame rotating at the mechanical frequency (i.e., $\hat{a} \to \hat{a}e^{i\omega_{\mathrm{m}}t}$, and so on for all operators), and making the rotating wave approximation, valid in the weak coupling ($g_0\sqrt{n_{\mathrm{c}}} \ll \kappa$) and sideband-resolved ($\kappa/\omega_{\mathrm{m}} \ll 1$) limit, we obtain

$$\dot{\hat{a}} = -\frac{\kappa}{2}\hat{a} + iG\hat{b}^\dagger + \sqrt{\kappa}\hat{a}_{\mathrm{in}}, \tag{5.1}$$

$$\dot{\hat{b}} = -\frac{\gamma}{2}\hat{b} + iG\hat{a}_{\mathrm{in}}^\dagger + \sqrt{\gamma}\hat{b}_{\mathrm{in}}, \tag{5.2}$$

where $G = g_0\sqrt{n_{\mathrm{c}}}$ and $\hat{a}_{\mathrm{in}}$, $\hat{b}_{\mathrm{in}}$ are the usual noise operators multiplied by $e^{-i\omega_{\mathrm{m}}t}$.

Since we are working in the weak-coupling limit ($G \ll \kappa$) we may use the adiabatic solution for $\hat{a}$ (i.e., $\dot{\hat{a}} \approx 0$). Moreover, we wish to include the effects of mechanical noise to lowest order. Considering the effect of a short optical pulse of duration τ, we consider the case $\gamma\tau, \gamma\int_0^\tau ds\; n_{\mathrm{b}}(s) \ll 1$, as well as $\gamma_{\mathrm{OM}}\tau \ll 1$ and $\gamma_{\mathrm{OM}} \gg \gamma$. Note that we have taken the mechanical bath occupancy n_{b} to be explicitly time dependent. As we will see in Chapter 4, the optically induced heating is best modeled by assuming a phenomenological time-dependence to the bath. Under these assumptions, as in Ref. [172], the noise term will be retained in the mechanical equation of motion only, since in the optical equations of motion it will acquire an extra factor of $\gamma_{\mathrm{OM}}\tau$ and thus can be neglected to lowest order. Furthermore, the $(\gamma/2)\,\hat{b}$ term will be neglected in favor of the much larger $(\gamma_{\mathrm{OM}}/2)\,\hat{b}$ term. With these approximations, we arrive at the approximate equations:

$$\hat{a} \approx i\sqrt{\frac{\gamma_{\mathrm{OM}}}{\kappa}}\hat{b}^\dagger + \frac{2}{\kappa}\hat{a}_{\mathrm{in}}, \tag{5.3}$$

$$\dot{\hat{b}} \approx \frac{\gamma_{\mathrm{OM}}}{2}\hat{b} + i\sqrt{\gamma_{\mathrm{OM}}}\hat{a}_{\mathrm{in}}^\dagger + \sqrt{\gamma}\hat{b}_{\mathrm{in}}. \tag{5.4}$$

Formally integrating the equation of motion for $\hat{b}$ over the duration of the pulse while using the optical input-output relation Eq. 1.6, we arrive at the approximate equations:

$$\hat{a}_{\mathrm{out}}(t) = \hat{a}_{\mathrm{in}}(t) + i\sqrt{\gamma_{\mathrm{OM}}}e^{\frac{\gamma_{\mathrm{OM}}t}{2}}\hat{b}^\dagger(0) + \gamma_{\mathrm{OM}}e^{\frac{\gamma_{\mathrm{OM}}t}{2}}\int_0^t ds\; e^{\frac{-\gamma_{\mathrm{OM}}s}{2}}\hat{a}_{\mathrm{in}}(s) \tag{5.5}$$

$$\hat{b}(t) = e^{\frac{\gamma_{\mathrm{OM}}t}{2}}\hat{b}(0) + i\sqrt{\gamma_{\mathrm{OM}}}e^{\frac{\gamma_{\mathrm{OM}}t}{2}}\int_0^t ds\; e^{\frac{-\gamma_{\mathrm{OM}}s}{2}}\hat{a}_{\mathrm{in}}^\dagger(s) + \sqrt{\gamma}e^{\frac{\gamma_{\mathrm{OM}}t}{2}}\int_0^t ds\; e^{\frac{-\gamma_{\mathrm{OM}}s}{2}}\hat{b}_{\mathrm{in}}(s). \tag{5.6}$$

We now introduce the following temporal modes [172]:

$$\hat{B}_{\rm in} = \hat{b}(0), \quad \hat{B}_{\rm out} = \hat{b}(\tau), \tag{5.7}$$

$$\hat{A}_{\rm in} = \sqrt{\frac{\gamma_{\rm OM}}{1-e^{-\gamma_{\rm OM}\tau}}} \int_0^\tau ds\, e^{\frac{-\gamma_{\rm OM} s}{2}} \hat{a}_{\rm in}(s), \tag{5.8}$$

$$\hat{A}_{\rm out} = \sqrt{\frac{\gamma_{\rm OM}}{e^{\gamma_{\rm OM}\tau}-1}} \int_0^\tau ds\, e^{\frac{\gamma_{\rm OM} s}{2}} \hat{a}_{\rm out}(s) \tag{5.9}$$

$$\hat{F} = \sqrt{\gamma} e^{\frac{\gamma_{\rm OM}\tau}{2}} \int_0^\tau ds\, e^{\frac{-\gamma_{\rm OM} s}{2}} \hat{b}_{\rm in}(s), \tag{5.10}$$

which allows us to rewrite the input/output equations at the end of the pulse ($t = \tau$):

$$\hat{A}_{\rm out} = e^{\frac{\gamma_{\rm OM}\tau}{2}} \hat{A}_{\rm in} + i\sqrt{e^{\gamma_{\rm OM}\tau}-1}\hat{B}^\dagger_{\rm in}, \tag{5.11}$$

$$\hat{B}_{\rm out} = e^{\frac{\gamma_{\rm OM}\tau}{2}} \hat{B}_{\rm in} + i\sqrt{e^{\gamma_{\rm OM}\tau}-1}\hat{A}^\dagger_{\rm in} + \hat{F}. \tag{5.12}$$

Note that $[\hat{A}_{\rm in}, \hat{A}^\dagger_{\rm in}] = [\hat{A}_{\rm out}, \hat{A}^\dagger_{\rm out}] = [\hat{B}_{\rm in}, \hat{B}^\dagger_{\rm in}] = 1$ as expected. As for the noise operator $\hat{F}$, the commutator and correlation functions are defined, using the known properties of $\hat{b}_{\rm in}(t)$ to lowest order as

$$[\hat{F}, \hat{F}^\dagger] \approx \gamma\tau + O(\gamma_{\rm OM}\gamma\tau^2), \tag{5.13}$$

$$\left\langle \hat{F}^\dagger \hat{F} \right\rangle = \gamma e^{\gamma_{\rm OM}\tau} \int_0^\tau ds\, e^{-\gamma_{\rm OM} s} n_{\rm b}(s), \quad \left\langle \hat{F}\hat{F}^\dagger \right\rangle = \gamma\tau + \left\langle \hat{F}^\dagger \hat{F} \right\rangle. \tag{5.14}$$

We now seek a unitary propagator $\hat{U}$ such that $\hat{A}_{\rm out} = \hat{U}^\dagger \hat{A}_{\rm in} \hat{U}$ and $\hat{B}_{\rm out} = \hat{U}^\dagger \hat{B}_{\rm in} \hat{U}$. Note that if we define $e^{\gamma_{\rm OM}\tau/2} = \cosh(r)$, $\sqrt{e^{\gamma_{\rm OM}\tau}-1} = \sinh(r)$, then the input/output equations have the form of two-mode squeezing between $\hat{A}_{\rm in}$ and $\hat{B}_{\rm in}$ in the absence of mechanical noise $\hat{F}$. Thus, we want $\hat{U} = \hat{U}_{\rm m}\hat{U}_0$, where

$$\hat{U}_0 = e^{ir\left(\hat{A}_{\rm in}\hat{B}_{\rm in} + \hat{A}^\dagger_{\rm in}\hat{B}^\dagger_{\rm in}\right)}, \tag{5.15}$$

$$\hat{U}^\dagger_{\rm m} \hat{B}_{\rm in} \hat{U}_{\rm m} = \hat{B}_{\rm in} + \hat{F}, \tag{5.16}$$

$$[\hat{U}_{\rm m}, \hat{A}_{\rm in}] = 0. \tag{5.17}$$

We find that a beam-splitter type interaction between $\hat{B}_{\rm in}$ and $\hat{F}$ satisfies this condition to lowest order. That is, if $\hat{U}_{\rm m} = e^{\hat{B}^\dagger_{\rm in}\hat{F} - \hat{B}_{\rm in}\hat{F}^\dagger}$, then using the above commutation relations for $\hat{F}$ we can show,

$$\begin{aligned} \hat{U}^\dagger_{\rm m} \hat{B}_{\rm in} \hat{U}_{\rm m} &= \cos\left(\sqrt{\gamma\tau}\right) \hat{B}_{\rm in} + (\gamma\tau)^{-1/2} \sin\left(\sqrt{\gamma\tau}\right) \hat{F} \\ &\approx \hat{B}_{\rm in} + \hat{F}. \end{aligned} \tag{5.18}$$

Given this propagator $\hat{U}$, which yields the approximate system evolution over the pulse interval

τ, and assuming an initial density matrix $\rho_0 = |0_\mathrm{A}\rangle\langle 0_\mathrm{A}| \otimes \rho_{\mathrm{B},0} \otimes \rho_\mathrm{F}$, where A, B, and F refer to the optical, mechanical, and bath subsystems, respectively, the density matrix for the mechanical resonator, ρ_B, conditioned upon detection of a single photon, is given by

$$\begin{aligned}\rho_\mathrm{B} &= \frac{\mathrm{Tr}_{\mathrm{A,F}}\left(|1_\mathrm{A}\rangle\langle 1_\mathrm{A}| \otimes I_\mathrm{B} \otimes I_\mathrm{F}\hat{U}\rho_0\hat{U}^\dagger\right)}{\mathrm{Tr}\left(|1_\mathrm{A}\rangle\langle 1_\mathrm{A}| \otimes I_\mathrm{B} \otimes I_\mathrm{F}\rho_0\right)} \\ &= P^{-1}\mathrm{Tr}_\mathrm{F}\left(\hat{U}_\mathrm{m}\langle 1_\mathrm{A}|\hat{U}_0\rho_0\hat{U}_0^\dagger|1_\mathrm{A}\rangle\hat{U}_\mathrm{m}^\dagger\right), \end{aligned} \tag{5.19}$$

where P is the overall probability of photodetection, given by

$$P = \mathrm{Tr}_\mathrm{B}\left(\langle 1_\mathrm{A}|\hat{U}_0|0_\mathrm{A}\rangle\langle 0_\mathrm{A}| \otimes \rho_{\mathrm{B},0}\hat{U}_0^\dagger|1_\mathrm{A}\rangle\right). \tag{5.20}$$

Denoting $\rho_{\mathrm{B,ideal}} = P^{-1}\hat{B}_\mathrm{in}^\dagger\rho_{\mathrm{B},0}\hat{B}_\mathrm{in}$ as the ideal conditional matrix in the absence of mechanical dissipation, and expanding $\hat{U}$ to lowest order in $\gamma_\mathrm{OM}\tau$, $\gamma\tau$ and $\left\langle\hat{F}^\dagger\hat{F}\right\rangle$, we find

$$P \approx \gamma_\mathrm{OM}\tau\left\langle\hat{B}_\mathrm{in}^\dagger\hat{B}_\mathrm{in}\right\rangle, \tag{5.21}$$

and

$$\begin{aligned}\rho_\mathrm{B} \approx &\left(1 - \left\langle\hat{F}^\dagger\hat{F}\right\rangle\right)\rho_\mathrm{B,ideal} + \left\langle\hat{F}^\dagger\hat{F}\right\rangle\hat{B}_\mathrm{in}^\dagger\rho_\mathrm{B,ideal}\hat{B}_\mathrm{in} + \left(\gamma\tau + \left\langle\hat{F}^\dagger\hat{F}\right\rangle\right)\hat{B}_\mathrm{in}\rho_\mathrm{B,ideal}\hat{B}_\mathrm{in}^\dagger \\ &- \frac{1}{2}\left(\gamma\tau + 2\left\langle\hat{F}^\dagger\hat{F}\right\rangle\right)\left(\hat{B}_\mathrm{in}^\dagger\hat{B}_\mathrm{in}\rho_\mathrm{B,ideal} + \rho_\mathrm{B,ideal}\hat{B}_\mathrm{in}^\dagger\hat{B}_\mathrm{in}\right). \end{aligned} \tag{5.22}$$

To evaluate the fidelity of the phonon addition subtraction we use the definition of fidelity between two quantum states ρ_1 and ρ_2, given by [173] $F = \mathrm{Tr}\left(\sqrt{\rho_1^{1/2}\rho_2\rho_1^{1/2}}\right)$, which is straightforward to evaluate for an initial thermal mechanical state, as both ρ_1 and ρ_2 will be diagonal in the Fock state basis.

Note that the only substantive difference in the case of phonon subtraction ($\Delta = \omega_\mathrm{m}$, red-detuned driving) is that the propagator $\hat{U}_0$ will have the form of a beam-splitter interaction rather than two mode squeezing. This will only change P and $\rho_\mathrm{B,ideal}$, which are now given by

$$P \approx \gamma_\mathrm{OM}\tau\left\langle\hat{B}_\mathrm{in}^\dagger\hat{B}_\mathrm{in}\right\rangle, \tag{5.23}$$

$$\rho_\mathrm{B,ideal} = \hat{B}_\mathrm{in}\rho_{\mathrm{B},0}\hat{B}_\mathrm{in}^\dagger, \tag{5.24}$$

while the definition of ρ_B in terms of $\rho_\mathrm{B,ideal}$ is unchanged assuming that the bath heating is only dependent on drive power and not drive detuning.

We assume the initial state of the mechanics to be a thermal state with average phonon number $\langle n\rangle$ given by Eq. 4.8, and evaluate $\left\langle\hat{F}^\dagger\hat{F}\right\rangle$ using the fit parameters of the effective time-dependent

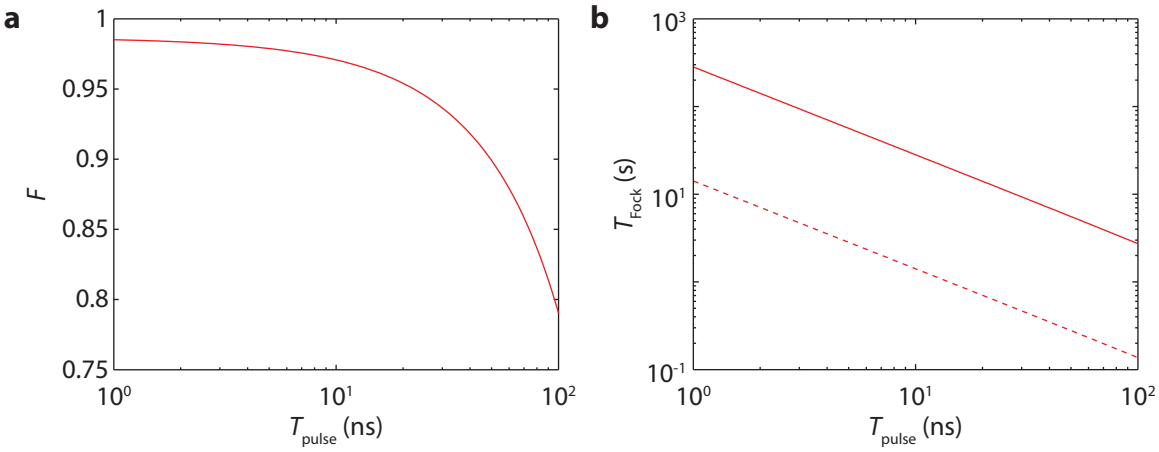

Figure 5.2: **Fock state generation.** **a**, Fidelity F of the generation of a single phonon Fock state versus pulse width T_{pulse} (pulse period $T_{\text{per}} = 1$ ms). **b**, Average time required to herald a Fock state T_{Fock} versus pulse width T_{pulse}. The solid line is calculated using the total detection efficiency measured in this work, while the dashed line uses the estimated ideal detection efficiency.

hot phonon bath.

Considering a range of pulse widths $0 < T_{\text{pulse}} \leq 100$ ns, which is sufficient to guarantee the necessary condition $\gamma_{\text{OM}} T_{\text{pulse}} \ll 1$, we find that for $T_{\text{per}} \geq 1$ ms there is no appreciable loss in fidelity due to long term heating causing an increase in $\langle n \rangle (t_0)$, and the primary loss in fidelity is due to transient heating during the pulse. Thus, we set $T_{\text{per}} = 1$ ms and calculate the fidelity of Fock state generation as a function of pulse width, displayed in Fig. 5.2a. For short ($T_{\text{pulse}} < 10$ ns) pulse widths, the fidelity approaches 98.5%, and remains above 80% for pulse widths up to 100 ns. However, it is equally important to quantify the expected time to herald such a state (i.e., the expected time before detection of a sideband photon), which in this case is given by $T_{\text{Fock}} = T_{\text{per}}/(\eta \gamma_{\text{OM}} T_{\text{pulse}} \langle n \rangle (t_0))$, where η is the total detection efficiency of the phonon counting measurement. This is shown versus T_{pulse} in Fig. 5.2b for two cases. The solid line is calculated the actual measured detection efficiency of the pulsed setup in Chapter 4, $\eta \approx 1\%$, while the dashed line is calculated using a realistic estimate of the ideal efficiency of our current measurement setup, $\eta \approx 5.5\%$. This latter figure is calculated using the measured efficiencies of the SPD ($\eta_{\text{SPD}} = 68\%$), the fiber-to-waveguide coupling efficiency ($\eta_{\text{cpl}} = 68\%$), and the waveguide-cavity coupling efficiency ($\kappa_{\text{e}}/\kappa = 0.5$), adding an additional 2 dB insertion loss per filter and 0.5 dB for the insertion loss of the optical circulator, which correspond to the highest efficiencies measured for these components in our lab. As shown in Fig. 5.2, even in the idealized case the expected time for the generation of a Fock state is $\gtrsim 100$ ms, which is much longer than the lifetime of the mechanical state. Thus, while it is feasible to use our current systems for the heralded, high-fidelity generation of non-classical mechanical states, it is still necessary to reduce heating effects (and thus allow for shorter pulse periods) in order to utilize this procedure for useful quantum information processing tasks (e.g., scalable entanglement distribution via the DLCZ

protocol [35, 36]).

5.2 2D Snowflake OMCs

While the Si nanobeam design used throughout this work has been optimized in terms of optomechanical coupling strength and optical quality, it represents one of the worst designs possible from the standpoint of heat dissipation, as any hot phonons generated within the cavity volume may only escape via a narrow, one-dimensional channel. A relatively simple way to mitigate some of the effects of heating is to move into a two-dimensional OMC structure such as the snowflake cavities described in section 1.3.6. Beyond improving the thermal properties, the planar nature of the snowflake structure would allow the development of a wider variety of optomechanical circuits composed of coupled arrays of OMC cavities. If successfully thermalized into the quantum ground state, such devices would form the basis for quantum phononic networks [174] and allow the study of quantum many-body physics in optomechanical metamaterials [166].

5.2.1 Heating

Estimating the specific heating behavior of such a cavity is non-trivial. In particular, as discussed in section 3.6 we have reason to believe the hot phonon bath formed in the nanobeam cavities is affected partly by a phonon bottleneck effects, and is thus dependent upon the specific cross-sectional dimensions of the structure. However, while we do not expect exactly the same effects in a 2D structure, the small thickness of the Si membrane (which is identical in both 1D and 2D cavities) should still provide a reduction of the phononic density of states compared to large structures, and thus a bottleneck effect should still arise. As the corresponding cutoff frequency should be determined by the lowest dimension of the structure, one would also expect the cutoff frequency to be similar.

Consequently, a crude estimate of the heating in the snowflake cavity can be made by assuming the nature of the hot phonon bath to be roughly the same as in the 1D case, with a similar cutoff frequency, and that the only substantive difference lies in the thermal conductance of the structure. As heat would still be carried predominantly by high-frequency phonons, one would still expect a thermal conductance scaling as T^3, so that $G_{\mathrm{th,2D}}/G_{\mathrm{th,1D}}$ is approximately constant. From a radiative model of heat transfer in the ballistic regime [175] one expects $G_{\mathrm{th}} \approx G_0 T^3$, where G_0 is a constant proportional to the effective cross-sectional area through which phonons may escape, A_{eff}, and otherwise dependent only on the material properties of Si. Then, assuming the absorbed power per photon is roughly the same as the nanobeam cavities we may relate the hot phonon temperature

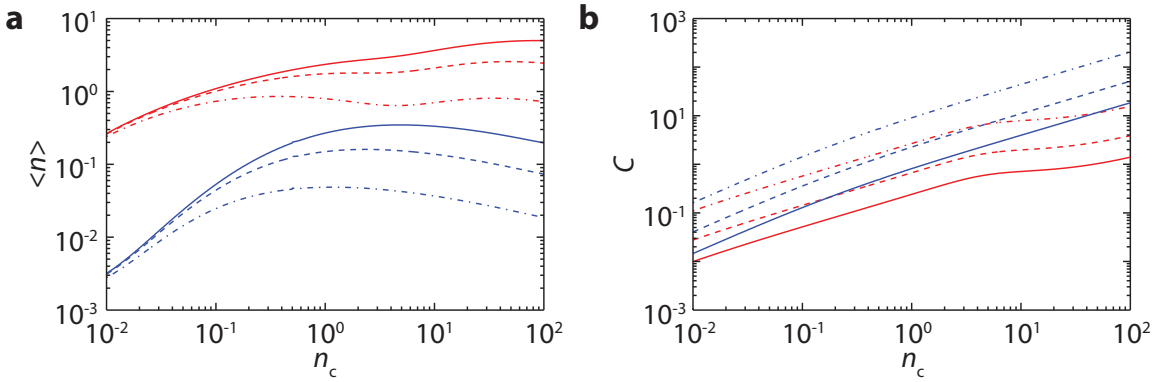

Figure 5.3: **Estimated heating in a snowflake OMC. a**, Estimated phonon occupancy $\langle n \rangle$ versus intracavity photon number n_c for the 2D snowflake cavity, assuming the same thermal conductance as for the 1D nanobeam (red) and the expected enhanced thermal conductance of the 2D slab (blue). Curves are calculated for $g_0/2\pi = 300$ kHz (solid), 500 kHz (dashed) and 1 MHz (dot-dashed). **b**, Cooperativity $C = \gamma_{OM}/(\gamma_0 + \gamma_p)$ versus n_c for the same sets of parameters as **a**.

T_p in the 1D and 2D cases:

$$\begin{aligned} P_{abs} &= \int_{T_0}^{T_{p,1D}} dT G_{th,1D}(T) \\ &= \frac{1}{4} G_{0,1D} \left(T_{p,1D}^4 - T_0^4 \right) \\ &= \frac{1}{4} G_{0,2D} \left(T_{p,2D}^4 - T_0^4 \right), \end{aligned}$$

which yields

$$T_{p,2D} = \left(\frac{A_{eff,1D}}{A_{eff,1D}} \left(T_{p,1D}^4 - T_0^4 \right) + T_0^4 \right)^{1/4}. \tag{5.25}$$

Given the empirical formula $T_{p,1D} = \left(D\, n_c + T_0^4 \right)^{1/4}$, we then have

$$T_{p,2D} = \left(\frac{A_{eff,1D}}{A_{eff,2D}} D\, n_c + T_0^4 \right)^{1/4}, \tag{5.26}$$

where D may be obtained from our fit to T_p in the 1D case (see Fig. 3.5b). Although the reduction in T_p will only scale as $(A_{eff,1D}/A_{eff,2D})^{1/4}$ for sufficiently large n_c, the exponential dependence of γ_p on T_p can lead to a significant change in the heating, as we will see.

For the 1D nanobeam $A_{eff,1D}$ is simply the cross-sectional area of the beam, while for the snowflake cavity we take $A_{eff,2D}$ to be the perimeter of the optical cavity region. For the heterostructure snowflake cavity shown in Fig. 1.13 this leads to $A_{eff,1D}/A_{eff,2D} \approx 0.03$. We take $\gamma_p(T_p)$ to have the same functional form and approximate cutoff temperature as shown in Fig. 3.6a, and assume $T_0 \approx 70$ mK, in accordance with our pulsed measurements of the minimum phonon occupancy of the nanobeams in the fridge. To determine $\gamma_{OM}(n_c)$, we take $g_0/2\pi = 300$ kHz, as was

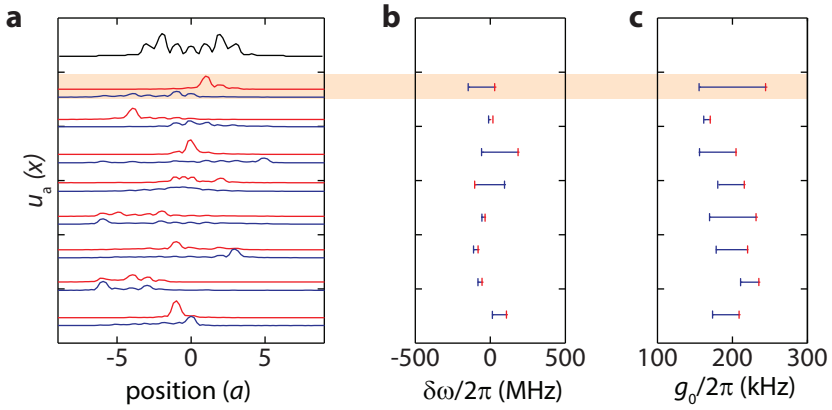

Figure 5.4: **Effects of disorder on the heterostructure snowflake cavity.** **a**, Plot of the linear acoustic energy density profile $u_{\mathrm{a}}(x)$, for the localized mechanical resonances with strong optomechanical coupling to the fundamental optical resonance. The top black curve corresponds to the unperturbed structure. Each pair of red and blue curves corresponds to the mechanical resonance with the largest and second largest magnitude of optomechanical coupling, respectively, for a different disordered structure; $u_{\mathrm{a}}(x)$ is computed by integrating the acoustic energy density across the transverse y and z dimensions. **b**, Plot of the corresponding mechanical frequency difference and **c**, optomechanical coupling rate g_0 for the two most strongly coupled mechanical resonances. Here, we show results for a representative eight of the simulated disordered structures.

simulated for the heterostructure snowflake cavity, and assume κ is roughly the same as the value measured for nanobeams in this work. We assume the same value of radiation-limited mechanical Q as was measured in the nanobeams ($Q_{\mathrm{m}} \approx 1.7\times10^7$), which for the higher frequency snowflake mode ($\omega_{\mathrm{m}}/2\pi = 9.5$ GHz) corresponds to $\gamma_0/2\pi \approx 560$ Hz. The corresponding estimated values of $\langle n\rangle$ and cooperativity $C = \gamma_{\mathrm{OM}}/(\gamma_0+\gamma_{\mathrm{p}})$ under CW drive conditions are shown versus n_{c} in Fig. 5.3a and b, respectively. The 2D estimates for the parameters stated above are shown as solid blue lines, while a solid red line shows the corresponding values assuming the same $T_{\mathrm{p}}(n_{\mathrm{c}})$ curve measured in the 1D case (i.e., no increase in thermal conductance). From this we see that even the relatively modest increase in thermal conductance leads to a drastic change in the heating, with $C \geq 1$ and $\langle n\rangle \lesssim 0.1$ achieved for $n_{\mathrm{c}} \geq 1$. While the snowflake design presented in section 1.3.6 has a relatively low value of g_0, there exists room for optimization. In particular, a different design presented in Ref. [30] possesses $g_0/2\pi \approx 500$ kHz. For comparison, the heating curves for $g_0 = 500$ kHz and $g_0 = 1$ MHz are shown in Fig. 5.3 as dashed and dot-dashed lines, respectively. Provided a snowflake cavity was designed with coupling comparable to the nanobeams, one could achieve $C \gg 1$ and $\langle n\rangle \sim 10^{-2}$ without using pulsed driving.

5.2.2 Experimental Demonstration

While the snowflake OMC designs are not as mature experimentally as the nanobeam cavities, initial fabrication and measurement of snowflake heterostructure cavities is detailed in Ref. [176]. While

this initial demonstration is promising, it reveals a fundamental flaw in the initial heterostructure design. In particular, when measuring the mechanical modes of the snowflake cavity, multiple closely spaced modes were observed near the expected resonance frequency of $\omega_{\rm m}/2\pi \approx 9.4$ GHz, each with a relatively weak optomechanical coupling.

This is explained by noting that the extremely flat dispersion of the acoustic waveguide mode from which the cavity is formed (see Fig. 1.12c) causes the spectrum of localized mechanical cavity modes to be highly sensitive to unavoidable fabrication disorder. To investigate this, the localized optical and mechanical modes for 50 different disordered structures were calculated numerically via FEM simulation, the results of which are summarized in Fig. 5.4. As with the simulations of mechanical Q described in section 1.3.5, disorder was introduced into the structures by varying the position, width, and radius of the snowflake holes according to a normal distribution with a given percent standard deviation about the nominal value. Based upon SEM images of the fabricated structures, we selected a standard deviation of $\sigma = 2\%$, which is the same degree of disorder which yields good correspondence between the simulated and measured $Q_{\rm m}$ at low temperature in the nanobeam cavities (see sections 1.3.5 and 4.2). Roughly 10% of the simulated disordered structures yield localized mechanical resonances with the frequency splitting and optomechanical coupling rates measured in the actual device.

As most of the simulated disordered structures possess values of g_0 well below the expected theoretical value of $g_0/2\pi \approx 300$ kHz, this flaw in the snowflake design presents a significant practical challenge to the goal of replacing nanobeam cavities. The high degree of sensitivity to disorder will need to be addressed in future generations of snowflake cavities by redesigning the $W1$ waveguide to yield larger dispersion in the desired mechanical band.

Appendix A

Mathematical Definition and Additional Derivations

A.1 Definitions

We use the following symmetric definition for the Fourier transform of an operator and the inverse transform:

$$\begin{aligned}\hat{A}(\omega) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} dt\; e^{i\omega t} \hat{A}(t), \\ \hat{A}(t) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} d\omega\; e^{-i\omega t} \hat{A}(\omega).\end{aligned} \tag{A.1}$$

Note that the same definition is used when taking the transform of the conjugate. That is,

$$\hat{A}^\dagger(\omega) = \frac{1}{\sqrt{2}} \int_{-\infty}^{\infty} dt\; e^{i\omega t} \hat{A}^\dagger(t),$$

so that $\hat{A}^\dagger(\omega) = \left(\hat{A}(-\omega)\right)^\dagger$. The power spectral density of an operator $\hat{A}$ is defined,

$$\begin{aligned}S_{\mathrm{AA}}(\omega) &= \int_{-\infty}^{\infty} d\tau e^{i\omega\tau} \left\langle \hat{A}^\dagger(t+\tau)\hat{A}(t) \right\rangle \\ &= \int_{-\infty}^{\infty} d\omega' \left\langle \hat{A}^\dagger(\omega)\hat{A}(\omega') \right\rangle,\end{aligned} \tag{A.2}$$

where $\left\langle \hat{A}^\dagger \hat{A} \right\rangle$ denotes the usual quantum expectation value and it is assumed we are dealing with stationary processes, where the two-time expectation value depends only on the time-difference τ, so that in general we may take $t = 0$. The *symmetrized* spectral density is defined as $\bar{S}_{\mathrm{AA}} = \frac{1}{2}\left(S_{\mathrm{AA}}(\omega) + S_{\mathrm{AA}}(-\omega)\right)$.

A.2 Input-Output Theory

In this section we briefly summarize the steps taken to derive the type of Heisenberg-Langevin equation shown in section 1.1.1. A full derivation can be found in Refs. [4] and [5]. We consider a single cavity with resonance frequency ω_c and annihilation (creation) operator $\hat{a}$ ($\hat{a}^\dagger$), defined by the system Hamiltonian:

$$\hat{H}_{\rm sys} = \hbar\omega_c \hat{a}^\dagger \hat{a}. \tag{A.3}$$

The external bath consists of a continuum of independent harmonic oscillators, with Hamiltonian

$$\hat{H}_{\rm bath} = \hbar \int_{-\infty}^{\infty} d\omega\ \omega\ \hat{A}^\dagger(\omega)\hat{A}(\omega), \tag{A.4}$$

where the bath creation and annihilation operators obey the commutation relations $\left[\hat{A}(\omega), \hat{A}^\dagger(\omega')\right] = \delta(\omega - \omega')$. Note that using $\omega = -\infty$ as the lower limit (rather than 0) is justified in the case when $\omega_c \gg 1$, as we only care about coupling to some frequency range in the vicinity of the cavity frequency. Assuming a generic linear interaction between the cavity and bath, and making a rotating wave approximation, the system-bath interaction Hamiltonian has the form

$$\hat{H}_{\rm int} = i\hbar \int_{-\infty}^{\infty} d\omega\ g(\omega) \left(\hat{A}^\dagger(\omega)\hat{a} - \hat{a}^\dagger \hat{A}(\omega)\right), \tag{A.5}$$

where $g(\omega)$ is the mode coupling rate [1]. The full Hamiltonian is given by $\hat{H} = \hat{H}_{\rm sys} + \hat{H}_{\rm bath} + \hat{H}_{\rm int}$. Using the Heisenberg equation of motion $\dot{\hat{A}} = \frac{i}{\hbar}\left[\hat{H}, \hat{A}\right] + \frac{\partial \hat{A}}{\partial t}$, we obtain the equations for both $\dot{\hat{a}}$ and $\dot{\hat{A}}(\omega)$. We may then formally integrate the equations of motion for $\dot{\hat{A}}(\omega)$ starting from initial time $t_0 < t$, and insert them into the equation for $\dot{\hat{a}}$ to arrive at

$$\dot{\hat{a}}(t) = -i\omega_c \hat{a}(t) - \int_{-\infty}^{\infty} d\omega\ g(\omega)\ e^{-i\omega(t-t_0)} \hat{A}(\omega, t_0) - \int_{-\infty}^{\infty} d\omega\ g(\omega)^2 \int_{t_0}^{t} dt' e^{-i\omega(t-t')} \hat{a}(t'), \tag{A.6}$$

Unfortunately, this equation for $\dot{\hat{a}}(t)$ depends on the state of the system at $t' \neq t$. In order to simplify, we assume that the bath coupling rate $g(\omega)$ is sufficiently slow-varying within the frequency range of interest (centered on ω_c) in the that is may be approximated as a constant $g(\omega) \approx \sqrt{\kappa/2\pi}$. This approximation, known as the Markov approximation, allows us to obtain the approximate equation of motion for the cavity field

$$\dot{\hat{a}} \approx -\left(i\omega_c + \frac{\kappa}{2}\right)\hat{a} + \sqrt{\kappa}\hat{a}_{\rm in}, \tag{A.7}$$

[1] It is not strictly necessary to assume the form given here for $\hat{H}_{\rm sys}$ and $\hat{H}_{\rm int}$. Rather, we may assume an arbitrary system Hamiltonian, with $\hat{H}_{\rm int}$ then taken as linear in some system operator $\hat{d}$. For simplicity of presentation we have chosen to explicitly present the relevant case for this work.

where the input noise operator is defined in terms of the initial state of the bath at t_0 as

$$\hat{a}_{\rm in} = \frac{-1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} d\omega e^{-i\omega(t-t_0)} \hat{A}(\omega, t_0), \tag{A.8}$$

which can be shown to obey the commutation relation $\left[\hat{a}_{\rm in}(t), \hat{a}_{\rm in}^{\dagger}(t')\right] = \delta(t-t')$. We may also write a time-reversed equation using the state of the bath at time $t_f > t$ as the initial condition, finding

$$\dot{\hat{a}} \approx -\left(i\omega_{\rm c} - \frac{\kappa}{2}\right)\hat{a} + \sqrt{\kappa}\hat{a}_{\rm out}, \tag{A.9}$$

where the output noise operator is defined as

$$\hat{a}_{\rm out} = \frac{-1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} d\omega e^{-i\omega(t-t_f)} \hat{A}(\omega, t_f). \tag{A.10}$$

Equating Eqs. A.7 and A.9 yields the input-output boundary condition

$$\hat{a}_{\rm out} = \hat{a}_{\rm in} - \sqrt{\kappa}\hat{a}. \tag{A.11}$$

This analysis can be performed for both the optical and mechanical resonators to obtain the equations of motion used in section 1.1.1. For the optical resonator we assume two completely uncorrelated baths, corresponding to the extrinsic and intrinsic decay channels. Although an input-output boundary condition may be defined for the intrinsic optical bath and the mechanical bath, by definition these baths are unmonitored and so such a relation is of no use in our analysis.

A.2.1 Bath Correlation Functions

We assume that the bath is initially in a thermal state. That is, every mode comprising the bath exists at a single well-defined temperature $T_{\rm b}$ at time t_0. This implies that

$$\left\langle \hat{A}^{\dagger}(\omega, t_0)\hat{A}(\omega', t_0)\right\rangle = n(\omega)\delta(\omega - \omega'), \tag{A.12}$$

$$\left\langle \hat{A}(\omega, t_0)\hat{A}^{\dagger}(\omega', t_0)\right\rangle = (n(\omega)+1)\delta(\omega - \omega'), \tag{A.13}$$

where $n(\omega) = \left(e^{\hbar\omega/k_{\rm B}T_{\rm b}} - 1\right)^{-1}$ is the thermal occupancy of the bath mode, and all other one- and two-time correlation functions are equal to zero. Using the definition of the input noise operator, it follows that

$$\left\langle \hat{a}_{\rm in}^{\dagger}(t)\hat{a}_{\rm in}(t')\right\rangle = \int_{-\infty}^{\infty} d\omega\ e^{i\omega(t-t')}\ n(\omega). \tag{A.14}$$

Assuming that we only care about a narrow bandwidth about the cavity frequency (i.e., $\kappa \ll \omega_{\rm c}$),

we may approximate $n(\omega) \approx n(\omega_c)$ which yields

$$\left\langle \hat{a}_{\rm in}^{\dagger}(t)\hat{a}_{\rm in}(t') \right\rangle \approx n(\omega_c)\delta(t-t'). \tag{A.15}$$

From this, $\left\langle \hat{a}_{\rm in}(t)\hat{a}_{\rm in}^{\dagger}(t') \right\rangle$ follows from the input noise operator commutation relation, and all other one- and two-time correlation functions of $\hat{a}_{\rm in}$ and $\hat{a}_{\rm in}^{\dagger}$ are easily shown to be equal to zero.

A.3 Spectral Density of the Mechanical Resonator

Consider first the Heisenberg-Langevin equation for the mechanical resonator in the absence of optomechanical coupling

$$\dot{\hat{b}} = -\left(i\omega_{\rm m} + \frac{\gamma_{\rm i}}{2}\right)\hat{b} + \sqrt{\gamma_{\rm i}}\hat{b}_{\rm in}. \tag{A.16}$$

Transforming into the frequency domain and computing the power spectral density of $\hat{b}$ according to the definition given in section A.1, we obtain

$$\begin{aligned} S_{\rm bb}(\omega) &= \int_{-\infty}^{\infty} d\omega' \frac{1}{-i(\omega_{\rm m}+\omega)+\gamma_{\rm i}/2}\frac{1}{i(\omega_{\rm m}-\omega')+\gamma_{\rm i}/2}\gamma_{\rm i}\left\langle \hat{b}_{\rm in}^{\dagger}(\omega)\hat{b}_{\rm in}(\omega') \right\rangle \\ &= \int_{-\infty}^{\infty} d\omega' \frac{1}{-i(\omega_{\rm m}+\omega)+\gamma_{\rm i}/2}\frac{1}{i(\omega_{\rm m}-\omega')+\gamma_{\rm i}/2}\gamma_{\rm i} n_{\rm b}\delta(\omega+\omega') \\ &= \frac{\gamma_{\rm i} n_{\rm b}}{(\omega_{\rm m}+\omega)^2+(\gamma_{\rm i}/2)^2}. \end{aligned} \tag{A.17}$$

As we will see, this function and its symmetrized counterpart will naturally arise in our derivation of the detected output noise spectrum of the optomechanical system. In general, we define the spectral density in terms of an arbitrary average phonon number and mechanical damping rate as

$$S_{\rm bb}(\omega;\langle n\rangle) = \frac{\gamma\langle n\rangle}{(\omega_{\rm m}+\omega)^2+(\gamma/2)^2}. \tag{A.18}$$

The corresponding spectral density for the creation operator is found to be

$$S_{\rm b^{\dagger}\,b^{\dagger}}(\omega;\langle n\rangle) = \frac{\gamma(\langle n\rangle+1)}{(\omega_{\rm m}-\omega)^2+(\gamma/2)^2}. \tag{A.19}$$

Note that these spectral densities have the important property

$$\begin{aligned} \int_{-\infty}^{\infty}\frac{d\omega}{2\pi}S_{\rm bb}(\omega;\langle n\rangle) &= \langle n\rangle, \\ \int_{-\infty}^{\infty}\frac{d\omega}{2\pi}S_{\rm b^{\dagger}\,b^{\dagger}}(\omega;\langle n\rangle) &= \langle n\rangle + 1. \end{aligned} \tag{A.20}$$

A.4 Phonon Occupancy

The steady-state phonon occupancy can be defined in terms of the frequency domain fluctuation operators as

$$\langle n \rangle = \left\langle \hat{b}^\dagger(t)\hat{b}(t) \right\rangle = \frac{1}{2\pi}\int_{-\infty}^{\infty} d\omega \int_{-\infty}^{\infty} d\omega' \; e^{-i(\omega+\omega')t} \; \left\langle \hat{b}^\dagger(\omega)\hat{b}(\omega') \right\rangle . \tag{A.21}$$

Inserting the full expression for $\hat{b}(\omega)$ given in Eq. 1.16, using the two-time correlation definitions given in section 1.1.1, and integrating over ω', we obtain

$$\begin{aligned} \langle n \rangle &= \frac{1}{2\pi}\int_{-\infty}^{\infty} d\omega |\chi_{\rm m}(\omega)|^2 \left(\gamma_{\rm i} n_{\rm b} + G^2\kappa |\chi_{\rm o}^*(-\omega)|^2\right) \\ &= \int_{-\infty}^{\infty} d\omega \frac{\gamma/2\pi}{(\omega_{\rm m}-\omega)^2 + (\gamma/2)^2}\left(\frac{\gamma_{\rm i} n_{\rm b}}{\gamma} + \frac{G^2\kappa}{\gamma}\frac{1}{(\Delta+\omega)^2+(\kappa/2)^2}\right). \end{aligned} \tag{A.22}$$

For $\kappa \ll \omega_{\rm m}$, $\Delta = \pm\omega_{\rm m}$, we can focus only on frequencies near $\omega_{\rm m}$ ($|\omega - \omega_{\rm m}| \ll \kappa$) and approximate $\Delta + \omega \approx \Delta + \omega_{\rm m}$. In the red-detuned ($\Delta = \omega_{\rm m}$) case we obtain

$$\begin{aligned} \langle n \rangle &\approx \frac{\gamma_{\rm i} n_{\rm b}}{\gamma} + \frac{\gamma_{\rm OM}}{\gamma}\left(\frac{\kappa}{4\omega_{\rm m}}\right)^2 \\ &= \frac{n_{\rm b}}{1+C} + \frac{C}{1+C}\left(\frac{\kappa}{4\omega_{\rm m}}\right)^2, \end{aligned} \tag{A.23}$$

where we have defined the cooperativity $C = \gamma_{\rm OM}/\gamma_{\rm i}$. Likewise, for blue-detuning ($\Delta = -\omega_{\rm m}$) we obtain

$$\begin{aligned} \langle n \rangle &\approx \frac{\gamma_{\rm i} n_{\rm b}}{\gamma} + \frac{\gamma_{\rm OM}}{\gamma} \\ &= \frac{n_{\rm b}}{1-C} + \frac{C}{1-C}. \end{aligned} \tag{A.24}$$

A.5 Scattering Matrix Elements

As described in section 1.1.3, for a sideband-resolved ($\kappa \ll \omega_{\rm m}$) system under red- or blue-detuned driving conditions ($\Delta = \pm\omega_{\rm m}$, respectively), the optical cavity output in the frequency domain may be related to the optical and mechanical noise operators by the equations

$$\hat{a}_{\rm out}(\omega)|_{\Delta=\omega_{\rm m}} \approx r(\omega;+)\hat{a}_{\rm in}(\omega) + n(\omega;+)\hat{a}_{\rm i}(\omega) + s(\omega;+)\hat{b}_{\rm in}(\omega), \tag{A.25}$$

and

$$\hat{a}_{\rm out}(\omega)|_{\Delta=-\omega_{\rm m}} \approx r(\omega;-)\hat{a}_{\rm in}(\omega) + n(\omega;-)\hat{a}_{\rm i}(\omega) + s(\omega;-)\hat{b}^\dagger_{\rm in}(\omega), \tag{A.26}$$

where the scattering matrix elements are given by

$$r(\omega;\pm) = 1 - \frac{2\kappa_{\mathrm{e}}}{\kappa} \pm \frac{\gamma_{\mathrm{OM}}\kappa_{\mathrm{e}}}{\kappa} \frac{1}{\pm i\left(\omega_{\mathrm{m}} \mp \omega\right) + \gamma/2}, \tag{A.27}$$

$$n(\omega;\pm) = \pm\frac{\sqrt{\kappa_{\mathrm{i}}\kappa_{\mathrm{e}}}}{\kappa} \left(\frac{\gamma_{\mathrm{OM}}}{\pm i\left(\omega_{\mathrm{m}} \mp \omega\right) + \gamma/2} \mp 2 \right), \tag{A.28}$$

and

$$s(\omega;\pm) = \sqrt{\frac{\kappa_{\mathrm{e}}}{\kappa}} \frac{i\sqrt{\gamma_{\mathrm{i}}\gamma_{\mathrm{OM}}}}{\pm i\left(\omega_{\mathrm{m}} \mp \omega\right) + \gamma/2}. \tag{A.29}$$

It can be shown that, in general, the scattering matrix elements have the property $|r(\omega;\pm)|^2 + |n(\omega;\pm)|^2 \pm |s(\omega;\pm)|^2 = 1$.

Appendix B

Coupled-Cavity Model for Butt-Couplers

The butt-coupler reflection spectrum, and the corresponding effects on the reflected cavity signal, can be modeled with reasonable accuracy by the system shown in Fig. B.1. The coupling waveguide, including the weak reflection from the fiber-to-nanowire interface and the high-reflectivity photonic crystal end-mirror, is modeled by a Fabry-Perot etalon, with front and back mirror reflectivities r_1 and r_2, respectively. We also assume a total round-trip power transmission efficiency η_{rt} and a phase shift θ_{rt}. The round-trip phase shift can also be written $\theta_{\mathrm{rt}} = 2kL_{\mathrm{eff}}$, where $k = \omega_{\mathrm{l}}/c$ is the free-space wavenumber of the incident drive laser at frequency ω_{l}, and L_{eff} is the effective length of the etalon. The effective length will be a little longer than the physical length of the coupling waveguide, due to the fact that the interior of the etalon includes the dielectric waveguide (i.e not vacuum). In practice the effective length can be determined empirically from the free-spectral range of the etalon (the spacing of the broad fringes in Fig. 2.6a). The front mirror is taken to have a real reflection coefficient $r_1 = \sqrt{R_{\mathrm{fib}}}$ (and corresponding transmission coefficient $t_1 = \sqrt{1 - R_{\mathrm{fib}}}$), where R_{fib} is the reflectivity of the fiber endface. We assume coupling to a high-finesse optical cavity (with total decay rate κ) at a rate κ_{e} via the back-mirror, which is otherwise perfectly reflective far from resonance. If this secondary cavity is high-finesse enough, we may consider only a single resonant mode, in which case the back mirror reflection coefficient can be given by the usual result

$$r_2 = 1 - \frac{\kappa_{\mathrm{e}}}{i\Delta + \kappa/2}, \tag{B.1}$$

where $\Delta = \omega_{\mathrm{c}} - \omega_{\mathrm{l}}$ is the cavity-drive detuning.

For an incident drive amplitude A_{in}, the reflected amplitude A_{out} may be written as a sum of an infinite number of partial reflections produced by multiple reflections between the two surfaces

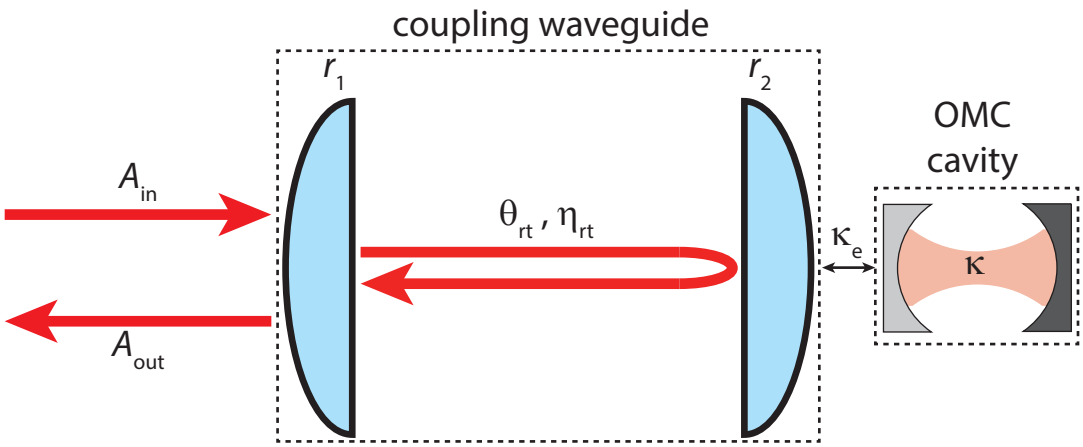

Figure B.1: **Coupled-cavity schematic.** The butt-coupled cavity can be modeled by a Fabry-Perot etalon, with round-trip phase shift and transmission efficiency $\theta_{\rm rt}$ and $\eta_{\rm rt}$, respectively, and front and back mirror reflectivities r_1 and r_2. This cavity is in turn coupled to the high-finesse OMC cavity at rate $\kappa_{\rm e}$. This coupling can be incorporated into the definition of the back mirror reflectivity r_2.

of the etalon. This results in the expression [177]

$$\begin{aligned} A_{\rm out} &= A_{\rm in}\left(r_1 - \left(1-r_1^2\right) r_2\sqrt{\eta_{\rm rt}}e^{-i\theta_{\rm rt}} \sum_{n=0}^{\infty}\left(r_1 r_2\sqrt{\eta_{\rm rt}}e^{-i\theta_{\rm rt}}\right)^n\right) \\ &= A_{\rm in}\left(r_1 - \frac{\left(1-r_1^2\right) r_2\sqrt{\eta_{\rm rt}}e^{-i\theta_{\rm rt}}}{1 - r_1 r_2\sqrt{\eta_{\rm rt}}e^{-i\theta_{\rm rt}}}\right) \\ &= A_{\rm in}\left(\frac{r_1 - r_2\sqrt{\eta_{\rm rt}}e^{-i\theta_{\rm rt}}}{1 - r_1 r_2\sqrt{\eta_{\rm rt}}e^{-i\theta_{\rm rt}}}\right). \end{aligned} \tag{B.2}$$

Far away from cavity resonance, we can approximate $r_2 \approx 1$. Then, inserting $r_1 = \sqrt{R_{\rm fib}}$ and $\sqrt{\eta_{\rm rt}} = \eta_{\rm cpl}$ into Eq. B.2, the total normalized reflection $R = |A_{\rm out}/A_{\rm in}|^2$ is given by

$$R = \frac{\left(\sqrt{R_{\rm fib}} - \eta_{\rm cpl}\right)^2 + 4\eta_{\rm cpl}\sqrt{R_{\rm fib}}\sin^2\left(\theta_{\rm rt}/2\right)}{\left(1 - \eta_{\rm cpl}\sqrt{R_{\rm fib}}\right)^2 + 4\eta_{\rm cpl}\sqrt{R_{\rm fib}}\sin^2\left(\theta_{\rm rt}/2\right)}, \tag{B.3}$$

from which we derive Eq. 2.34 using $R_{\rm min} = R\left(\theta_{\rm rt} = 0\right)$ and $R_{\rm max} = R\left(\theta_{\rm rt} = \pi\right)$. The full expression including the cavity, obtained by combining Eqs. B.1 and B.2, can be used to fit the cavity lineshape in the presence of interference from the etalon, and is used to produce the plots in Fig. 2.7.

Appendix C

Superconducting Nanowire Single Photon Detectors

The SPDs used in this work are amorphous WSi-based superconducting nanowire single-photon detectors developed in collaboration between the Jet Propulsion Laboratory and NIST, designed for high-efficiency detection of individual photons in the wavelength range $\lambda = 1520 - 1610$ nm with maximum count rates of about 2.5×10^7 counts per second (c.p.s) [178]. The SPDs are mounted on the still stage of the dilution refrigerator at ~ 700 mK. Single-mode optical fibers (Corning SMF-28) are passed into the refrigerator through vacuum feedthroughs and coupled to the SPDs via a fiber sleeve attached to each SPD mount. Proper alignment of the incoming fiber with the 15 μm $\times$ 15 μm square area of the SPD nanowire is ensured by a self-aligned mounting system incorporated into the design of the SPD [178]. The radio-frequency output of each SPD is amplified by a cold-amplifier mounted on the 50 K stage of the refrigerator as well as a room-temperature amplifier, then read out by a triggered PicoQuant PicoHarp 300 time-correlated single photon counting module. By systematically isolating the input optical fiber from environmental light sources and filtering out long wavelength blackbody radiation inside the fridge we have achieved dark count rates of ~ 4 (c.p.s.). At just below the switching current of the detectors, we have measured a peak detection efficiency of $\eta_{\mathrm{SPD}} = 68\%$, with $\lesssim 20\%$ variability depending on photon polarization.

Appendix D

Publications

S. M. Meenehan, J. D. Cohen, G. S. MacCabe, and O. Painter, "Pulsed excitation dynamics of an optomechanical crystal resonator near its quantum ground-state of motion," *in preparation* (2015).

J. D. Cohen, S. M. Meenehan, G. S. MacCabe, A. H. Safavi-Naeini, F. Marsili, M. D. Shaw, and O. Painter, "Phonon counting and intensity interferometry of a nanomechanical resonator," *Nature*, in press (arXiV:1410.1047) (2015).

S. M. Meenehan, J. D. Cohen, S. Gröblacher, J. T. Hill, A. H. Safavi-Naeini, M. Aspelmeyer, and O. Painter, "Silicon optomechanical crystal resonator at millikelvin temperatures," *Phys. Rev. A*, vol. 90, 011803(R) (2014).

A. H. Safavi-Naeini, J. T. Hill, S. Meenehan, J. Chan, S. Gröblacher, and O. Painter, "Two-dimensional phononic-photonic bandgap optomechanical crystal cavity," *Phys. Rev. Lett.*, vol. 112, 153603 (2014).

S.-P. Yu, J. D. Hood, J. A. Muniz, M. J. Martin, R. Norte, C.-L. Hung, S. M. Meenehan, J. D. Cohen, O. Painter, and H. J. Kimble, "Nanowire photonic crystal waveguides for single-atom trapping and strong light-matter interactions," *Appl. Phys. Lett.*, vol. 104, 111103 (2014).

C.-L. Hung, S. M. Meenehan, D. E. Chang, O. Painter, and H. J. Kimble, "Trapped atoms in one-dimensional photonic crystals," *New J. Phys.*, vol. 15, 083026 (2013).

J. D. Cohen, S. M. Meenehan, and O. Painter, "Optical coupling to nanoscale optomechanical cavities for near quantum-limited motion transduction," *Opt. Express*, vol. 21, 11227 (2013).

J. Chan, A. H. Safavi-Naeini, J. T. Hill, S. Meenehan, and O. Painter, "Optimized optomechanical

crystal cavity with acoustic radiation shield," *Appl. Phys. Lett.*, vol. 101, 081115 (2012).

M. Winger, T. D. Blasius, T. P. Mayer Alegre, A. H. Safavi-Naeini, S. Meenehan, J. Cohen, S. Stobbe, and O. Painter, "A chip-scale integrated cavity-electro-optomechanics platform," *Opt. Express*, vol. 19, 24905 (2011).

R. Perahia, J. D. Cohen, S. Meenehan, T. P. Mayer Alegre, and O. Painter, "Electrostatically tunable optomechanical "zipper" cavity laser," *Appl. Phys. Lett.*, vol. 97, 191112 (2010).

Bibliography

[1] J. Chan, *Laser cooling of an optomechanical crystal resonator to its quantum ground state of motion*. PhD thesis, California Institute of Technology, 2012.

[2] A. H. Safavi-Naeini, *Quantum optomechanics with silicon nanostructures*. PhD thesis, California Institute of Technology, 2013.

[3] C. K. Law, "Effective Hamiltonian for the radiation in a cavity with a moving mirror and a time-varying dielectric medium," *Phys. Rev. A*, vol. 49, pp. 433–437, 1994.

[4] C. W. Gardiner and M. J. Collett, "Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation," *Phys. Rev. A*, vol. 31, pp. 3761–3774, 1985.

[5] C. W. Gardiner and P. Zoller, *Quantum Noise*. Springer Series in Synergetics, 2004.

[6] M. Ludwig, B. Kubala, and F. Marquardt, "The optomechanical instability in the quantum regime," *New J. Phys.*, vol. 10, p. 095013, 2008.

[7] P. Rabl, "Photon blockade effect in optomechanical systems," *Phys. Rev. Lett.*, vol. 107, p. 063601, 2011.

[8] A. Nunnenkamp, K. Borkje, and S. M. Girvin, "Single-photon optomechanics," *Phys. Rev. Lett.*, vol. 107, p. 063602, 2011.

[9] F. Marquardt, J. G. E. Harris, and S. M. Girvin, "Dynamical multistability induced by radiation pressure in high-finesse micromechanical optical cavities," *Phys. Rev. Lett.*, vol. 96, p. 103901, 2006.

[10] D. A. Rodrigues and A. D. Armour, "Amplitude noise suppression in cavity-driven oscillations of a mechanical resonator," *Phys. Rev. Lett.*, vol. 104, p. 053601, 2010.

[11] H. P. Yuen and V. W. S. Chan, "Noise in homodyne and heterodyne detection," *Opt. Lett.*, vol. 8, pp. 177–179, 1983.

[12] B. L. Schumaker, "Noise in homodyne detection," *Opt. Lett.*, vol. 9, no. 5, pp. 189–191, 1984.

[13] A. H. Safavi-Naeini, J. Chan, J. T. Hill, T. P. M. Alegre, A. Krause, and O. Painter, "Observation of quantum motion of a nanomechanical resonator," *Phys. Rev. Lett.*, vol. 108, p. 033602, 2012.

[14] A. J. Weinstein, C. U. Lei, E. E. Wollman, J. Suh, A. Metelmann, A. A. Clerk, and K. C. Schwab, "Observation and interpretation of motional sideband asymmetry in a quantum electromechanical device," *Phys. Rev. X*, vol. 4, p. 041003, 2014.

[15] D. Lee, M. Underwood, D. Mason, A. B. Shkarin, K. Børkje, S. M. Girvin, and J. G. E. Harris, "Observation of quantum motion in a nanogram-scale object," *arXiv:1406.7254*, 2014.

[16] T. P. Purdy, P.-L. Yu, N. S. Kampel, R. W. Peterson, K. Cicak, R. W. Simmonds, and C. A. Regal, "Optomechanical Raman-ratio thermometry," *arXiv:1406.7247*, 2014.

[17] L. Diósi, "Laser linewidth hazard in optomechanical cooling," *Phys. Rev. A*, vol. 78, p. 021801, 2008.

[18] A. Schliesser, R. Rivière, G. Anetsberger, O. Arcizet, and T. J. Kippenberg, "Resolved-sideband cooling of a micromechanical oscillator," *Nature Phys.*, vol. 4, pp. 415–419, 2008.

[19] P. Rabl, C. Genes, K. Hammerer, and M. Aspelmeyer, "Phase-noise induced limitations on cooling and coherent evolution in optomechanical systems," *Phys. Rev. A*, vol. 80, p. 063819, 2009.

[20] G. A. Phelps and P. Meystre, "Laser phase noise effects on the dynamics of optomechanical resonators," *Phys. Rev. A*, vol. 83, p. 063838, June 2011.

[21] M. Poggio, C. L. Degen, H. J. Mamin, and D. Rugar, "Feedback cooling of a cantilever's fundamental mode below 5 mK," *Phys. Rev. Lett.*, vol. 99, p. 017201, 2007.

[22] T. Rocheleau, T. Ndukum, C. Macklin, J. B. Hertzberg, A. A. Clerk, and K. C. Schwab, "Preparation and detection of a mechanical resonator near the ground state of motion," *Nature*, vol. 463, pp. 72–75, 2010.

[23] 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 J. Phys.*, vol. 15, p. 035007, 2013.

[24] M. L. Gorodetsky and I. S. Grudinin, "Fundamental thermal fluctuations in microspheres," *J. Opt. Soc. Am. B*, vol. 21, no. 4, pp. 697–705, 2004.

[25] A. Schliesser, G. Anetsberger, R. Rivière, O. Arcizet, and T. J. Kippenberg, "High-sensitivity monitoring of micromechanical vibration using optical whispering gallery mode resonators," *New J. Phys.*, vol. 10, no. 9, p. 095015, 2008.

[26] J. Komma, C. Schwarz, G. Hofmann, D. Heinert, and R. Nawrodt, "Thermo-optic coefficient of silicon at 1550 nm and cryogenic temperatures," *Appl. Phys. Lett.*, vol. 101, p. 041905, 2012.

[27] A. H. Safavi-Naeini, S. Gröblacher, J. T. Hill, J. Chan, M. Aspelmeyer, and O. Painter, "Squeezed light from a silicon micromechanical resonator," *Nature*, vol. 500, pp. 185–189, 2013.

[28] H. J. Kimble, "The quantum internet," *Nature*, vol. 453, pp. 1023–1030, 2008.

[29] K. Stannigel, P. Rabl, A. S. Sørensen, P. Zoller, and M. D. Lukin, "Optomechanical Transducers for Long-Distance Quantum Communication," *Phys. Rev. Lett.*, vol. 105, p. 220501, 2010.

[30] A. H. Safavi-Naeini and O. Painter, "Proposal for an optomechanical traveling wave phonon-photon translator," *New J. Phys.*, vol. 13, p. 013017, 2011.

[31] R. J. Schoelkopf and S. M. Girvin, "Wiring up quantum systems," *Nature*, vol. 451, pp. 664–669, 2008.

[32] M. R. Vanner, M. Aspelmeyer, and M. S. Kim, "Quantum state orthogonalization and a toolset for quantum optomechanical phonon control," *Phys. Rev. Lett.*, vol. 110, p. 010504, 2013.

[33] K. Børkje, A. Nunnenkamp, and S. M. Girvin, "Proposal for entangling remote micromechanical oscillators via optical measurements," *Phys. Rev. Lett.*, vol. 107, p. 123601, 2011.

[34] C. Galland, N. Sangouard, N. Piro, N. Gisin, and T. J. Kippenberg, "Heralded single phonon preparation, storage and readout in cavity optomechanics," *Phys. Rev. Lett.*, vol. 112, p. 143602, 2014.

[35] L. M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, "Long-distance quantum communication with atomic ensembles and linear optics," *Nature*, vol. 414, pp. 413–318, 2001.

[36] N. Sangouard, C. Simon, H. de Riedmatten, and Nicolas, "Quantum repeaters based on atomic ensembles and linear optics," *Rev. Mod. Phys.*, vol. 83, p. 33, 2011.

[37] A. D. O'Connell, M. Hofheinz, M. Ansmann, R. C. Bialczak, M. Lenander, E. Lucero, M. Neeley, D. Sank, H. Wang, M. Weides, J. Wenner, J. M. Martinis, and A. N. Cleland, "Quantum ground state and single-phonon control of a mechanical resonator," *Nature*, vol. 464, pp. 697–703, 2010.

[38] M. H. Devoret and R. J. Schoelkopf, "Superconducting circuits for quantum information: an outlook," *Science*, vol. 339, pp. 1169–1174, 2013.

[39] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O'Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis, "Superconducting quantum circuits at the surface code threshold for fault tolerance," *Nature*, vol. 508, pp. 500–503, 2014.

[40] J. T. Hill, A. H. Safavi-Naeini, J. Chan, and O. Painter, "Coherent optical wavelength conversion via cavity optomechanics," *Nat. Commun.*, vol. 3, p. 1196, June 2012.

[41] 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," *Nat. Phys.*, vol. 10, pp. 321–326, 2014.

[42] Y.-D. Wang and A. A. Clerk, "Using interference for high fidelity quantum state transfer in optomechanics," 2011.

[43] L. Tian, "Adiabatic state conversion and pulse transmission in optomechanical systems," *Phys. Rev. Lett.*, vol. 108, p. 153604, 2012.

[44] J. T. Hill, *Nonlinear optics and wavelength translation via cavity optomechanics*. PhD thesis, California Institute of Technology, 2013.

[45] 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*, vol. 478, pp. 89–92, 2011.

[46] 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*, vol. 475, pp. 359–363, 2011.

[47] H. Flayac and V. Savona, "Heralded preparation and readout of entangled phonons in a photonic crystal cavity," *Phys. Rev. Lett.*, vol. 113, p. 143603, 2014.

[48] B. Julsgaard, J. Sherson, J. I. Cirac, J. Fiurášek, and E. S. Polzik, "Experimental demonstration of quantum memory for light," *Nature*, vol. 432, pp. 482–486, 2004.

[49] C. W. Chou, H. de Riedmatten, D. Felinto, S. V. Polyakov, S. J. van Enk, and H. J. Kimble, "Measurement-induced entanglement for excitation stored in remote atomic ensembles," *Nature*, vol. 438, pp. 828–832, 2005.

[50] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, "Entanglement of single-atom quantum bits at a distance," *Nature*, vol. 449, pp. 68–71, 2007.

[51] 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*, vol. 334, pp. 1253–1256, 2011.

[52] B. Pepper, R. Ghobadi, E. Jeffrey, C. Simon, and D. Bouwmeester, "Optomechanical superpositions via nested interferometry," *Phys. Rev. Lett.*, vol. 109, p. 023601, 2012.

[53] S. Pirandola, D. Vitali, P. Tombesi, and S. Lloyd, "Macroscopic entanglement by entanglement swapping," *Phys. Rev. Lett.*, vol. 97, p. 150403, 2006.

[54] S. Mancini and P. Tombesi, "High-sensitivity force measurement using entangled probes," *Europhys. Lett.*, vol. 61, pp. 8–14, 2003.

[55] J. D. Cohen, *Fiber-optic integration and efficient detection schemes for optomechanical resonators*. PhD thesis, California Instituite of Technology, 2015.

[56] J. D. Cohen, S. M. Meenehan, G. S. MacCabe, S. Gröblacher, A. H. Safavi-Naeini, F. Marsili, M. D. Shaw, and O. Painter, "Phonon counting and intensity interferometry of a nanomechanical resonator," *Nature, in press, (arXiv:1410.1047)*, 2015.

[57] T. Corbitt, C. Wipf, T. Bodiya, D. Ottaway, D. Sigg, N. Smith, S. Whitcomb, and N. Mavalvala, "Optical dilution and feedback cooling of a gram-scale oscillator to 6.9 mK," *Phys. Rev. Lett.*, vol. 99, p. 160801, 2007.

[58] O. Arcizet, P.-F. Cohadon, T. Briant, M. Pinard, and A. Heidmann, "Radiation-pressure cooling and optomechanical instability of a micromirror," *Nature*, vol. 444, pp. 71–74, 2006.

[59] S. Gigan, H. R. Böhm, M. Paternostro, F. Blaser, G. Langer, J. B. Hertzberg, K. C. Schwab, D. Bäuerle, M. Aspelmeyer, and A. Zeilinger, "Self-cooling of a micromirror by radiation pressure," *Nature*, vol. 444, pp. 67–70, 2006.

[60] 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*, vol. 452, pp. 72–75, 2008.

[61] T. P. Purdy, R. W. Peterson, P.-L. Yu, and C. A. Regal, "Cavity optomechanics with Si_3N_4 membranes at cryogenic temperatures," *New J. Phys.*, vol. 14, p. 115021, 2012.

[62] L. Ding, C. Baker, P. Senellart, A. Lemaitre, S. Ducci, G. Leo, and I. Favero, "Wavelength-sized GaAs optomechanical resonators with gigahertz frequency," *Appl. Phys. Lett.*, vol. 98, p. 113108, 2011.

[63] T. Carmon and K. Vahala, "Modal spectroscopy of optoexcited vibrations of a micron-scale on-chip resonator at greater than 1 GHz frequency," *Phys. Rev. Lett.*, vol. 98, p. 123901, 2007.

[64] A. Naik, O. Buu, M. D. LaHaye, A. D. Armour, A. A. Clerk, M. P. Blencowe, and K. C. Schwab, "Cooling a nanomechanical resonator with quantum back-action," *Nature*, vol. 443, pp. 193–196, 2006.

[65] J. D. Teufel, J. W. Harlow, C. A. Regal, and K. W. Lehnert, "Dynamical backaction of microwave fields on a nanomechanical oscillator," *Phys. Rev. Lett.*, vol. 101, p. 197203, 2008.

[66] J. D. Teufel, D. Li, M. S. Allman, K. Cicak, A. J. Sirois, J. D. Whittaker, and R. W. Simmonds, "Circuit cavity electromechanics in the strong-coupling regime," *Nature*, vol. 471, pp. 204–208, 2011.

[67] J. Suh, A. J. Weinstein, C. U. Lei, E. E. Wollman, S. K. Steinke, P. Meystre, A. A. Clerk, and K. C. Schwab, "Mechanically detecting and avoiding the quantum fluctuations of a microwave field," *Science*, vol. 344, pp. 1262–1265, 2014.

[68] K. W. Murch, K. L. Moore, S. Gupta, and D. M. Stamper-Kurn, "Observation of quantum-measurement backaction with an ultracold atomic gas," *Nature Phys.*, vol. 4, pp. 561–564, 2008.

[69] S. Combrié, A. D. Rossi, Q. V. Tran, and H. Benisty, "GaAs photonic crystal cavity with ultrahigh Q: microwatt nonlinearity at 1.55 μ m," *Opt. Lett.*, vol. 33, pp. 1908–1910, 2008.

[70] H. Sekoguchi, Y. Takahashi, T. Asano, and S. Noda, "Photonic crystal nanocavity with a Q-factor of $\sim$ 9 million," *Opt. Express*, vol. 22, pp. 916–924, 2014.

[71] E. Yablonovitch, T. Gmitter, J. P. Harbison, and R. Bhat, "Extreme selectivity in the lift-off of epitaxial GaAs films," *Appl. Phys. Lett.*, vol. 51, pp. 2222–2224, 1987.

[72] J. D. Joannopolous, S. G. Johnson, J. N. Winn, and R. D. Meade, *Photonic Crystals: Molding the Flow of Light, 2nd edition*. Princeton University Press, 2008.

[73] K. Sakoda, *Optical Properties of Photonic Crystals*. Springer, 2005.

[74] M. Tinkham, *Group Theory and Quantum Mechanics*. Dover Publications, 2003.

[75] S. G. Johnson and J. D. Joannopolous, "Block-iterative frequency-domain methods for Maxwell's equations in a planewave basis," *Opt. Express*, vol. 8, pp. 173–190, 2001.

[76] MIT Photonic Bands 1.5, http://ab-initio.mit.edu/wiki/index.php/MIT_ Photonic_ Bands.

[77] R. E. Camley, B. Djafari-Rouhani, L. Dobrzynski, and A. A. Maradudin, "Transverse elastic waves in periodically layered infinite and semi-infinite media," *Phys. Rev. B*, vol. 27, pp. 7318–7329, 1983.

[78] M. S. Kushwaha, P. Halevi, L. Dobrzynski, and B. Djafari-Rouhani, "Acoustic band structure of periodic elastic composites," *Phys. Rev. Lett.*, vol. 71, pp. 2022–2025, 1993.

[79] S. Mohammadi, A. Eftekhar, A. Khelif, and A. Adibi, "Simultaneous two-dimensional phononic and photonic band gaps in opto-mechanical crystal slabs," *Opt. Express*, vol. 18, pp. 9164–9172, 2010.

[80] A. H. Safavi-Naeini and O. Painter, "Design of optomechanical cavities and waveguides on a simultaneous bandgap phononic-photonic crystal slab," *Opt. Express*, vol. 18, pp. 14926–14943, 2010.

[81] B. A. Auld, *Acoustic Fields and Waves in Solids, vol. I and II.* Wiley Interscience, 1973.

[82] COMSOL Multiphysics 3.5, http://www.comsol.com/.

[83] C. Kittel and A. H. Mitchell, "Theory of donor and acceptor states in silicon and germanium," *Phys. Rev.*, vol. 96, pp. 1488–1493, 1954.

[84] G. Wannier, "Dynamics of band electrons in electric and magnetic fields," *Rev. Mod. Phys.*, vol. 34, pp. 645–655, 1962.

[85] M. Charbonneau-Lefort, E. Istrate, M. Allard, J. Poon, and E. Sargent, "Photonic crystal heterostructures: Waveguiding phenomena and methods of solution in an envelope function picture," *Phys. Rev. B*, vol. 65, p. 125318, 2002.

[86] E. Istrate, M. Charbonneau-Lefort, and E. Sargent, "Theory of photonic crystal heterostructures," *Phys. Rev. B*, vol. 66, p. 075121, 2002.

[87] O. Painter, K. Srinivasan, and P. E. Barclay, "Wannier-like equation for the resonant cavity modes of locally perturbed photonic crystals," *Phys. Rev. B*, vol. 68, p. 035214, 2003.

[88] K. Srinivasan and O. Painter, "Momentum space design of high-Q photonic crystal optical cavities," *Opt. Express*, vol. 10, pp. 670–684, 2002.

[89] Y. Akahane, T. Asano, B.-S. Song, and S. Noda, "High-Q photonic nanocavity in a two-dimensional photonic crystal," *Nature*, vol. 425, pp. 944–947, 2003.

[90] C. Cohen-Tannoudji, B. Diu, and F. Laloe, *Quantum Mechanics, vol. I and II.* Wiley, 1991.

[91] S. G. Johnson, M. Ibanescu, M. A. Skorobogatiy, O. Weisberg, J. D. Joannopolous, and Y. Fink, "Perturbation theory for Maxwell's equations with shifting material boundaries," *Phys. Rev. E*, vol. 65, p. 066611, 2002.

[92] A. Yariv and P. Yeh, *Optical Waves in Crystals*. Wiley-Interscience, 1983.

[93] D. K. Biegelsen, "Frequency dependence of the photoelastic coefficients of silicon," *Phys. Rev. B*, vol. 12, pp. 2427–2431, 1975.

[94] J. A. Nelder and R. Mead, "A simplex method for function minimization," *The Computer Journal*, vol. 7, pp. 308–313, 1965.

[95] J. Chan, A. H. Safavi-Naeini, J. T. Hill, S. Meenehan, and O. Painter, "Optimized optomechanical crystal cavity with acoustic radiation shield," *Appl. Phys. Lett.*, vol. 101, p. 081115, 2012.

[96] A. Akhieser, "On the absorption of sound in solids," *J. Phys. (USSR)*, vol. 1, p. 277, 1939.

[97] R. Lifshitz and M. L. Roukes, "Thermoelastic damping in micro- and nanomechanical systems," *Phys. Rev. B*, vol. 61, pp. 5600–5609, Feb 2000.

[98] M. Eichenfield, J. Chan, A. H. Safavi-Naeini, K. J. Vahala, and O. Painter, "Modeling dispersive coupling and losses of localized optical and mechanical modes in optomechanical crystals," *Opt. Express*, vol. 17, pp. 20078–20098, 2009.

[99] D. Bindel and S. Govindjee, "Elastic PMLs for resonator anchor loss simulation," *Int. J. Numer. Meth. Eng.*, vol. 64, pp. 789–818, 2005.

[100] B. S. Song, S. Noda, T. Asano, and Y. Akahane, "Ultra-high-Q photonic double-heterostructure nanocavity," *Nat. Mater.*, vol. 4, pp. 207–210, 2005.

[101] E. Kuramochi, M. Notomi, S. Mitsugi, A. Shinya, T. Tanabe, and T. Watanabe, "Ultrahigh-Q photonic crystal nanocavities realized by the local width modulation of a line defect," *Appl. Phys. Lett.*, vol. 88, p. 041112, 2006.

[102] A. Chutinan and S. Noda, "Waveguides and wavguide bends in two-dimensional photonic crystal slabs," *Phys. Rev. B*, vol. 82, pp. 4488–4492, 2000.

[103] S. G. Johnson, P. R. Villeneuve, S. Fan, and J. D. Joannopolous, "Linear waveguides in photonic-crystal slabs," *Phys. Rev. B*, vol. 62, pp. 8212–8222, 2000.

[104] M. Borselli, T. J. Johnson, and O. Painter, "Measuring the role of surface chemistry in silicon microphotonics," *App. Phys. Lett.*, vol. 88, p. 131114, 2006.

[105] A. H. Safavi-Naeini, T. P. M. Alegre, J. Chan, M. Eichenfield, M. Winger, Q. Lin, J. T. Hill, D. Chang, and O. Painter, "Electromagnetically induced transparency and slow light with optomechanics," *Nature*, vol. 472, pp. 69–73, 2011.

[106] C. P. Michael, M. Borselli, T. J. Johnson, C. Chrystal, and O. Painter, "An optical fiber-taper probe for wafer-scale microphotonic device characterization," *Opt. Express*, vol. 15, pp. 4745–4752, 2007.

[107] P. E. Barclay, K. Srinivasan, and O. Painter, "Design of photonic crystal waveguides for evanescent coupling to optical fiber tapers and integration with high-Q cavities," *J. Opt. Soc. Am. B*, vol. 20, pp. 2274–2284, Nov. 2003.

[108] P. E. Barclay, K. Srinivasan, M. Borselli, and O. Painter, "Efficient input and output fiber coupling to a photonic crystal waveguide," *Opt. Lett.*, vol. 29, pp. 697–699, 2004.

[109] S. Gröblacher, J. T. Hill, A. H. Safavi-Naeini, J. Chan, and O. Painter, "Highly efficient coupling from an optical fiber to a nanoscale silicon optomechanical cavity," *Appl. Phys. Lett.*, vol. 103, p. 181104, 2013.

[110] A. J. Manninen, J. P. Pekola, G. M. Kira, J. P. Ruutu, A. V. Babkin, H. Alles, and O. V. Lounasmaa, "First optical observations of superfluid ^{3}He," *Phys. Rev. Lett.*, vol. 69, p. 2392, 1992.

[111] A. J. R. MacDonald, G. G. Popowich, B. D. Hauer, P. H. Kim, A. Fredrick, X. Rojas, P. Doolin, and J. P. Davis, "Optical microscope and tapered fiber coupling apparatus for a dilution refrigerator," *Rev. Sci. Inst.*, vol. 86, p. 013107, 2015.

[112] D. Taillaert, W. Bogaerts, P. Bienstman, T. F. Krauss, P. V. Daele, I. Moerman, S. Verstuyft, K. D. Mesel, and R. Baets, "An out-of-plane grating coupler for efficient butt-coupling between compact planar waveguides and single-mode fibers," *IEEE J. Quantum Electron.*, vol. 38, pp. 949–955, 2002.

[113] D. Taillaert, P. Bienstman, and R. Baets, "Compact efficient broadband grating coupler for silicon-on-insulator waveguides," *Opt. Lett.*, vol. 29, pp. 2749–2751, 2004.

[114] D. Taillaert, F. V. Laere, M. Ayre, W. Bogaerts, D. V. Thourhout, P. Bienstman, and R. Baets, "Grating couplers for coupling between optical fibers and nanophotonic waveguides," *Jpn. J. of Appl. Phys.*, vol. 45, pp. 6071–6077, 2006.

[115] D. Vermeulen, S. Selvaraja, P. Verheyen, G. LePage, W. Bogaerts, P. Absil, D. V. Thourhout, and G. Roelkens, "High-efficiency fiber-to-chip grating couplers realized using an advanced cmos-compatible silicon-on-insulator platform," *Opt. Express*, vol. 18, pp. 18278–18283, 2010.

[116] Y. Tang, Z. Wang, L. Wosinski, U. Westergren, and S. He, "Highly-efficient nonuniform grating coupler for silicon-on-insulator nanophotonic circuits," *Opt. Lett.*, vol. 35, pp. 1290–1292, 2010.

[117] F. V. Laere, G. Roelkens, M. Ayre, J. Schrauwen, D. Taillaert, D. V. Thourhout, T. F. Krauss, and R. Baets, "Compact and highly efficient grating couplers between optical fiber and nanophotonic waveguides," *J. Lightwave Technol.*, vol. 25, pp. 151–156, 2007.

[118] Z. Cheng, X. Chen, C. Y. Wong, K. Xu, and H. K. Tsang, "Apodized focusing subwavelength grating couplers for suspended membrane waveguides," *Appl. Phys. Lett.*, vol. 101, p. 101104, 2012.

[119] O. Mitomi, K. Kasaya, and H. Miyazawa, "Design of a single-mode tapered waveguide for low-loss chip-to-fiber coupling," *IEEE J. Quantum Elect.*, vol. 30, pp. 1787–1793, 1994.

[120] V. R. Almeida, R. R. Panepucci, and M. Lipson, "Nanotaper for compact mode conversion," *Opt. Lett.*, vol. 28, pp. 1302–1304, 2003.

[121] S. McNab, N. Moll, and Y. Vlasov, "Ultra-low loss photonic integrated circuit with membrane-type photonic crystal waveguides," *Opt. Express*, vol. 11, pp. 2927–2939, 2003.

[122] L. Chen, C. R. Doerr, Y.-K. Chen, and T.-Y. Liow, "Low-loss and broadband cantilever couplers between standard cleaved fibers and high-index-contrast Si_3N_4 or Si waveguides," *IEEE Photonic. Tech. Lett.*, vol. 22, pp. 1744–1746, 2010.

[123] A. W. Snyder and J. D. Love, *Optical Waveguide Theory.* Springer, 1983.

[124] D. Marcuse, *Theory of Dielectric Optical Waveguides.* Academic Press, 1991.

[125] C. Zener, "Non-adiabatic crossing of energy levels," *Proc. R. Soc. A*, vol. 137, pp. 696–702, 1932.

[126] A. Messiah, *Quantum Mechanics: Vol. II.* Wiley, 1976.

[127] D. G. Dalgoutte, R. B. Smith, G. Achutaramayya, and J. H. Harris, "Externally mounted fibers for integrated optics interconnections," *Appl. Opt.*, vol. 14, pp. 1860–1865, 1975.

[128] A. W. Snyder, "Mode propagation in a non-uniform cylindrical medium," *I.E.E.E. Trans. Microwave Theory Tech.*, vol. 19, pp. 402–403, 1971.

[129] S. G. Johnson, P. Bienstman, M. A. Skorobogatiy, M. Ibanescu, E. Lidorikis, and J. D. Joannopolous, "Adiabatic theorem and continuous coupled-mode theory for efficient taper transitions in photonic crystals," *Phys. Rev. E*, vol. 66, p. 066608, 2002.

[130] J. Chan, M. Eichenfield, R. Camacho, and O. Painter, "Optical and mechanical design of a "zipper" photonic crystal optomechanica cavity," *Opt. Express*, vol. 17, pp. 3802–3817, 2009.

[131] M. Eichenfield, R. Camacho, J. Chan, K. J. Vahala, and O. Painter, "A picogram- and nanometre-scale photonic-crystal optomechanical cavity," *Nature*, vol. 459, pp. 550–555, 2009.

[132] V. Braginsky and A. Manukin, *Measurement of weak forces in physics experiments*. Univ. of Chicago Press, 1977.

[133] A. G. Krause, M. Winger, T. D. Blasius, Q. Lin, and O. Painter, "A high-resolution microchip optomechanical accelerometer," *Nature Photonics*, vol. 6, pp. 768–772, 2012.

[134] D. Kleckner and D. Bouwmeester, "Sub-kelvin optical cooling of a micromechanical resonator," *Nature*, vol. 444, pp. 75–78, 2006.

[135] C. Genes, D. Vitali, P. Tombesi, S. Gigan, and M. Aspelmeyer, "Ground-state cooling of a micromechanical oscillator: comparing cold damping and cavity-assisted cooling schemes," *Phys. Rev. A*, vol. 77, p. 033804, 2008.

[136] J. D. Cohen, S. M. Meenehan, and O. Painter, "Optical coupling to nanoscale optomechanical cavities for near quantum-limited motion transduction," *Opt. Express*, vol. 21, pp. 11227–11236, 2013.

[137] Datasheet of Corning SMF-28 optical fiber.

[138] Lumerical Solutions Inc. http://www.lumerical.com/tcad-products/fdtd/.

[139] E. Desurvire, D. Bayart, B. Desthieux, and S. Bigo, *Erbium-Doped Fiber Amplifiers, Device and System Developments*. Wiley-Interscience, 2002.

[140] W. B. Gardner, "Microbending loss in optical fibers," *Bell Syst. Tech. J.*, vol. 54(2), pp. 457–465, 1975.

[141] Y. T. Yang, C. Callegari, X. L. Feng, and M. L. Roukes, "Surface adsorbate fluctuations and noise in nanoelectromechanical systems," *Nano Lett.*, vol. 11, pp. 1753–1759, 2011.

[142] J. Gao, *The Physics of Superconducting Microwave Resonators*. PhD thesis, California Institute of Technology, 2008.

[143] J. J. Olivero and R. L. Longbothum, "Empirical fits to a Voigt line width: A brief review," *J. Quant. Spectrosc. Radiat. Transfer*, vol. 17, pp. 233–236, 1977.

[144] M. G. Holland, "Analysis of lattice thermal conductivity," *Phys. Rev.*, vol. 132, p. 2461, 1963.

[145] P. S. Zyryanov and G. G. Taluts, "On the theory of sound absorption in solids," *JETP*, vol. 22, pp. 1326–1330, June 1966.

[146] W. A. Phillips, "Tunneling states in amorphous solids," *Journal of Low Temperature Physics*, vol. 7, pp. 351–360, 1972.

[147] P. W. Anderson, B. I. Halperin, and C. M. Varma, "Anomalous low-temperature thermal properties of glasses and spin glasses," *Philos. Mag.*, vol. 25, pp. 1–9, 1972.

[148] M. Borselli, T. J. Johnson, C. P. Michael, M. D. Henry, and O. Painter, "Surface encapsulation for low-loss silicon photonics," *Appl. Phys. Lett.*, vol. 91, p. 131117, 2007.

[149] G. P. Srivastava, *The Physics of Phonons*. Taylor and Francis Group, 1990.

[150] A. Stesmans, "Passivation of P_{b0} and P_{b1} interface defects in thermal (100) Si/SiO_2 with molecular hydrogen," *App. Phys. Lett.*, vol. 68, 1996.

[151] R. Rao, J. Menendez, C. D. Poweleit, and A. M. Rao, "Anharmonic phonon lifetimes in carbon nanotubes: Evidence for a one-dimensional phonon decay bottleneck," *Phys. Rev. Lett.*, vol. 99, p. 047403, 2007.

[152] M. Steiner, M. Freitag, V. Perebeinos, J. C. Tsang, J. P. Small, M. Kinoshita, D. Yuan, J. Liu, and P. Avouris, "Phonon populations and electrical power dissipation in carbon nanotube transistors," *Nat. Nanotechnol.*, vol. 4, pp. 320–324, 2009.

[153] D. Singh, J. Y. Murthy, and T. S. Fisher, "Spectral phonon conduction and dominant scattering pathways in graphene," *J. Appl. Phys.*, vol. 110, p. 094312, 2011.

[154] R. Chen, A. I. Hochbaum, P. Murphy, J. Moore, P. Yang, and A. Majumdar, "Thermal conductance of thin silicon nanowires," *Phys. Rev. Lett.*, vol. 101, p. 105501, 2008.

[155] K. Schwab, E. A. Henriksen, J. M. Worlock, and M. L. Roukes, "Measurement of the quantum of thermal conductance," *Nature*, vol. 404, pp. 974–977, 2000.

[156] N. Mingo, "Calculation of Si nanowire thermal conductivity using complete phonon dispersion relations," *Phys. Rev. B*, vol. 68, p. 113308, 2003.

[157] J. B. Hertzberg, M. Aksit, O. O. Otelaja, D. A. Stewart, and R. D. Robinson, "Direct measurements of surface scattering in Si nanosheets using a microscale phonon spectrometer: implications for Casimir-limit predicted by Ziman theory," *Nano Lett.*, vol. 14, pp. 403–415, 2014.

[158] W. A. Phillips, "Two-level states in glasses," *Rep. Prog. Phys.*, vol. 50, pp. 1657–1708, 1987.

[159] W. A. Phillips, "Comment on "Two-level systems observed in the mechanical properties of single-crystal silicon at low temperatures"," *Phys. Rev. Lett.*, vol. 61, p. 2632, 1988.

[160] R. N. Kleiman, G. Agnolet, and D. J. Bishop, "Two-level systems observed in the mechanical properties of single-crystal silicon at low temperatures," *Phys. Rev. Lett.*, vol. 59, pp. 2079–2082, 1987.

[161] P. Mohanty, D. A. Harrington, K. L. Ekinci, Y. T. Yang, M. J. Murphy, and M. L. Roukes, "Intrinsic dissipation in high-frequency micromechanical resonators," *Phys. Rev. B*, vol. 66, p. 085416, 2002.

[162] A. H. Safavi-Naeini, J. T. Hill, S. Meenehan, J. Chan, S. Gröblacher, and O. Painter, "Two-dimensional phononic-photonic bandgap optomechanical crystal cavity," *Phys. Rev. Lett.*, vol. 112, p. 153603, Apr. 2014.

[163] M. Schmidt, M. Ludwig, and F. Marquardt, "Optomechanical circuits for nanomechanical continuous variable quantum state processing," *New J. Phys.*, vol. 14, p. 125005, Dec. 2012.

[164] A. Tomadin, S. Diehl, M. D. Lukin, P. Rabl, and P. Zoller, "Reservoir engineering and dynamical phase transitions in optomechanical arrays," *Phys. Rev. A*, vol. 86, p. 033821, Sept. 2012.

[165] M. Ludwig and F. Marquardt, "Quantum many-body dynamics in optomechanical arrays," *Phys. Rev. Lett.*, vol. 111, p. 073603, Aug. 2013.

[166] M. Schmidt, V. Peano, and F. Marquardt, "Optomechanical metamaterials: Dirac polaritons, gauge fields, and instabilities," *arXiv:1311.7095*, 2013.

[167] F. Diedrich, J. C. Bergquist, W. M. Itano, and D. J. Wineland, "Laser cooling to the zero-point energy of motion," *Phys. Rev. Lett.*, vol. 62, pp. 403–406, Jan. 1989.

[168] F. Y. Khalili, H. Miao, H. Yang, A. H. Safavi-Naeini, O. Painter, and Y. Chen, "Quantum back-action in measurements of zero-point mechanical oscillations," *Phys. Rev. A*, vol. 86, p. 033840, 2012.

[169] M. Davanco, J. Chan, A. H. Safavi-Naeini, O. Painter, and K. Srinivasan, "Slot-mode-coupled optomechanical crystals," *Opt. Express*, vol. 20, pp. 24394–24410, 2012.

[170] R. A. Norte, *Nanofabrication for on-chip optical levitation, atom-trapping, and superconducting quantum circuits*. PhD thesis, California Institute of Technology, 2015.

[171] A. Pitani, J. M. Fink, A. H. Safavi-Naeini, J. T. Hill, C. U. Lei, A. Tredicucci, and O. Painter, "Strong opto-electro-mechanical coupling in a silicon photonic crystal cavity," *Opt. Express*, vol. 23, pp. 3196–3208, 2015.

[172] S. G. Hofer, W. Wieczorek, M. Aspelmeyer, and K. Hammerer, "Quantum entanglement and teleportation in pulsed cavity optomechanics," *Phys. Rev. A*, vol. 84, p. 052327, 2011.

[173] M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[174] S. J. M. Habraken, K. Stannigel, M. D. Lukin, P. Zoller, and P. Rabl, "Continuous mode cooling and phonon routers for phononic quantum networks," *New J. Phys.*, vol. 14, p. 115004, 2012.

[175] T. Klitsner, J. E. VanCleve, H. E. Fischer, and R. O. Pohl, "Phonon radiative heat transfer and surface scattering," *Phys. Rev. B*, vol. 38, pp. 7576–7594, 1988.

[176] A. H. Safavi-Naeini, J. T. Hill, S. Meenehan, J. Chan, S. Gröblacher, and O. Painter, "Two-dimensional phononic-photonic band gap optomechanical crystal cavity," *Phys. Rev. Lett.*, vol. 112, p. 153603, 2014.

[177] A. Yariv and P. Yeh, *Photonics: Optical Electronics in Modern Communications, 6th edition*. Oxford University Press, 2007.

[178] 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," *Nat. Photon.*, vol. 7, pp. 210–214, 2013.