

# Nonlinear Optics and Wavelength Translation Via Cavity-Optomechanics

Thesis by  
Jeffrey Thomas Hill

In Partial Fulfillment of the Requirements  
for the Degree of  
Doctor of Philosophy



California Institute of Technology  
Pasadena, California

2013  
(Defended May 22, 2013)



© 2013

Jeffrey Thomas Hill

All Rights Reserved

# Acknowledgments

To Oskar Painter who made this all possible. You provided all the tools and guidance, and always let us do what we found interesting and fun. I can't imagine a better setting to explore an ever increasingly interesting field.

To my friends and surrogate family that is the Painter group. This would not have been nearly as enjoyable without you, and I can't believe how lucky we've been to work together. To Amir Safavi-Naeini and Alex Krause, you've been with me since the beginning and provided constant friendship, support, and intellectual stimulation—more debt and gratitude than these words can convey. How I got so lucky I will never know. To Jasper Chan, how it took us so long to start working together and hanging out I don't know, but it was as if you always belonged. To Simon Gröblacher for bringing your cool, common sense approach and knowledge at just the right time and making our circle of friends that much bigger and better. To everyone in the group: Justin Cohen, Seán Meenehan, Tim Blasius, Taofiq Paraiso, Johannes Fink, and Alessandro Pitanti. And of course: HI Richard! (Norte). Barry Baker and Michelle Aldecua for making our lives so much easier. A special thanks to Qiang Lin who took me under his wing and with incredible patience and knowledge taught me many things. Former group members: Raviv Perahia, Ryan Camacho, Matt Eichenfield, Jessie Rosenberg, Thiago Alegre, and Martin Winger. Outside of the group: Ari Weinstein you are as much a part of the family as anyone. Dan Turner-Evans you have been an amazing roommate, friend, and constant inspiration on the correct way to keep a balanced, sane lifestyle at Caltech.

To my friends back home who I've failed to keep in touch with as much as you deserve: Mike Gevaert, Nino Zahirovic, and August Claxton. I couldn't have made it here without you.

To Camille Boucher-Veronneau, I wouldn't be here if it wasn't for you, literally. When I came to Caltech I was told you can have work and one thing. You are the one and words can only fall short.

To my parents. You've stood behind me and encouraged me as long as I can remember. To my sister Julie who continued to keep in touch with me despite my terrible ability to respond and always made sure I never forgot an important family event. To my brothers Ryan and Adam. Friend and brother-in-law Derek, always having an open house, beer, and camaraderie. We all have shaped each other and I wouldn't be here if it wasn't for the influence you all had on me.

# Abstract

The field of cavity-optomechanics explores the interaction of light with sound in an ever increasing array of devices. This interaction allows the mechanical system to be both sensed and controlled by the optical system, opening up a wide variety of experiments including the cooling of the mechanical resonator to its quantum mechanical ground state and the squeezing of the optical field upon interaction with the mechanical resonator, to name two.

In this work we explore two very different systems with different types of optomechanical coupling. The first system consists of two microdisk optical resonators stacked on top of each other and separated by a very small slot. The interaction of the disks causes their optical resonance frequencies to be extremely sensitive to the gap between the disks. By careful control of the gap between the disks, the optomechanical coupling can be made to be quadratic to first order which is uncommon in optomechanical systems. With this quadratic coupling the light field is now sensitive to the energy of the mechanical resonator and can directly control the potential energy trapping the mechanical motion. This ability to directly control the spring constant without modifying the energy of the mechanical system, unlike in linear optomechanical coupling, is explored.

Next, the bulk of this thesis deals with a high mechanical frequency optomechanical crystal which is used to coherently convert photons between different frequencies. This is accomplished via the engineered linear optomechanical coupling in these devices. Both classical and quantum systems utilize the interaction of light and matter across a wide range of energies. These systems are often not naturally compatible with one another and require a means of converting photons of dissimilar wavelengths to combine and exploit their different strengths. Here we theoretically propose and experimentally demonstrate coherent wavelength conversion of optical photons using photon-phonon translation in a cavity-optomechanical system. For an engineered silicon optomechanical crystal nanocavity supporting a 4 GHz localized phonon mode, optical signals in a 1.5 MHz bandwidth are coherently converted over a 11.2 THz frequency span between one cavity mode at wavelength 1460 nm and a second cavity mode at 1545 nm with a 93% internal (2% external) peak efficiency. The thermal and quantum limiting noise involved in the conversion process is also analyzed and, in terms of an equivalent photon number signal level, are found to correspond to an internal noise level of only 6 and  $4 \times 10^{-3}$  quanta, respectively.

We begin by developing the requisite theoretical background to describe the system. A significant amount of time is then spent describing the fabrication of these silicon nanobeams, with an emphasis on understanding the specifics and motivation. The experimental demonstration of wavelength conversion is then described and analyzed. It is determined that the method of getting photons into the cavity and collected from the cavity is a fundamental limiting factor in the overall efficiency. Finally, a new coupling scheme is designed, fabricated, and tested that provides a means of coupling greater than 90% of photons into and out of the cavity, addressing one of the largest obstacles with the initial wavelength conversion experiment.

# Contents

|                                                                      |             |
|----------------------------------------------------------------------|-------------|
| <b>Acknowledgments</b>                                               | <b>iv</b>   |
| <b>Abstract</b>                                                      | <b>v</b>    |
| <b>List of Figures</b>                                               | <b>xi</b>   |
| <b>List of Tables</b>                                                | <b>xiii</b> |
| <b>Preface</b>                                                       | <b>xiv</b>  |
| <b>1 Quadratic Coupling With Double Disks</b>                        | <b>1</b>    |
| 1.1 Quadratic Optomechanical Coupling . . . . .                      | 1           |
| 1.2 Fabrication . . . . .                                            | 4           |
| 1.2.1 Material Growth . . . . .                                      | 4           |
| 1.2.1.1 PECVD of Si and SiO <sub>2</sub> . . . . .                   | 4           |
| 1.2.1.2 Anneal . . . . .                                             | 5           |
| 1.2.2 Electron Beam Lithography . . . . .                            | 6           |
| 1.2.3 Etching . . . . .                                              | 6           |
| 1.2.4 Release . . . . .                                              | 8           |
| 1.2.5 Independent Engineering of the Mechanical Properties . . . . . | 9           |
| 1.3 Measurement . . . . .                                            | 9           |
| 1.3.1 Tuning to an Anticrossing . . . . .                            | 10          |
| 1.3.1.1 Laser Tuning . . . . .                                       | 10          |
| 1.3.1.2 Thermal-Mechanical Tuning . . . . .                          | 14          |
| 1.3.2 Linear Transduction of Mechanical Motion . . . . .             | 15          |
| 1.3.3 Suppression of Linear Coupling . . . . .                       | 15          |
| 1.4 Results . . . . .                                                | 17          |
| <b>2 Theory of Sideband-Resolved Optomechanical Systems</b>          | <b>21</b>   |
| 2.1 Model of an Optomechanical System . . . . .                      | 21          |

|          |                                                                                                              |           |
|----------|--------------------------------------------------------------------------------------------------------------|-----------|
| 2.2      | Equations of Motion for Linearly Coupled Optomechanical Systems . . . . .                                    | 23        |
| 2.2.1    | Linearization Around a Strong Coherent Field . . . . .                                                       | 24        |
| 2.2.2    | Optomechanical Modification of the Mechanical Frequency and<br>Damping–Weak Coupling Approximation . . . . . | 25        |
| 2.3      | Theory of Optomechanical Wavelength Conversion . . . . .                                                     | 27        |
| 2.3.1    | Efficiency, Bandwidth, and Noise . . . . .                                                                   | 32        |
| <b>3</b> | <b>Fabrication of Nanobeam Optomechanical Crystals</b>                                                       | <b>36</b> |
| 3.1      | Design . . . . .                                                                                             | 37        |
| 3.1.1    | Design for Fabrication . . . . .                                                                             | 38        |
| 3.2      | Electron Beam Lithography . . . . .                                                                          | 40        |
| 3.2.1    | Parameter Sweeps . . . . .                                                                                   | 40        |
| 3.2.2    | Design Compensation . . . . .                                                                                | 41        |
| 3.2.3    | File Fracturing . . . . .                                                                                    | 43        |
| 3.2.4    | Spinning and Baking of Electron Beam Resist . . . . .                                                        | 48        |
| 3.2.5    | Electron Beam Writing . . . . .                                                                              | 49        |
| 3.2.6    | Develop and Rinse . . . . .                                                                                  | 50        |
| 3.3      | Etching . . . . .                                                                                            | 51        |
| 3.3.1    | ICP RIE Overview . . . . .                                                                                   | 51        |
| 3.3.1.1  | SF <sub>6</sub> - Sulfur Hexafluoride . . . . .                                                              | 52        |
| 3.3.1.2  | C <sub>4</sub> F <sub>8</sub> - Octafluorocyclobutane . . . . .                                              | 53        |
| 3.3.1.3  | ICP Power . . . . .                                                                                          | 53        |
| 3.3.1.4  | RF Power and DC Bias . . . . .                                                                               | 53        |
| 3.3.1.5  | Chamber Pressure . . . . .                                                                                   | 54        |
| 3.3.1.6  | Helium Backing . . . . .                                                                                     | 54        |
| 3.3.1.7  | Etch Table Temperature . . . . .                                                                             | 55        |
| 3.3.1.8  | Matching Capacitors (AMU1 and AMU2) . . . . .                                                                | 55        |
| 3.3.1.9  | Etch Time . . . . .                                                                                          | 56        |
| 3.3.2    | Securing the Chip to the Carrier Wafer for Etching . . . . .                                                 | 57        |
| 3.3.3    | Silicon Etching and Optimization . . . . .                                                                   | 57        |
| 3.3.3.1  | Original Si Etch . . . . .                                                                                   | 57        |
| 3.3.3.2  | Si Etch Reoptimization Steps . . . . .                                                                       | 59        |
| 3.3.3.3  | Key Points and Lessons Learned . . . . .                                                                     | 64        |
| 3.3.4    | Silicon Nitride Etching and Optimization . . . . .                                                           | 65        |
| 3.3.4.1  | Key Points and Lessons Learned . . . . .                                                                     | 68        |
| 3.4      | Wet Chemistry . . . . .                                                                                      | 68        |

|          |                                                                                    |            |
|----------|------------------------------------------------------------------------------------|------------|
| 3.4.1    | Standard Solvent Clean . . . . .                                                   | 68         |
| 3.4.2    | Resist Strip . . . . .                                                             | 68         |
| 3.4.2.1  | Removing the Etched Chip from the Carrier Wafer . . . . .                          | 71         |
| 3.4.2.2  | TCE Bath or ZDMAC Clean . . . . .                                                  | 72         |
| 3.4.2.3  | Piranha . . . . .                                                                  | 73         |
| 3.4.3    | HF Undercut . . . . .                                                              | 74         |
| 3.4.4    | Surface Preparation . . . . .                                                      | 75         |
| <b>4</b> | <b>Experimental Demonstration of Coherent Optomechanical Wavelength Conversion</b> | <b>76</b>  |
| 4.1      | Optomechanical Wavelength Conversion . . . . .                                     | 77         |
| 4.2      | Experimental Setup and Device Characterization . . . . .                           | 80         |
| 4.2.1    | Experimental Setup . . . . .                                                       | 80         |
| 4.2.2    | Device Characterization . . . . .                                                  | 82         |
| 4.2.3    | Generation and Conversion of the Input Signal . . . . .                            | 85         |
| 4.2.4    | Calibration of the Input Signal . . . . .                                          | 86         |
| 4.2.5    | Power-Dependent Optical Cavity Loss . . . . .                                      | 87         |
| 4.3      | Experimental Results . . . . .                                                     | 88         |
| 4.3.1    | Conversion Efficiency . . . . .                                                    | 89         |
| 4.3.2    | Bandwidth . . . . .                                                                | 90         |
| 4.3.3    | Noise Considerations . . . . .                                                     | 91         |
| 4.4      | Towards Single-Photon Wavelength Conversion . . . . .                              | 92         |
| <b>5</b> | <b>Adiabatically Tapered Waveguide Coupling</b>                                    | <b>94</b>  |
| 5.1      | Motivation . . . . .                                                               | 94         |
| 5.2      | Adiabatically Tapered Waveguide-Fiber Coupling . . . . .                           | 95         |
| 5.3      | Single-Sided Coupling of the Waveguide to the Device . . . . .                     | 97         |
| 5.4      | Experimental Results . . . . .                                                     | 99         |
| 5.5      | Discussion . . . . .                                                               | 102        |
| <b>A</b> | <b>Mathematical Tools and Definitions</b>                                          | <b>103</b> |
| A.1      | Fourier Transforms and Spectral Densities . . . . .                                | 103        |
| A.2      | Noise Correlations . . . . .                                                       | 103        |
| A.3      | Some Useful Commutator Relationships . . . . .                                     | 104        |
| A.4      | The Mechanical Mode: A Bosonic Harmonic Oscillator . . . . .                       | 104        |

|                                                            |            |
|------------------------------------------------------------|------------|
| <b>B Mathematical Description of an Optical Cavity</b>     | <b>106</b> |
| B.1 A Single Optical Cavity . . . . .                      | 106        |
| B.2 Input/Output Coupling and Dissipation . . . . .        | 107        |
| B.2.1 Single and Double-Sided Coupling . . . . .           | 107        |
| B.2.2 Solution to the Cavity Equation of Motion . . . . .  | 108        |
| <b>C Power Spectral Density of the Mechanical Mode</b>     | <b>112</b> |
| C.1 Mechanical Mode Coupled to Thermal Bath . . . . .      | 112        |
| C.2 Mechanical Mode with Optomechanical Coupling . . . . . | 114        |
| C.2.1 Red-Detuned Pump ( $\Delta = \omega_m$ ) . . . . .   | 115        |
| C.2.2 On-Resonance Pump ( $\Delta = 0$ ) . . . . .         | 115        |
| C.2.3 Blue-Detuned Pump ( $\Delta = -\omega_m$ ) . . . . . | 116        |
| <b>D Publications</b>                                      | <b>117</b> |
| <b>Bibliography</b>                                        | <b>119</b> |



# List of Figures

|      |                                                                                             |    |
|------|---------------------------------------------------------------------------------------------|----|
| 1.1  | Tuning a double disk . . . . .                                                              | 2  |
| 1.2  | Fabrication of double disks . . . . .                                                       | 4  |
| 1.3  | Examples of the double-spun resist . . . . .                                                | 6  |
| 1.4  | Examples of good and bad double disk etches . . . . .                                       | 7  |
| 1.5  | Experimental setup for quadratic coupling . . . . .                                         | 11 |
| 1.6  | Tuning the double disks into an anticrossing . . . . .                                      | 12 |
| 1.7  | Extraction of quadratic coupling from anticrossing used . . . . .                           | 14 |
| 1.8  | Suppression of the linear coupling . . . . .                                                | 16 |
| 1.9  | Intracavity photon number dependence of the static spring effect . . . . .                  | 18 |
| 1.10 | Detuning from the quadratic coupling point dependence of the static spring effect . . . . . | 19 |
| 2.1  | Canonical optomechanical system . . . . .                                                   | 22 |
| 2.2  | System model for optomechanical wavelength conversion . . . . .                             | 27 |
| 2.3  | Pumping scheme for wavelength conversion . . . . .                                          | 28 |
| 3.1  | Overview of silicon processing for the nanobeams . . . . .                                  | 38 |
| 3.2  | Comparison of rectangular vs. elliptical holes . . . . .                                    | 39 |
| 3.3  | Comparison of a small etched slot and pulled in beams . . . . .                             | 40 |
| 3.4  | Unit cell of the nanobeam and phononic shield . . . . .                                     | 42 |
| 3.5  | Nanobeam fracture flow . . . . .                                                            | 44 |
| 3.6  | Fracture test patterns and results . . . . .                                                | 46 |
| 3.7  | Example .dxf and .gpf files . . . . .                                                       | 47 |
| 3.8  | Example of under and overdosing the resist . . . . .                                        | 50 |
| 3.9  | Schematic of a generic ICP/RIE etch system . . . . .                                        | 52 |
| 3.10 | Etch log template . . . . .                                                                 | 56 |
| 3.11 | Development of original Si etch . . . . .                                                   | 57 |
| 3.12 | Original silicon etch . . . . .                                                             | 58 |
| 3.13 | Progression of the Si etch reoptimization . . . . .                                         | 60 |
| 3.14 | Etch test pattern . . . . .                                                                 | 64 |

|      |                                                                                     |     |
|------|-------------------------------------------------------------------------------------|-----|
| 3.15 | Original silicon nitride etch . . . . .                                             | 65  |
| 3.16 | Progression of the silicon nitride etch optimization . . . . .                      | 67  |
| 3.17 | Progression of the silicon nitride etch optimization with 600 W ICP power . . . . . | 70  |
| 3.18 | Comparison of resist stripping . . . . .                                            | 71  |
| 4.1  | System model . . . . .                                                              | 78  |
| 4.2  | Schematic of the experimental setup . . . . .                                       | 81  |
| 4.3  | Optomechanical crystal nanobeam . . . . .                                           | 83  |
| 4.4  | Optical spectroscopy of the device . . . . .                                        | 84  |
| 4.5  | Characterization of the optomechanical coupling . . . . .                           | 85  |
| 4.6  | Calibration tones . . . . .                                                         | 86  |
| 4.7  | Effect of large intracavity photon population . . . . .                             | 88  |
| 4.8  | Wavelength conversion efficiency . . . . .                                          | 89  |
| 4.9  | Output tones for input tone detunings . . . . .                                     | 90  |
| 4.10 | Bandwidth of the wavelength conversion . . . . .                                    | 90  |
| 4.11 | Converted tone on thermal noise . . . . .                                           | 91  |
| 4.12 | Output noise . . . . .                                                              | 92  |
| 5.1  | Adiabatically Tapered Waveguide-Optical Fiber Coupling . . . . .                    | 97  |
| 5.2  | Waveguide-Device Coupling . . . . .                                                 | 98  |
| 5.3  | Waveguide Coupled Zipper Nanobeams and Single Nanobeam . . . . .                    | 99  |
| 5.4  | Experimental Setup . . . . .                                                        | 100 |
| 5.5  | Experimental Results . . . . .                                                      | 101 |
| B.1  | Schematic of single and double-sided coupling . . . . .                             | 108 |
| B.2  | Transmission depth on resonance for single and double-sided coupling . . . . .      | 110 |

# List of Tables

|     |                                                              |    |
|-----|--------------------------------------------------------------|----|
| 1.1 | Double disk deposition recipe . . . . .                      | 5  |
| 1.2 | Double disk etch parameters . . . . .                        | 8  |
| 1.3 | Double disk SF <sub>6</sub> dry release parameters . . . . . | 9  |
| 3.1 | Overview of fabrication process . . . . .                    | 37 |
| 3.2 | Design compensation for silicon nanobeams . . . . .          | 42 |
| 3.3 | Design compensation for silicon nitride nanobeams . . . . .  | 43 |
| 3.4 | Layout Beamer fracturing options . . . . .                   | 45 |
| 3.5 | Si etch re-optimization parameters . . . . .                 | 62 |
| 3.6 | SiN etch optimization parameters . . . . .                   | 69 |

# Preface

I started at Caltech in the Fall of 2008 with a background in computer engineering and experimental low-temperature physics, and only a vague idea of wanting to learn interesting things and do exciting research. Knowing nothing about photonics or optics, I was given a generous offer by Oskar to begin working in his lab and explore the research being done. First year Applied Physics, however, had a different plan for me which included many classes and a lot of homework. Having not taken a class in several years, this experience is perhaps best equated to treading water, in the ocean, far from shore, in a storm—not drowning, but never really making it closer to shore. As the year progressed I remembered how to take classes and managed to start to explore the research of the Painter group. Little did I know how lucky my timing was as research in optomechanics was just beginning in the group and would turn out to be a very exciting research area where we could make significant contributions.

I first started working with a postdoc Ryan Camacho and shortly after a new graduate student Alex Krause. Our goal was to use the new optomechanical structures developed in the group, the so-called zipper cavities, to see the effects of radiation pressure shot-noise on the mechanical motion of the device. It was during these months that I was introduced to many of the skills that would form the foundation of my graduate career. Our progress was at times slow, and our goals lofty, but Ryan was a patient teacher who trained us well. Around this time the UCSB group demonstrated a macroscopic mechanical resonator in the ground state for the first time—exciting the field and setting the stage for the rapid pace of results to follow.

It became clear our goal was more ambitious than previously thought and I moved to work on a different project with low-hanging fruit. This new project involved me working with another postdoc, Qiang Lin, and on a new optomechanical system, the double disks. This project involved attempting to create an extremely strong quadratic optomechanical coupling, fundamentally different than the linear optomechanical coupling of most devices studied. With this project the devices were already designed, fabricated, and well understood, and had been used quite successfully in our group before. How to create the quadratic coupling was also understood and the experiments we wished to do with them already planned out—it seemed like all that was left was the mechanics of taking the data. Things of course never go exactly to plan and take longer than you think. Many problems

emerged, many details needed to be considered further. The experience lasted almost a year, and my time on this project with Qiang was incredibly exciting, educational, and fulfilling. We worked very hard and very long, and managed to explore these system quite thoroughly. Qiang was extremely knowledgeable, tirelessly patient, and a great teacher, experimentalist and theorist. We took data right up until the day Qiang left to start his faculty position.

At this point I began working with Amir, who had joined the lab the same year as me. We quickly became a good team and worked on several interesting projects simultaneously. At the same time Jasper and Thiago had been working on the ground state cooling of a high-frequency optomechanical nanobeam which would be a great result if possible and was also being attempted by many other groups around the world. Amir and I joined Jasper and Thiago and our super-subgroup was born. After months of intense work we managed to cool the 4 GHz mode of the nanobeam to less than 1 phonon—the first result involving laser cooling to the ground state. We quickly followed up this work with an experiment observing the zero-point motion of the mechanical resonator. Next we did the wavelength conversion that is the main thrust of this thesis.

During this time Thiago took up his faculty position in Brazil, and we gained Simon, and continued to do exciting research. The last experiment before writing this thesis involved the squeezing of light from an optomechanical resonator. This squeezing is a result of the driving of the mechanical resonator by the radiation pressure shot-noise, which appropriately brings us full circle to the first experiment I was involved in when joining the group. Better late than never.

This thesis first details some of the work on the quadratic coupling with the double disk resonators, and then focuses on the wavelength conversion work done with the high-frequency nanobeams, and ends with an improved coupling scheme providing a means to drastically boost the wavelength conversion efficiency.

## Chapter 1

# Quadratic Coupling With Double Disks

### 1.1 Quadratic Optomechanical Coupling

As will be shown in section 2.1, one of the easiest ways to understand the optomechanical interaction is to consider an optical cavity with one movable mirror, making its resonance frequency dependent on the position of the mirror. Taylor expanding the frequency gives,

$$\omega_o(\hat{x}) = \omega_o + \frac{d\omega_o(\hat{x})}{d\hat{x}}\hat{x} + \frac{1}{2!} \frac{d^2\omega_o(\hat{x})}{d\hat{x}^2} + \dots \quad (1.1)$$

with  $\omega_o$  representing the unperturbed optical frequency. In almost all systems the linear term dominates and the optomechanical coupling is between the optical mode and the position of the mechanical mode. This is the type of coupling responsible for the effects described in the rest of this thesis. In this section we will focus on the situation that occurs when the linear optomechanical coupling vanishes, that is  $g_{OM} = d\omega_o(\hat{x})/d\hat{x} \rightarrow 0$  and the dominate coupling is the quadratic coupling,

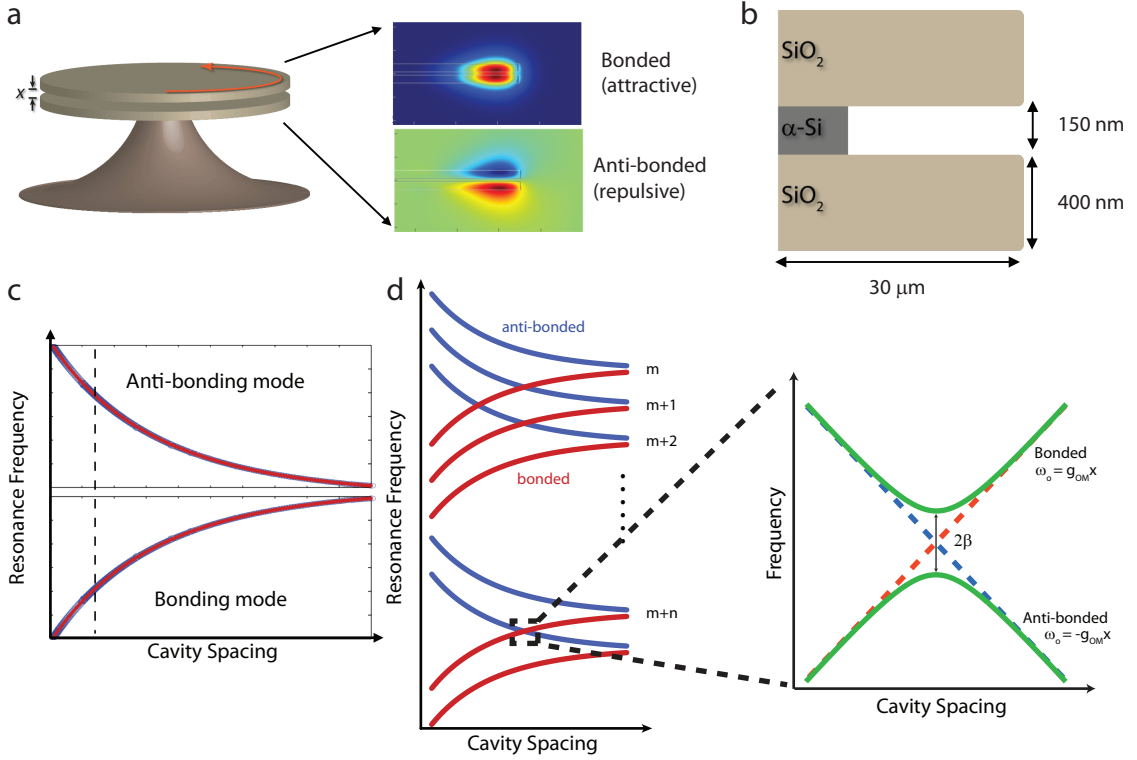
$$g' = \frac{d^2\omega_o(\hat{x})}{d\hat{x}^2}. \quad (1.2)$$

Writing down the Hamiltonian for this system we have,

$$\begin{aligned} H &= \hbar\omega_o(x)\hat{a}^\dagger\hat{a} + \hbar\omega_m\hat{b}^\dagger\hat{b} \\ &= \hbar(\omega_o + \frac{1}{2}g'\hat{x}^2)\hat{a}^\dagger\hat{a} + \hbar\omega_m\hat{b}^\dagger\hat{b} \\ &= \hbar\omega_o\hat{a}^\dagger\hat{a} + \frac{\hat{x}^2}{2}(k_0 + \hbar g'\hat{a}^\dagger\hat{a}) + \frac{\hat{p}^2}{2m}, \end{aligned} \quad (1.3)$$

where we have used the Taylor expansion from equation 1.1 and the definition of the quadratic coupling from equation 1.2 and  $k_0$  is the mechanical spring constant. Here the coupling between the light and the mechanics is dependent on the square of the position,  $\hat{x}^2$ , not simply the position  $\hat{x}$ .

This is a significant difference and means that the signal read out by the optical field is proportional to the energy of the mechanical mode as opposed to its position. It also means that instead of influencing the mechanical mode through an optical force (which would be the case for the linear coupling to  $\hat{x}$ ), the optical field now modifies the potential energy of the mechanical mode and acts as a photon number controlled static spring constant (we refer to it as a static spring to highlight the difference between the dynamical optical spring common in linear optomechanics).



**Figure 1.1: Tuning a double disk.** (a) Model of the double disk structure with two silica disks, separated by a gap  $x$  sitting atop a pedestal. FEM simulations of the bonded and anti-bonded modes are shown. Optical power in the bonded modes has an attractive force on the disks, decreasing their separation. The opposite is true for the anti-bonded modes. (b) Schematic of the device showing the two silica ( $\text{SiO}_2$ ) disks each 400 nm thick separated by a 150 nm amorphous silicon ( $\alpha\text{-Si}$ ) spacer partial etched away. (c) Diagram showing how the resonance frequency of the bonded and anti-bonded mode tune as the disk spacing is changed. The bonded mode decreases in resonance frequency as the gap is decreased while the anti-bonded mode increases in frequency. (d) The combination of many azimuthal orders of optical resonances found in disk (here represented by azimuthal order number  $m$ ) and different tuning directions of the bonded and anti-bonded modes provides a means to tune modes of different order into an anticrossing as depicted on the right.

The fact that the optical field is now sensitive to the energy of the mechanical oscillator has garnered a lot of interest with the hope that it will provide a means of doing a quantum non-demolition measurement on the mechanical mode [1]. Quadratic coupling has previously been studied in the context of superconducting nano-electromechanics [2, 3] and also in optomechanics where the

canonical quadratic coupling system is the membrane in the middle, studied extensively in the Harris group at Yale [4, 5], and has received much theoretical attention [6–10].

As mentioned above, most optomechanical systems have only a linear optomechanical coupling, and it is more difficult to design a system with quadratic coupling. The system we use does not initially have a quadratic coupling but can be made to have one as will be shown below. The device we use consists of two microdisk optical resonators fabricated out of silicon dioxide and stacked on top of each other with a very small separation (see Figure 1.1b). Each disk supports its own host of high-Q optical resonances and when the disks are brought near each other the modes renormalize into bonded and anti-bonded pairs whose frequency is very sensitive to the separation of the two disks, as shown in Figure 1.1a. The bonded mode has significant field between the disks, and as such the frequency of the bonded modes decreases as the disk gap decreases (and the effective index seen by the mode increases). The anti-bonded mode's field has a node between the disks, and as such when the disk's gap decreases this increases the frequency of the mode. The dependence of the resonances' frequencies on the spacing between the disks can be seen in Figure 1.1c. The bonded and anti-bonded modes can also be used to control the gap between the disks. Increasing the energy in a bonded mode (i.e., increasing the intracavity photon population) will cause the gap to decrease, lowering the overall energy of the system. Similarly, increasing the intracavity photon population of an anti-bonded mode will cause the disk gap to increase.

One of the keys to using this system for quadratic coupling is that there are many pairs of bonded and anti-bonded modes corresponding to different azimuthal orders (and radial orders) and the mode families have a free spectral range of around 5 nm. This means that we can modify the gap between the two disks and push bonded and anti-bonded modes of different azimuthal orders into an anticrossing, as shown in Figure 1.1d. Here we can see that as these modes are pushed towards each other by modifying the disk gap they begin to interact, and anti-cross giving rise to the quadratic optomechanical coupling. We can see that at the center of the anticrossing the resulting optical modes (green curves in the zoomed in plot of Figure 1.1d on the right) have no dependence on the cavity spacing, to first order. However, the resonance frequency does have a dependence on the square of the position,  $\omega_o = g'\hat{x}^2/2$ . It is through this method that we are able to create the quadratic coupling in the double disks. The strength of this coupling is determined by the slopes of linear couplings of the two anticrossing modes as well as their separation  $2\beta$  and is given by

$$g' = \frac{g_{\text{OM}}^2}{\beta}, \quad (1.4)$$

where  $g_{\text{OM}}$  is the linear optomechanical coupling and represents the shift in the optical resonance frequency, given a change in the position of the mechanical resonator.



## 1.2 Fabrication




Figure 1.2: **Fabrication of double disks.** (a) Stackup showing the two 400 nm silicon dioxide (SiO<sub>2</sub>) layers grown on the silicon substrate. The layers are separated by a 150 nm grown layer of amorphous silicon (α-Si). (b) A SEM of the ZEP mask layer used for etching. (c) A SEM of the disks after being etched. The two silica layers can be seen to be separated by the brighter silicon layer. (d) A floppy double disk (referred to as a spider web) after being under cut by an isotropic SF<sub>6</sub> plasma. (e) Zoomed in view of the released double disks showing a ~150 nm disk gap. (f) A different design that is much stiffer and has a higher mechanical frequency. The optical resonance resides in the outer ~6 μm of the disk, providing a lot of ability to engineer the mechanical properties separately as seen with (d).

To create two disks stacked vertically on top of each other and free to move requires the material stackup to be grown ourselves with a suitable sacrificial layer separating the two disks. The disks are then defined lithographically and etched anisotropically before being released in an isotropic etch step. The steps are outlined in Figure 1.2 and will be explained in detail below.

### 1.2.1 Material Growth

#### 1.2.1.1 PECVD of Si and SiO<sub>2</sub>

The fabrication of the double disks begins with the growth of the material stackup as seen in Figure 1.2a. The double disk structure requires two silica, optical microdisk resonators to be vertically stacked with an intermediate, sacrificial layer acting as a spacer between the two disks. In order to have exactly the stackup we want with the correct silica disk thicknesses and separation, we grow our own material using plasma enhanced chemical vapor deposition (PECVD). We need to grow a silica layer followed by the silicon layer and finally the last silica layer, and each step can be done in the same chamber sequentially with the exact recipes outlined in Table 1.1 and explained in detail below. The PECVD machine accepts 100 mm wafers and many large Si substrate chips can

be packed closely (this is important to avoid edge effects) together on a single wafer and all grown at the same time. After the wafer is transferred into the PECVD chamber it is let to sit for 20–30 minutes to allow the wafer and chips to thermalize to the table which is kept at 350°C.

The first step is the silica ( $\text{SiO}_2$ ) growth of what will become the bottom disk resonator. The silica growth is very straightforward and the parameters are listed on the first line of Table 1.1. The growth rate of the  $\text{SiO}_2$  is 50–60 nm/minute with some variance across the chip. Edge effects, that is, how close to the edge of the chip, can change the growth rate and, as such, many chips are packed together and there is an understanding that the material at the edge will likely not be as good to use. The growth step produces a much thicker layer of silica than we want but is calibrated to produce a  $\sim 400$  nm thick layer after the anneal step described below.

After the material for the bottom disk is grown, the sacrificial Si spacer layer can be immediately grown on top using the recipe in Table 1.1. This step is slightly more complicated as the plasma will not strike (the gas will not be ionized) for the parameters we actually wish to use to deposit the material. As such, we set parameters for a strike step which are immediately changed to the deposition parameters after the plasma is struck and the plasma glow is seen through the view port. For the Si deposition the HF power must be set to 20 W for the strike step and then immediately changed to 8 W after the plasma is struck. This deposition is again calibrated to give  $\sim 150$  nm of Si after the anneal step. The final step is depositing silica for the top disk which is done identically to the first, bottom disk.

| Material       | $\text{SiH}_4$<br>(sccm) | $\text{N}_2\text{O}$<br>(sccm) | $P_{\text{ch}}$<br>mTorr | HF<br>W   | T<br>°C   | Time<br>(mm:ss) | Description              |
|----------------|--------------------------|--------------------------------|--------------------------|-----------|-----------|-----------------|--------------------------|
| $\text{SiO}_2$ | 42 [36]                  | 838 [822]                      | 1000 [994]               | 20 [20/0] | 350 [350] | 9:32            | 400 nm bot-<br>tom layer |
| Si             | 300 [292]                | –                              | 2000 [1992]              | 20        | 350 [350] | –               | Si strike step           |
|                | 300 [292]                | –                              | 2000 [1992]              | 8 [8/0]   | 350 [350] | 6:46            | 150 nm spacer<br>layer   |
| $\text{SiO}_2$ | 42 [36]                  | 838 [822]                      | 1000 [994]               | 20 [20/0] | 350 [350] | 9:32            | 400 nm top<br>layer      |

Table 1.1: **Double disk deposition recipe.** Recipes to grow the stackup for the double disks. First, a layer of  $\text{SiO}_2$  is deposited. The deposition rate is roughly  $\sim 25$ – $30$  nm/second giving  $\sim 475$ – $570$  nm of  $\text{SiO}_2$ . After the anneal step (also called densification), this layer will be  $\sim 400$  nm thick and is what the deposition is calibrated to. The Si deposition has a strike step and the actual deposition. Immediately after the plasma is struck the HF power is changed from 20 W to 8 W. This Si layer will act as a spacer between the two disk resonators and is a sacrificial layer which will mostly be removed later. A final  $\text{SiO}_2$  layer is deposited the same as the first, bottom layer.

### 1.2.1.2 Anneal

After the material is grown it is annealed in a furnace at 1050°C for 10 hours. This step increases the density of the material and decreases absorption due to unwanted chemical groups (Si-NH, Si-

H, Si-OH, etc. within the material) [11, 12]. As the material's density is increased the thickness also decreases resulting in the  $\sim 400$  nm thickness we desire. The increase in the density (and thus refractive index) and removal of unwanted absorption peaks produces a material that can support much higher quality factor optical resonators. Having a high quality factor is always important, but for the quadratic coupling it is especially important as the linewidth (inverse of the quality factor) of the two optical modes anticrossing determines the smallest possible separation ( $\beta$ ) that can be distinguished, and thus limits the size of the quadratic coupling possible (see Equation 1.4).

### 1.2.2 Electron Beam Lithography

The next step is beam-writing the pattern for the etch mask. As the double disks are quite large (typical diameter between  $30\ \mu\text{m}$  and  $60\ \mu\text{m}$ ), with no small features, this step is very straightforward and requires no optimization. However, the resist layer does need to be very thick as we need to etch through  $800\ \text{nm}$  of  $\text{SiO}_2$  and  $150\ \text{nm}$  of Si. To create a very thick layer we double spin the electron beam resist (ZEP 520A). The first layer is spun at  $2000\ \text{rpm}$  for  $60$  seconds and then baked at  $180^\circ\text{C}$  for  $5$  minutes and left to sit for  $5$  minutes to cool. The chip is then spun again at  $2000\ \text{rpm}$  for  $60$  seconds and baked for  $15$  minutes again at  $180^\circ\text{C}$ . After the beamwrite the chips are developed for  $3$  minutes in the developer and  $30$  seconds in the rinse. A developed resist layer ready to be used as an etch mask is shown in Figure 1.2b and Figure 1.3.




Figure 1.3: **Examples of the double-spun resist.** (a) An example of the double-spun resist after development showing the mask that will be used to etch the disk. (b) Zoomed in view of the resist shown in (a). Some striations can be seen in the sidewalls of the resist which could transfer into the device layers.

### 1.2.3 Etching

A long pseudo-Bosch etch is used to etch through the entire stackup in one etch with the parameters shown in Table 1.2. This is a difficult task and careful optimization is required to be able to get

all the way through both layers of silica without completely removing the resist layer and beginning to etch the silica layer that is the top disk. Results from various etches can be seen in Figure 1.4. The first SEM in Figure 1.4a shows an example of a very deep over-etch of the double disks. Here both top and bottom silica disk layers can be seen etched through and the silicon substrate below is further etched  $\sim 500$  nm—far too deep. To etch this deeply requires a combination of a very aggressive etch or a very long etch time (or both) and can be seen to have warped the resist layer which is beginning to peel off the top of the disks. This is generally a result of extreme temperatures and leads to poor etch results as the resist is no longer doing its job as an etch mask. In contrast, Figure 1.4b shows an under-etched device. The resist on top of the disks can be seen clearly and looks good. The top disk and spacer silicon layer are etched through, however, the bottom silica disk is only half etched through.




Figure 1.4: **Examples of good and bad double disk etches.** (a) An example of over-etching. The silicon below the disks is etched nearly 500 nm, which is much too deep. The resist layer is seen to be peeling back from the top of the disk due to the long, harsh etch. (b) An example of under-etching. The bottom disk is barely etched halfway through. (c) An example of pitting of the silica layer and striations due to the etching. This rough disk surface would cause significant optical scattering leading to a low  $Q$  device which would not be usable. (d) A pretty good etch that is clearly through both disk layers and shows almost no pitting or striations. This is the device used for the experiments: DSP06A4.

Another problem arises from damaging the resist at the edge of the pattern which can lead to

striations being etching into the outside edge of the disks, causing unwanted optical scattering. An example of this can be seen in Figure 1.4c which shows striated resist on top of the disks and striations which have been transferred into the device layer. A significant amount of pitting, especially around the spacer silicon layer, can also be seen. Much of the etch optimization was focused on making the sidewalls of the disks as smooth as possible to minimize unwanted optical scattering. The image in Figure 1.4d shows a pretty good etch with smooth sidewalls and a reasonable depth into the silicon substrate. This device is the one used in the experiments described later in the chapter and the parameters for this etch are listed in Table 1.2

| Time<br>(mm:ss) | C <sub>4</sub> F <sub>8</sub><br>(sccm) | SF <sub>6</sub><br>(sccm) | C <sub>4</sub> F <sub>8</sub> /SF <sub>6</sub> | RF Forward<br>(W) | DC Bias<br>(V) | ICP<br>(W)    | He Backing<br>(Torr; sccm) |
|-----------------|-----------------------------------------|---------------------------|------------------------------------------------|-------------------|----------------|---------------|----------------------------|
| 19:00           | 22 [24]                                 | 12 [12]                   | 2.0                                            | 15.5 [17/5]       | 41/42          | 1300 [1263/2] | 10 [10/5.6]                |

Table 1.2: **Double disk etch parameters.** The parameters used for the double disk etching. These parameters are from the etch used on device DSP06A4 which is the device the experiment was performed on.

#### 1.2.4 Release

The final step in preparing the double disks is the release of the disks from each other and the substrate below. Since both the spacer layer and the substrate are silicon, this can be done in one step. Due to the spacing between the disks, wet etching the silicon between them with a chemical like KOH (potassium chloride) would cause the surface tension to pull the disks together as the liquid evaporated causing them to stick permanently. To avoid this, a critical point dryer could be used to avoid any evaporation by continuously moving the liquid the disks are immersed in around the critical point (i.e., no phase transition and no evaporation). Another option is to use a dry etch similar to the etching step used to define the disks themselves. The difference here is that we need a selective etch that is isotropic. This can be accomplished by striking a SF<sub>6</sub> plasma in the etch chamber and then not accelerating it towards the sample (i.e., RF forward power of 0 W). The SF<sub>6</sub> plasma will chemically attack the silicon preferentially, and as such will provide an isotropic undercut of the disks and remove the silicon spacer between them. The parameters for this etch are outlined in Table 1.3. Proper timing is required to remove enough of the silicon spacer between the disks without completely removing the silicon underneath the disks, leaving a pedestal to support them. Through trial and error it was found that the best solution was to do several 3 minute release steps followed by 15-20 minutes to let the sample cool back down and a total of 4-5 of these steps was required based on the specific design used. The etch is very exothermic and a continuous etch is too aggressive at removing the substrate (it is a much larger open area) relative to the silicon spacer and results in devices completely removed from the substrate. Figure 1.2e shows a zoomed



in SEM of the gap between the two released disks. The upper disk is noticeably slanted and for this particular device it was designed to be this way.

| Step     | RF Forward<br>(W) | ICP<br>(W)  | SF <sub>6</sub><br>(sccm) | P <sub>ch</sub><br>mTorr | He Backing<br>(Torr; sccm) | T<br>°C    | Time<br>(mm:ss) |
|----------|-------------------|-------------|---------------------------|--------------------------|----------------------------|------------|-----------------|
| Strike   | 10                | 950         | 100                       | 15                       | 20                         | 10         | –               |
| Release  | 0                 | 950 [924/3] | 300 [300]                 | [102.3]                  | 20 [20/19.5]               | 10 [15,16] | 3:00            |
| Cooldown | –                 | –           | –                         | –                        | –                          | –          | ≥15:00          |

Table 1.3: **Double disk SF<sub>6</sub> dry release parameters.** The parameters used for the SF<sub>6</sub> dry release of the double disk devices. The chamber pressure is not set, but instead the angle of the valve is what leads to a constant pumping rate and variable pressure. The release step produces a large amount of heating and is limited to 3 minutes. After at least 15 minutes of waiting to let the sample cool down, the release can be repeated. To fully release the structures 5 release steps are often required.

### 1.2.5 Independent Engineering of the Mechanical Properties

We are generally only interested in the first order (and possibly the second order) optical resonances of the disk. The mode profile of these resonances is at the very edge of the disk and does not extend inward more than 6  $\mu\text{m}$  or so (higher order radial modes extend further inwards radially). As such, we can modify the inner structure of the disks to engineer a specific stiffness and thus mechanical frequency as long as we leave the outer rim untouched. The disk in Figure 1.2d has a spring constant of  $k = 1.63 \text{ N/m}$  while the structure in Figure 1.2f has a spring constant of 128 N/m, almost two orders of magnitude stiffer. The disk shown in Figure 1.2d is incredibly floppy and can be difficult to handle without it collapsing. To put its stiffness in perspective, a small piece of aluminum foil of size 4x1 mm has a spring constant of  $k = 1 \text{ N/m}$ , and the double disks are only 140 nm apart.

## 1.3 Measurement

The measurements on the double disks for the quadratic coupling can be quite involved and require a setup with the ability to use up to three separate lasers simultaneously as well as optical detectors, optical modulators, and spectrum analyzers. This setup is shown in Figure 1.5. As mentioned, three separate lasers are required and can lead to confusion in properly identifying them. They each have specific tasks and operate at different wavelengths, and will be labeled accordingly. The first laser is the 980 laser which will be used to probe and interact with the two optical modes involved in the anticrossing producing the quadratic coupling. This laser operates over the wavelength range of  $\sim 960\text{--}995 \text{ nm}$ , which is not nearly as popular a range as the 1520–1570 nm range, and equipment here is much more difficult to obtain. It is the 980 laser that can be

considered the science laser, while the others are more of helper lasers. There are two lasers in the 1500 nm wavelength range which are combined by a wavelength division multiplexer (WDM) that limits their operation to a 20 nm window around their respective center frequencies: 1530 nm and 1550 nm. The first laser, the 1530 laser, is used to probe the mechanical motion of the system by side-locking to an optical mode and via the linear optomechanical transduction monitoring the position of the mechanical resonator and will be discussed in more detail in Section 1.3.2. The last laser is the 1550 laser which is used to tune the spacing between the disks and push the 980 modes into an anticrossing, and this will be explained in Section 1.3.1.1.

### 1.3.1 Tuning to an Anticrossing

The double disks is not a system that inherently has a quadratic coupling, in fact they have a very strong linear optomechanical coupling similar to that of the zipper optomechanical system and have been used extensively for this purpose [13–15]. To create the quadratic coupling, we need to be able to change the disk gap which tunes the bonded and anti-bonded modes in the 980 nm band in opposite directions, causing them to anti-cross each other. There are two methods we employ to do this; the first uses the linear optomechanical coupling to apply a static optical force to the disks and causes the gap between them to decrease tuning the modes in the 980 nm band. This method is very fast, repeatable, and easy to control, but doesn't work well in vacuum due to the large optical powers required to significantly change the disk gap and tune the optical modes. The second method involves thermal-mechanically tuning the disks. The disks are heated or cooled which causes the gap between them to either increase or decrease respectively which then tunes the optical modes. This method is a bit slower and more difficult to control but works well in vacuum. Both these methods will be described in more detail below.

#### 1.3.1.1 Laser Tuning

The linear optomechanical interaction provides a straightforward means to change the gap between the two disks and thus tune the bonded and anti-bonded optical resonances of the system. This can be most easily understood by considering how the linear optomechanical coupling is proportional to the position of the mechanical mode, and thus force-like. In fact, the static optical force on the mechanical system is given by  $n_c \hbar g_{\text{OM}}$  from which the corresponding static displacement is given by  $\Delta x_{\text{static}} = n_c \hbar g_{\text{OM}} / k$ , which follows from Hooke's law, with  $n_c$  the intracavity photon number. The change in the optical resonance frequency is then simply given by

$$\Delta\omega_o = g_{\text{OM}} \Delta x_{\text{static}} = \frac{n_c \hbar g_{\text{OM}}^2}{k}. \quad (1.5)$$

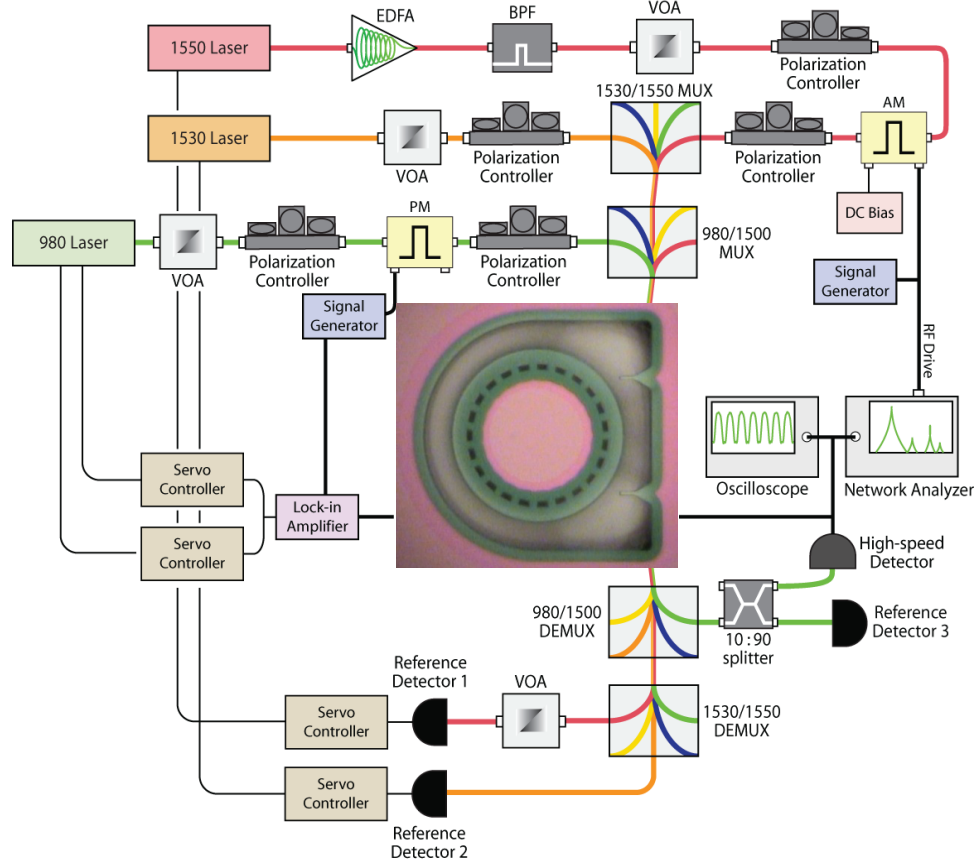
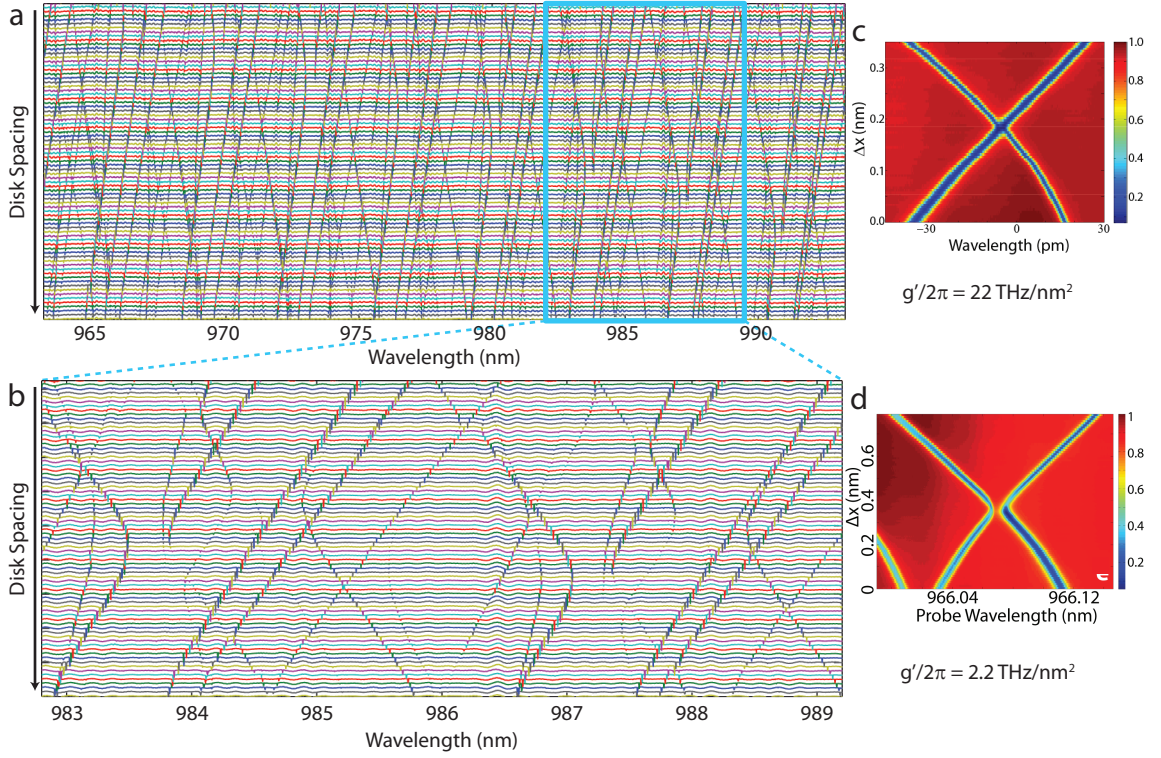


Figure 1.5: **Experimental setup for quadratic coupling.** Three lasers are available for this experiment. The 1550 laser is used to tune the gap spacing by pumping a bonded or anti-bonded mode with a large intracavity power. It is amplified by an erbium doped fiber amplifier (EDFA) and then sent through a bandpass filter (BPF) to remove excess noise. The power is set by a variable optical attenuator (VOA) and the polarization is set and it is then sent through an amplitude modulator (AM) before being combined with the other lasers in a wavelength division multiplexer (WDM) filter and sent to the device. The 1530 laser is used to probe the mechanical mode via the linear optomechanical coupling. The 980 laser probes the quadratic coupling. After the device, the light is again split using a WDM filter and the 1530 light is sent to a photodetector and spectrum analyzer to monitor the mechanical mode. The 980 light can similarly be sent to a spectrum analyzer or oscilloscope. The transmission of all the lasers can be monitored to characterize the optical resonances and used to lock the lasers to the side of the cavity. The 980 laser can also be locked to the center of the quadratic coupling by using the Pound-Drever-Hall (PDH) technique.

As such, by increasing the intracavity photon number of an optical mode with a linear optomechanical coupling we can put a static force on the mechanical system, changing the position of the resonator and thus tuning the optical modes. This method is very precise as we have very good control over the amount of power we can put in the system and allows for the control of the disk gap as fast as we can put photons into or let them decay from the optical cavity—the lifetime of the optical resonance. The tuning range depends on the intracavity photon number, and thus the power dropped into



the optical cavity. A tuning range of 4.4 nm has been demonstrated with a dropped power of 1.7 mW [14], which corresponded to a static mechanical displacement of 17.7 nm.



**Figure 1.6: Tuning the double disks into an anticrossing.** (a) A plot of transmissions scans across the full laser wavelength range showing the optical resonances of the double disks. Each scan represents a different disk spacing decreasing from top to bottom. Decreasing the disk spacing tunes the bonded modes to higher wavelengths (right) and the antibonded modes to lower wavelengths (left). From this plot many crossing and anticrossings can be seen. (b) Zoomed in view of the plot over a smaller wavelength span showing the crossings and anticrossings in greater detail. (c) An example anticrossing that is barely resolved and has a quadratic coupling of  $g'/2\pi = 22 \text{ THz/nm}^2$ . (d) An example anticrossing that is more clearly resolved and has a quadratic coupling of  $g'/2\pi = 2.2 \text{ THz/nm}^2$ .

In order to achieve such large tuning ranges, a very large amount of optical power is required—more than what is provided by the lasers themselves. The laser used for the optical tuning is the 1550 laser and it is first sent to an erbium doped fiber amplifier (EDFA), which is used to boost the available power (see Figure 1.5). The amplifier works in the wavelength range of 1530–1560 nm, and after amplification is sent through a band-pass filter (BPF) to remove excess noise during the amplification process and then to a variable optical attenuator (VOA), providing a means to control the power incident on the device. The polarization is adjusted using fiber polarization controllers and then combined with the other lasers using a WDM. This WDM limits the channel width to 20 nm centered around 1550 nm which is not a problem given the large number of optical resonances

in these structure. However, in this wavelength range (and generally true for  $\lambda \gtrsim 1350$  nm), only bonded modes are present providing only for a means to decrease the disk gap, not increase it as would be possible with anti-bonded modes.

In the 980 nm band there are both bonded and anti-bonded resonances present, and to tune them into an anticrossing the 1550 laser is locked to the side of a bonded mode and the power is slowly increased, decreasing the disk gap and tuning the bonded modes to lower wavelength and the anti-bonded modes to high a wavelength. As the 1550 laser's optical power is increased, the 980 laser is scanned across the bonded and anti-bonded resonances present in that wavelength range, and their tuning with respect to the tuning laser (1550) power is recorded. An example of this tuning versus optical power is shown in Figure 1.6a where each optical scan is shown offset in the vertical direction with the initial zero tuning scan shown at the top of the plot, and the optical power in the tuning mode is increased, and disk spacing decreased moving towards the bottom of the plot. The tuning of the optical modes is quite clear in this plot and a zoomed in section showing a smaller wavelength span is shown in Figure 1.6b with several anticrossings present and easy to identify. Since the resonance frequencies are being tuned optically, we can quickly characterize a device to find anticrossings that will lead to large quadratic coupling ( $g'$ ). When a promising anticrossing is found we can do a fine scan around it to determine exactly what the quadratic coupling will be and in Figure 1.6c we see a fine scan of an anticrossing that is barely resolvable and gives a quadratic coupling of  $g'/2\pi = 22$  THz/nm<sup>2</sup>. Anticrossings with smaller separations have higher quadratic coupling, however, to be able to resolve the anticrossing the linewidth of the optical mode needs to be smaller than the separation (otherwise we can't see the separation of the anticrossing). In these devices the linewidths are on the order of 1 pm. Another anticrossing with a better resolved anticrossing can be see in Figure 1.6d where the quadratic coupling here is  $g'/2\pi = 2.2$  THz/nm<sup>2</sup>.

The actual anticrossing used is shown in Figure 1.7a and is well resolved. At the anticrossing the resonances frequencies of the upper ( $\omega_+$ ) and lower ( $\omega_-$ ) branch of the anticrossing are,

$$\omega_{\pm} = \omega_o + \delta g_{OM}x \pm \sqrt{\beta^2 + \bar{g}_{OM}^2 x^2}, \quad (1.6)$$

where  $\omega_o$  is the center frequency of the anticrossing,  $\beta$  is half the splitting, and  $\delta g_{OM}$  is the difference between the linear optomechanical coupling of the two resonances (i.e., the tuning rate of the bonded and anti-bonded modes) and  $\bar{g}_{OM}^2$  is their average linear coupling. To determine the quadratic coupling of this anticrossing we can fit the crossing to Equation 1.6 and we see that this anticrossing has a quadratic coupling of  $g'/2\pi = 3.4$  THz/nm<sup>2</sup>. This quadratic coupling is very large and was six orders of magnitude larger than any other system at the time [4, 5] and still more than two orders of magnitude larger than the system with the next largest quadratic coupling of 20 GHz/nm<sup>2</sup> [16].

The laser tuning method is very fast, precise, and repeatable, making it a great tool to be

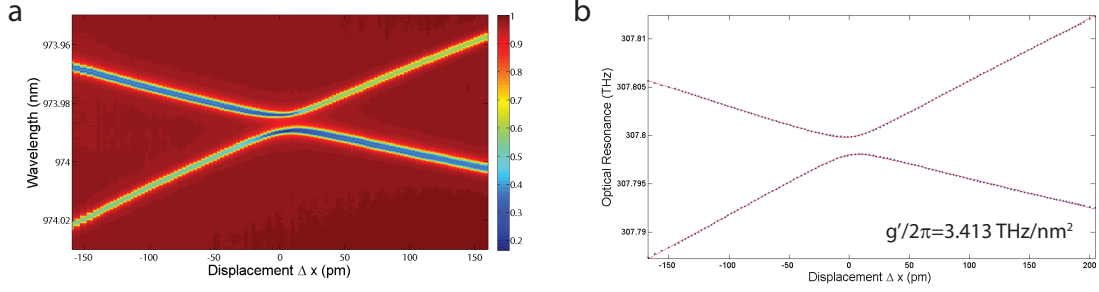


Figure 1.7: **Extraction of quadratic coupling from anticrossing used.** (a) The anticrossing used in the experiment. (b) A fit of the anticrossing to theory. The extracted quadratic coupling was  $g'/2\pi = 3.413 \text{ THz/nm}^2$ .

used to find and characterize quadratic couplings in the  $\text{N}_2$  purged test setups we use for device characterization. For the laser tuning method described above to work, as mentioned, we must put a very large number of photons into the cavity to cause any appreciable tuning of the optical frequencies. This poses a problem for the laser tuning method in vacuum as the lower mechanical frequency and higher mechanical quality factors, due to the absence of the squeezed film damping, lead to the onset of mechanical self-oscillation (lasing) before any appreciable tuning can occur. Because we have a great deal of freedom with respect to engineering of the mechanical properties of the disks independently of the optical properties, attempts were made to modify the design to increase the stiffness of the disks and find a trade-off between stiff enough to be tunable by the laser method, while not being so stiff that we couldn't carry out the experiment (which involves exploring the regime of the static optical spring dominating the intrinsic mechanical spring constant). While exploring different designs to find a sweet-spot for the intrinsic mechanical stiffness of the disks, another simpler and more elegant tuning method was discovered: thermal-mechanical tuning.

### 1.3.1.2 Thermal-Mechanical Tuning

The disk gap spacing is sensitive to thermal changes: when the disks are heated the gap increases, when cooled the gap decreases. This then provides a method for tuning the gap not requiring the 1550 tuning laser at all, simply a TE cooler with a reasonable temperature tuning range. It has the further advantage of being able to increase and decrease the gap spacing, providing access to previously inaccessible anticrossings. To this end a new mount was made for the vacuum chamber (in this case, a modified Janis ST-500 cryostat [17]) with a thermo-electric (TE) cooler [18] attached allowing for several nanometers of optical tuning in both directions (opening and closing the gap). With this TE cooler the anti-bonded mode tuned  $\sim 20 \text{ pm/K}$  and the bonded mode tuned  $\sim 2 \text{ pm/K}$ . This asymmetry is due to the fact that there are two thermal components to the tuning, thermo-mechanical (changing the disk gap) and thermo-optic [19]. For the anti-bonded mode these

two tuning components had the same sign and added, while for the bonded mode they had opposite sign and almost canceled each other out. This is not a problem for the actual experiment as it doesn't matter how we tune the modes into the anticrossing, however, it doesn't look as nice and makes it more difficult to properly fit the anticrossing to theory, as well as being much slower and less repeatable. For these reasons the characterization of the anticrossing was still accomplished with the optical tuning while the experiments in the vacuum chamber used the thermal tuning.

### 1.3.2 Linear Transduction of Mechanical Motion

The 1530 laser is used to probe the optical and mechanical properties of the system. Similar to the 1550 tuning laser, the 1530 laser is passed through a polarization controller and attenuator before being combined with the other lasers in a WDM, which again limits its bandwidth to 20 nm centered around 1530 nm (see Figure 1.5). This laser has two main roles: to monitor an optical resonance's shift due to the tuning action described above, and to monitor the position of the mechanical resonator through a linear optomechanical coupling. In the first case it is simply used for calibration. As the disk spacing is changed, it changes the optical resonances' frequencies and that change is recorded by the 1530 laser in a series of wavelength scans. Knowing the linear optomechanical coupling ( $g_{OM}$ ) of the mode being tracked by the 1530 laser, the change in the disk gap ( $\Delta x$ ) can be calibrated.

The 1530 laser is also used to transduce the mechanical motion of the resonator through the linear optomechanical coupling. When the 980 modes are tuned into an anticrossing they are no longer sensitive to the position of the mechanical mode, only its energy. By locking the 1530 laser to the side of an optical mode it will transduce, through the linear optomechanical coupling, the mechanical motion of the resonator which can be seen by looking at the RF spectrum of the transmitted 1530 laser light. In this way the mechanical position spectrum can always be monitored independently of the quadratic coupling in the 980 modes.

### 1.3.3 Suppression of Linear Coupling

Now that we have the anticrossing and have determined the quadratic coupling strength, it is important to verify that at the center of the quadratic coupling we have suppressed the linear optomechanical coupling, leaving the lowest order coupling quadratic. This can be seen quite remarkably by first tuning the gap spacing so the two optical modes involved in the quadratic coupling are not interacting yet and we are far from the quadratic coupling point while looking at a wavelength scan of the optical modes. This can be seen in Figure 1.8a. The bottom panel shows the map of the quadratic coupling with a blue line indicating the location the 980 wavelength scan was taken, which is far from the quadratic coupling point. The top panel shows the 980 wavelength scan

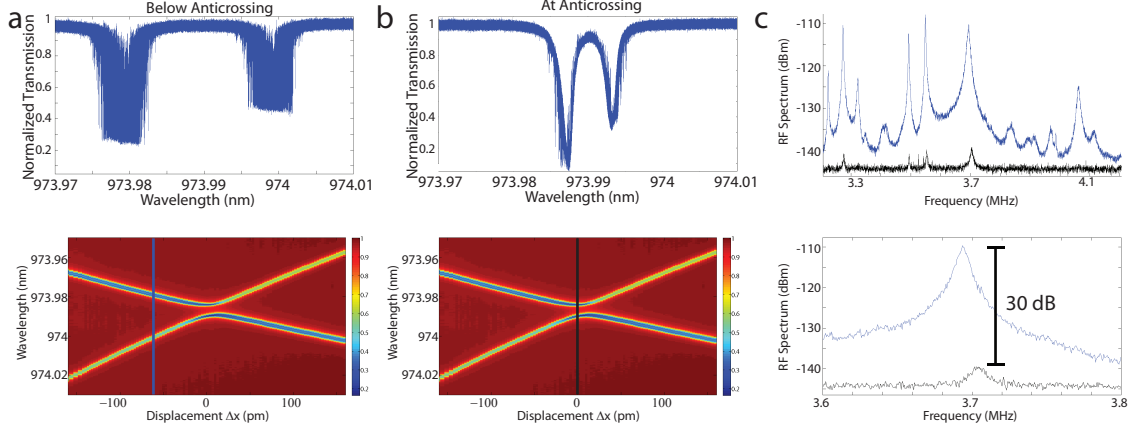


Figure 1.8: **Suppression of the linear coupling.** (a) A wavelength scan (top) of the two optical modes slightly away from anticrossing (bottom). The optical resonances are very noisy as a result of the linear optomechanical transduction of the thermal motion. (b) An optical scan (top) of the two optical modes at the anticrossing (bottom). The resonances appear much cleaner due to the suppression of the linear optomechanical coupling at the anticrossing. The thermal motion is no longer being transduced well. (c) The RF spectrum (top) of the optical signal showing the transduction of the mechanical motion (blue) when the system is not at the anticrossing as in (a) and when it is at the anticrossing as in (b) (black). At the anticrossing the linear optomechanical coupling is suppressed and, as such, the thermal motion is barely transduced. The bottom is a zoomed-in plot of the fundamental mechanical mode showing 30 dB of suppression.

where the optical modes have yet to interact and appear quite noisy. This noise is produced by the linear optomechanical coupling sensing the large amplitude of thermal motion associated with the low mechanical frequency. The linear coupling then converts this mechanical motion into a change in the wavelength of the optical modes, smearing them out considerably in the wavelength scan. By tuning the disk into the quadratic coupling point we should suppress the linear coupling and not see this noise on the optical scan as shown in Figure 1.8b. Here the bottom panel reflects the position that the 980 wavelength scan was taken, with the system tuned into the quadratic coupling. The top panel shows a 980 scan over the same wavelength range as Figure 1.8a, however, this time the modes are basically noiseless due to the suppression of the linear coupling as expected. The noise that does remain can be understood as follows: The top branch of the quadratic coupling (see the bottom panel of Figure 1.8b) corresponds to the left optical mode (lower wavelength; see the top panel of Figure 1.8b). Any thermal motion can only move the disks around the quadratic coupling point and for the top branch and left mode this means decreasing the wavelength of the mode. As such, the noise in the wavelength scan always shows the noise of the left mode decreasing the wavelength, as shown in the top panel of Figure 1.8b. The opposite is true for the bottom branch (right optical mode), where any thermal motion leads to an increase in optical wavelength as seen in the noise in the right optical mode of the top panel of Figure 1.8b. A smaller quadratic coupling, and hence shallower quadratic curve (see bottom panel Figure 1.8b) and wider quadratic coupling

point, would reduce this but we would then have a smaller quadratic coupling.

Another way to see this is to look at the RF spectrum of the transmitted 980 light, which through the linear optomechanical coupling, is a transduced spectrum of the mechanical motion. With the disks tuned far from the quadratic coupling point (as in Figure 1.8a) the linear transduction is very strong and many mechanical modes can be seen in the blue trace of the top panel of Figure 1.8c. At the quadratic coupling point the RF spectrum looks very different as shown in the black trace of the top panel of Figure 1.8c, with almost no transduction of the position of the mechanical modes due to near complete suppression of the linear coupling. The bottom panel of Figure 1.8c shows a zoomed-in view of the fundamental mechanical mode suppression of 30 dB.

## 1.4 Results

What we are interested in exploring is the ability the quadratic coupling gives us to directly control the mechanical stiffness. From Equation 1.3 we know the interaction Hamiltonian of the quadratically coupled system is given by

$$H_{\text{int}} = \frac{x^2}{2}(k_0 + \hbar g' \hat{a}^\dagger \hat{a}). \quad (1.7)$$

The frequency of the mechanical resonator is then given by  $\omega_m = \sqrt{(k_0 + k_{\text{opt}})/m}$  where  $m$  is the mass of the resonator and the static optical spring constant is given by

$$k_{\text{opt}} = \hbar g' n_c, \quad (1.8)$$

and  $n_c$  is the intracavity photon number. We can see that we have a component of the stiffness of the mechanical resonator that is due to the quadratic optomechanical coupling and whose size is controllable by the intracavity photon number. The quadratic coupling of our system is  $g' = 3.4 \text{ THz/nm}^2$  giving a per-photon spring constant of  $k_{\text{opt}}/\text{photon} = 2.3 \text{ mN/m}$ . This means we have the ability to directly control the mechanical rigidity of the system and we can do this as fast as we can get photons into and out of the optical cavity, which is much faster than the mechanical frequency, and thus allows for the possibility of classical squashing of the mechanical motion. By dynamically stiffening the system we can also suppress the perturbations due to thermal noise without changing the energy of the system.

We wish to explore our ability to dynamically and elastically control the mechanical stiffness. To do this, we tune the disks into the center of the anticrossing while monitoring the linear optomechanical transduction, allowing us to determine the point of maximum quadratic coupling. Next, the 980 laser is locked using the Pound-Drever-Hall (PDH) technique [20] to the center of the upper branch of the quadratic coupling (see Figure 1.7a). A second probe laser (the 1530 laser) is then side-locked

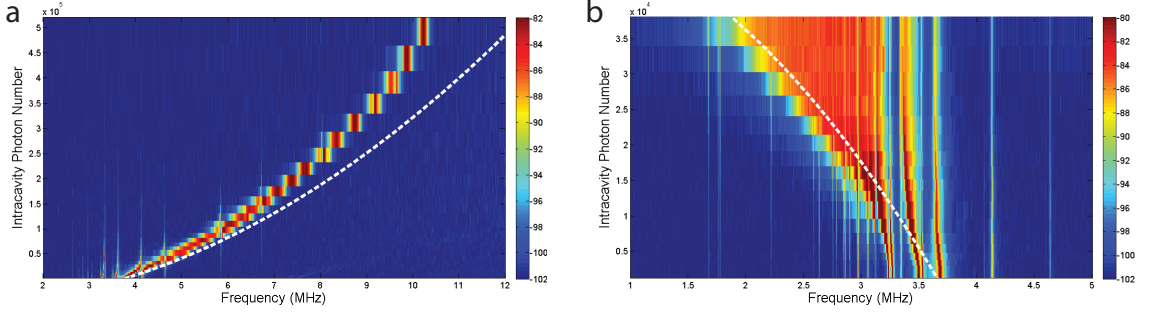


Figure 1.9: **Intracavity photon number dependence of the static spring effect.** (a) Plot showing the mechanical mode spectrum as a function of the intracavity photon number with the laser locked to the center of the anticrossing on the upper, stable branch. The dashed white line is the theoretical mechanical frequency as a function of intracavity photon number. (b) Plot showing the mechanical mode spectrum when the laser is locked to the lower, unstable branch of the anticrossing. As the power and thus intracavity photon number increases, the mechanical mode mixes with a forest of lower frequency modes. It is also difficult to lock to this branch (it is unstable) and, as such, the exact tuning of the laser becomes uncertain, with increasing power contributing to the smeared out mechanical spectrum.

to an optical mode to be used to transduce the mechanical motion via the linear optomechanical coupling—only the two modes at the anticrossing have quadratic coupling while all other modes maintain their linear optomechanical coupling. The laser power of the 980 mode (the one probing the quadratic coupling) is now increased as the mechanical motion is transduced by the 1530 laser (probing the linear coupling). The mechanical spectra is shown in Figure 1.9a where the power is increased along the y-axis and each slice shows the mechanical spectrum at that intracavity photon number. The dashed white line is a theoretical curve and the frequency of the mechanical motion of the fundamental mechanical mode can be seen to increase in good correspondence with theory. The maximum power dropped into the cavity is  $241 \mu\text{W}$  corresponding to an intracavity photon number of  $5.58 \times 10^5$  and giving a static optical spring constant of  $k_{\text{opt}} = 1255 \text{ N/m}$ . With an intrinsic optical spring constant of  $k_0 = 128 \text{ N/m}$ , we have increased the stiffness of the mechanical resonator by a factor of 10 and the mechanical frequency by a factor of 3 from 3.8 MHz to 12.1 MHz. We can also ask what happens when we lock the 980 laser to the lower, unstable branch of the quadratic coupling. Here the sign of the quadratic coupling rate is negative, leading to a negative static optical spring and a decreasing of mechanical stiffness with increase intracavity photon number. Figure 1.9b shows the results similar to Figure 1.9a, but now with the frequency of the mechanical mode decreasing with applied power. The dashed white line is again a theoretical curve. Two important differences should be noted. First, the spectrum of the mechanical mode appears to broaden significantly as the power is increased. This is due to two effects, the first of which is that there are many mechanical modes at lower frequencies than the fundamental and as we tune the mode through these they hybridize and merge into each other. This didn't happen when locked to the upper branch



due to the lack of mechanical modes at higher frequencies. The second important difference is that we can put in much less power with a maximum dropped power of  $19.2 \mu\text{W}$  corresponding to an intracavity photon number of  $4.45 \times 10^4$ . This is because the lower branch is the unstable branch and any deviation from the very center of the quadratic coupling causes it to run away instead of being trapped. As such, the lock of the laser to the very center has to compensate for any deviation and at higher powers this becomes more difficult, leading to an uncertainty in the lock point causing further broadening of the mechanical mode's spectrum. This leads to a maximum dropped power before the system is too unstable to lock. It is shown here only to demonstrate that it is possible to both increase and decrease the intrinsic spring constant.

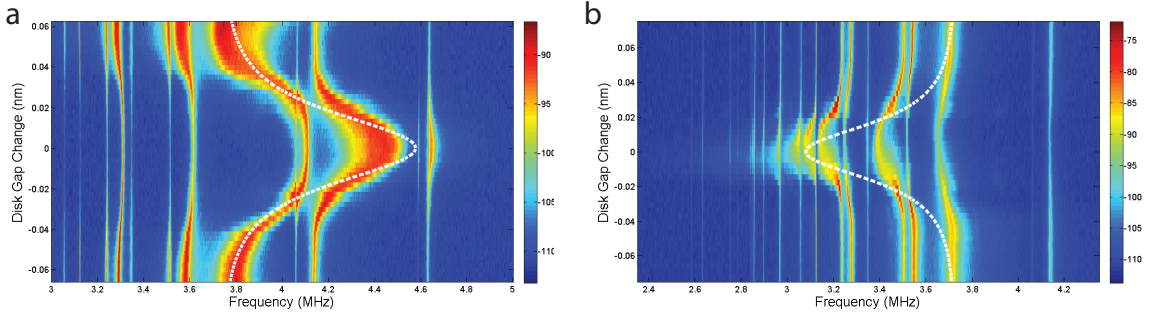


Figure 1.10: **Detuning from the quadratic coupling point dependence of the static spring effect.** (a) Plot of the detuning dependence of the static spring effect. Here the power was kept constant and laser locked to the stable branch while the disk spacing was tuned through the anti-crossing. The dashed white line is again the theoretical prediction of the mechanical frequency. (b) Same as in (a) except now the laser was locked to the unstable branch.

As the static optical spring constant is given by  $k_{\text{opt}} = \hbar g' n_c$ , and we know that we must tune the disks into the quadratic coupling, we can investigate the detuning dependence of  $g'$  on  $k_{\text{opt}}$ . Using Equation 1.6 we can write the quadratic coupling rate as a function of the position of the disks with  $x = 0$  being the quadratic coupling point,

$$g'(x) = \frac{\pm g_{\text{OM}}^2 / \beta}{(1 + g_{\text{OM}}^2 x^2 / \beta^2)^{3/2}}. \quad (1.9)$$

Here the effects of a change in the separation between the disks around the optical quadratic coupling point is made explicit. In the previous experiment we considered the case of  $x = 0$ . To see and verify this effect, we begin by tuning the disks away from the optimal quadratic coupling point and lock the 980 laser to the mode that leads into the upper branch, while keeping its power, and thus the intracavity photon number, constant. We then slowly tune the disks' separation through the optimal point while again monitoring the linear transduction of the mechanical motion with the 1530 laser. The results are shown in Figure 1.10a with the dashed white line again theory. In these spectra we have zoomed in on the fundamental mechanical mode and the correspondence with theory is very



good. Also evident in these scans is the mixing of the fundamental mode with other slightly higher frequency modes as it is tuned through them. The 980 power was much lower than the experiment looking at the power dependence, so the lock would be more stable. We can repeat this experiment but this time locking to the optical mode that will, when tuned through the quadratic coupling, lead into the lower, unstable branch. A similar amount of optical power is used and the results are shown in Figure 1.10b with good correspondence to theory (dashed white line). As expected, the mechanical mode is tuned to a lower frequency and mode mixing is again observed.

## Chapter 2

# Theory of Sideband-Resolved Optomechanical Systems

This chapter will develop the basic theoretical framework to understand the optomechanical interaction specifically in the sideband resolved regime. A thorough development of the theory leading up to what is required to understand optomechanical wavelength conversion will be presented.

### 2.1 Model of an Optomechanical System

The canonical optomechanical system consists of a Fabry-Pérot optical cavity [21] as shown in Figure 2.1. This cavity consists of two reflective mirrors a distance  $L_0$  apart. Due to interference effects, only optical fields with resonant wavelengths can exist in this cavity. Specifically, the wavelength of the light is  $\lambda_0/2 = L_0/n$ , where  $n$  is a positive number labeling the order of the optical mode in the cavity. We now make the assumption that the size of the optical cavity is not fixed, i.e., the distance between the two mirrors  $L_0$  varies. This can be model by assuming one of the mirrors is fixed and the other is attached to a spring providing a stiffness  $k$  as in Figure 2.1. In this case we can write down a Hamiltonian of our system consisting of the energy of the optical and mechanical mode

$$H = \hbar\omega_o(\hat{x})\hat{a}^\dagger\hat{a} + \hbar\omega_m\hat{b}^\dagger\hat{b} \quad (2.1)$$

where  $\hat{a}$  ( $\hat{a}^\dagger$ ) and  $\hat{b}$  ( $\hat{b}^\dagger$ ) are the annihilation (creation) operators for the optical and mechanical mode, respectively. The frequency of the optical mode is given by  $\omega_o(x)$  which is now dependent on the position ( $\hat{x}$ ) of the mechanical resonator. To understand the effect the position-dependent optical resonance frequency has on the system we can assume that the change in the optical frequency is small (compared to the unperturbed frequency), and expand  $\omega_o(\hat{x})$  in a Taylor series giving

$$\omega_o(\hat{x}) = \omega_o + \frac{d\omega_o(\hat{x})}{d\hat{x}}\hat{x} + \frac{1}{2!}\frac{d^2\omega_o(\hat{x})}{d\hat{x}^2} + \dots \quad (2.2)$$




Figure 2.1: **Canonical optomechanical system.** The canonical optomechanical system is a Fabry-Pérot cavity with length  $L_0$  and nominal frequency  $\omega_o$ . One end mirror is fixed and cannot move while the other is free to move as a mass (with effective mass  $m$ ) on a spring (with intrinsic spring constant  $k$ ) and frequency  $\omega_m = \sqrt{k/m}$ . The quantum annihilation (creation) operators for the optical and mechanical systems are  $\hat{a}$  ( $\hat{a}^\dagger$ ) and  $\hat{b}$  ( $\hat{b}^\dagger$ ) respectively, with the optical cavity input and output given by  $\hat{a}_{\text{in}}$  and  $\hat{a}_{\text{out}}$ . Modified figure from [22].

with  $\omega_o$  representing the unperturbed optical frequency.

From the second term in Equation 2.2 we can see that the optical frequency has a dependence on the position ( $\hat{x}$ ) of the mechanical resonator, with a coupling coefficient given by  $g_{\text{OM}} = d\omega_o(\hat{x})/d\hat{x}$  representing the shift in the optical resonance frequency given a change in the position of the mechanical resonator. The third term has a dependence on the position-squared ( $\hat{x}^2$ ) of the mechanical resonator and is referred to as a quadratic coupling. The expansion contains all the higher-order couplings, however, in almost every case the linear position-dependent ( $\hat{x}$ ) coupling dominates. It is possible to engineer a system with a vanishing linear coupling leaving the quadratic coupling as the dominate perturbation. This is an active area of research and quite interesting as the optical system couples to  $\hat{x}^2$ , a coupling that provides a means of sensing the energy of the mechanical mode directly as well as exerting direct optical control of the mechanical potential (i.e., through an optical stiffness) [4, 5, 7–10, 23, 24].

Keeping only the linear term in the expansion,

$$\begin{aligned} H &= \hbar(\omega_o + g_{\text{OM}}\hat{x})\hat{a}^\dagger\hat{a} + \hbar\omega_m\hat{b}^\dagger\hat{b} \\ &= \hbar\omega_o\hat{a}^\dagger\hat{a} + \hbar\omega_m\hat{b}^\dagger\hat{b} + \hbar g\hat{a}^\dagger\hat{a}(\hat{b}^\dagger + \hat{b}) \end{aligned} \quad (2.3)$$

where we have expanded the position operator in terms of the quadrature representation of the creation and annihilation operators,  $\hat{x} = x_{\text{zpf}}(\hat{b}^\dagger + \hat{b})$ , and defining the single photon coupling rate  $g = g_{\text{OM}}x_{\text{zpf}}$  with  $x_{\text{zpf}} = \sqrt{\hbar/2m\omega_m}$  the root-mean-squared (RMS) zero-point fluctuation in the position quadrature associated with ground/vacuum state of the mechanical mode  $\hat{b}$  (see Appendix A.4). From Equation 2.3 we can identify the Hamiltonian of the optical system, mechanical system, and the interaction. This coupling between the optical and mechanical modes means that the optical mode can extract information about the mechanical mode and also influence it. The

opposite is also true and the mechanical mode can extract information about the optical mode and also influence it. From an experimental point of view we use the optical system (our lasers) to probe and manipulate the mechanical system, and our language reflects this. There are many reasons for this asymmetric viewpoint, foremost being the technological resources allowing for the generation, manipulation, and detection of photons as compared to phonons. Another important reason to consider is the relative energy scales: kilohertz-gigahertz for the mechanical systems and terahertz for the optical systems. At these energy scales thermal noise is prominent in the mechanical systems at room temperature while the optical system is quiet, giving us a quantum limited system to probe the mechanics (consider  $n_b = \hbar\omega/k_B T$  the thermal occupancy of the bath at a temperature  $T$  and frequency  $\omega$ ). It is easy to see from Equation 2.3 that the position  $\hat{x}$  of the mechanical mode couples to the optical mode in a way that changes the frequency of the optical mode. To see this consider that the position is coupled to the energy operator ( $\hat{a}^\dagger \hat{a}$ ) of the optical system. The equation of motion for the energy operator will not see this term as the energy operator will commute with itself (recall  $\hbar \dot{\hat{A}} = i[\hat{H}, \hat{A}]$ ). However, this will not be the case for the operator for the opposite quadrature, the phase quadrature, thus leading to a coupling to the frequency. Similarly, the coupling to the mechanical system is through the position operator ( $\hat{x}$ ) and thus will show up in the equation of motion for the momentum operator and look like a force dependent on the photon number.

## 2.2 Equations of Motion for Linearly Coupled Optomechanical Systems

To further understand the system we can study the time evolution of its operators; namely the annihilation (creation) operators of our optical and mechanical modes:  $\hat{a}^\dagger$  ( $\hat{a}$ ) and  $\hat{b}^\dagger$  ( $\hat{b}$ ). The Heisenberg-Langevin equations of motion can be immediately obtained from  $\hbar \dot{\hat{A}} = i[\hat{H}, \hat{A}]$  giving,

$$\begin{aligned}\dot{\hat{a}} &= -(i\Delta - \frac{\kappa}{2})\hat{a} - ig\hat{a}(\hat{b}^\dagger + \hat{b}) - \sqrt{\eta\kappa}\hat{a}_{\text{in}} - \sqrt{\kappa'}\hat{a}_{\text{in},i} \\ \dot{\hat{b}} &= -(i\omega_{\text{m}0} - \frac{\gamma_i}{2})\hat{b} - ig\hat{a}^\dagger\hat{a} - \sqrt{\gamma_i}\hat{b}_{\text{in}},\end{aligned}\tag{2.4}$$

where the loss rates and input couplings have been added heuristically. For the optical mode: the total energy loss (decay) rate  $\kappa$ , input/output coupling  $\eta\kappa$  to an external source (waveguide)  $\hat{a}_{\text{in}}$ , and coupling  $\kappa'$  which denotes all optical loss channels that go undetected (information lost) and associated noise input  $\hat{a}_{\text{in},i}$ . The equation of motion for the optical mode has been moved into a frame rotating at the frequency of the laser drive such that  $\Delta = \omega_o - \omega_\ell$ , allowing us to focus on the slowly varying response around this frequency (for details see Appendix B.2.2). Similarly for the mechanical mode: intrinsic loss rate  $\gamma_i$  and the noise input due to the coupling to the thermal bath  $\hat{b}_{\text{in}}$ . Details regarding the input/output coupling, noise, and dissipation can be found in Appendix B.2.2.

### 2.2.1 Linearization Around a Strong Coherent Field

The interaction Hamiltonian ( $H_{\text{int}} = \hbar g \hat{a}^\dagger \hat{a} (\hat{b}^\dagger + \hat{b})$ ) is intrinsically nonlinear with the strength of the interaction depending on the intensity of the light ( $\hat{a}^\dagger \hat{a}$ ), leading to nonlinear equations of motion. Linearizing the equations of motion will make them easier to solve, and, more importantly, aid in providing physical insight. To do so we use a common technique in quantum optics which is to split the operators into a large coherent displacement (the classical component) and fluctuations (the quantum component). This is accomplished by making the transformation  $\hat{a} \rightarrow \alpha + \hat{a}$  and  $\hat{b} \rightarrow \beta + \hat{b}$  where  $\alpha$  and  $\beta$  represent the large classical field (i.e., the amplitude of a coherent field displaced from the vacuum state) and operators  $\hat{a}$  and  $\hat{b}$  now represent the vacuum fluctuations. It is instructive to quickly look at what the number operator looks like in this approximation,

$$\begin{aligned} \hat{n} &= \hat{a}^\dagger \hat{a} \\ &= (\alpha^* + \hat{a}^\dagger)(\alpha + \hat{a}) \\ &= |\alpha|^2 + \alpha^* \hat{a} + \alpha \hat{a}^\dagger + \hat{a}^\dagger \hat{a}. \end{aligned} \tag{2.5}$$

With the assumption that the fluctuations are small we can neglect all but the first order terms (i.e. ignore  $\hat{a}^\dagger \hat{a}$ ). Further, if we assume  $\alpha$  to be real, we can write  $\hat{n} = \alpha^2 + \alpha \hat{X}_1$  with  $\hat{X}_1 = \hat{a}^\dagger + \hat{a}$  being the in-phase quadrature of the optical mode. Since a photodetector naturally measures the photon number (flux) we can identify this quadrature as the amplitude quadrature.

Making these substitutions in Equation 2.4 along with a similar substitution for the input  $\hat{a}_{\text{in}} \rightarrow \alpha_{\text{in}} + \hat{a}_{\text{in}}$  and splitting the equations into classical steady state ( $\dot{\alpha} = \dot{\beta} = 0$ ) and fluctuating terms gives for the classical amplitudes,

$$\begin{aligned} \alpha &= \frac{-\sqrt{\eta\kappa}\alpha_{\text{in}}}{i[\Delta + g(\beta^* + \beta)] + \kappa/2} \\ \beta &= \frac{-ig|\alpha|^2}{i\omega_{\text{m}} + \gamma_{\text{i}}/2}. \end{aligned} \tag{2.6}$$

Similar equations for the fluctuations of the optical and mechanical mode can be found giving,

$$\begin{aligned} \dot{\hat{a}} &= -(i\Delta + \kappa/2)\hat{a} - iG(\hat{b}^\dagger + \hat{b}) - \sqrt{\eta\kappa}\hat{a}_{\text{in}} - \sqrt{\kappa'}\hat{a}_{\text{in},\text{i}} \\ \dot{\hat{b}} &= -(\omega_{\text{m}} + \gamma_{\text{i}}/2)\hat{b} - iG(\hat{a}^\dagger + \hat{a}) - \sqrt{\gamma_{\text{i}}}\hat{b}_{\text{in}}, \end{aligned} \tag{2.7}$$

where  $G = g|\alpha|$  is the enhanced/parametric optomechanical coupling rate, proportional to the square-root of the intracavity photon number ( $\sqrt{n_{\text{c}}} = |\alpha|$ ).

We now have a set of linear differential equations and can use the tools available for linear, time-invariant system (transfer functions). To do this we begin by transforming our equations to

the frequency domain via the Fourier transform (see Appendix A.1),

$$\begin{aligned} -i\omega\hat{a}(\omega) &= -(i\Delta + \kappa/2)\hat{a}(\omega) - iG(\hat{b}^\dagger(\omega) + \hat{b}(\omega)) - \sqrt{\eta\kappa}\hat{a}_{\text{in}}(\omega) - \sqrt{\kappa'}\hat{a}_{\text{in},i}(\omega) \\ -i\omega\hat{b}(\omega) &= -(\omega_m + \gamma_i/2)\hat{b}(\omega) - iG(\hat{a}^\dagger(\omega) + \hat{a}(\omega)) - \sqrt{\gamma_i}\hat{b}_{\text{in}}(\omega). \end{aligned} \quad (2.8)$$

Before moving on with the solutions to the linearized equations of motion it is instructive to consider the linearized interaction Hamiltonian. Starting with the nonlinear interaction Hamiltonian,  $H_{\text{int}}$ , and using the substitutions above we see,

$$\begin{aligned} H_{\text{int}} &= \hbar g \left[ |\alpha|^2(\beta^* + \beta) + |\alpha|^2(\hat{b}^\dagger + \hat{b}) + (\beta^* + \beta)(\alpha^*\hat{a} + \alpha\hat{a}^\dagger) + \hat{a}^\dagger\hat{a}(\beta^* + \beta) \right. \\ &\quad \left. + \hat{a}^\dagger\hat{a}(\hat{b}^\dagger + \hat{b}) \right. \\ &\quad \left. + \alpha^*\hat{a}(\hat{b}^\dagger + \hat{b}) + \alpha\hat{a}^\dagger(\hat{b}^\dagger + \hat{b}) \right] \\ &\approx \hbar g |\alpha|(\hat{a}^\dagger + \hat{a})(\hat{b}^\dagger + \hat{b}), \end{aligned} \quad (2.9)$$

where we have ignored the terms responsible for the static response (first line) and the 2nd order contribution (second line), keeping only the terms linear in  $\alpha$  which we assume to be real. We can collect pair combinations of the creation and annihilation operators and see

$$H_{\text{int}} = \hbar G(\hat{a}^\dagger\hat{b} + \hat{a}\hat{b}^\dagger) + \hbar G(\hat{a}^\dagger\hat{b}^\dagger + \hat{a}\hat{b}), \quad (2.10)$$

where the first term is called the beam-splitter-like Hamiltonian as it coherently exchanges photons for phonons at a rate  $G$ , the parametrically enhanced optomechanical coupling rate, while the second term acts to entangle photons and phonons and drive the system. It is this first term, the beam-splitter-like Hamiltonian, which is the on-resonant term when we sit a mechanical frequency red-detuned with our strong, linearizing coherent pump beam. It is this interaction which leads to the coherent exchange of photons and phonons ( $\hbar G\hat{a}^\dagger\hat{b}$ : creation of a photon and destruction of a phonon at a rate  $G$ ;  $\hbar G\hat{a}\hat{b}^\dagger$ : creation of a phonon and destruction of a photon at a rate  $G$ ) giving rise to electromagnetically induced transparency (EIT) [25, 26], the ground state cooling of the mechanical mode [27–30], as well as the physics of coherent wavelength conversion in a system with two optical modes [31–33].

### 2.2.2 Optomechanical Modification of the Mechanical Frequency and Damping–Weak Coupling Approximation

We can simplify the equations and gain powerful insight into the modification to the mechanical mode due to the optomechanical coupling by solving for  $\hat{a}(\omega)$  while ignoring the input operators

( $\hat{a}_{in}(\omega)$  and  $\hat{a}_{in,i}(\omega)$ ) giving,

$$\hat{a}(\omega) = \frac{-iG(\hat{b}^\dagger(\omega) + \hat{b}(\omega))}{i(\Delta - w) + \kappa/2} \quad \hat{a}^\dagger(\omega) = \frac{iG(\hat{b}^\dagger(\omega) + \hat{b}(\omega))}{-i(\Delta + w) + \kappa/2}, \quad (2.11)$$

using the fact that  $(\hat{A}(\omega))^\dagger = \hat{A}^\dagger(-\omega)$  (see Appendix A.1). Substituting these into  $\hat{b}(\omega)$  and ignoring  $\hat{b}_{in}(\omega)$  for now gives

$$-i\omega\hat{b}(\omega) = -i\omega_m\hat{b}(\omega) - (\gamma_i/2)\hat{b}(\omega) - iG\left(\frac{iG(\hat{b}^\dagger(\omega) + \hat{b}(\omega))}{-i(\Delta + w) + \kappa/2} + \frac{-iG(\hat{b}^\dagger(\omega) + \hat{b}(\omega))}{i(\Delta - w) + \kappa/2}\right). \quad (2.12)$$

The spectrum of  $\hat{b}(\omega)$  is centered around  $+\omega_m$ , and as such, the contributions from  $\hat{b}^\dagger(\omega)$ , which are centered around  $-\omega_m$ , can be ignored (we work in the regime where  $\omega_m \gg \gamma_i$ , i.e.  $Q_m \gg 1$ ). This allows us then to collect and identify the real and imaginary contributions from the optomechanical coupling, allowing us to write

$$-i\omega\hat{b}(\omega) = -i(\omega_m + \delta\omega_m)\hat{b}(\omega) - (\gamma_i/2 + \gamma_{OM}/2)\hat{b}(\omega) \quad (2.13)$$

where

$$\begin{aligned} \delta\omega_m &= |G|^2 \text{Im} \left( \frac{1}{i(\Delta - \omega_m) + \kappa/2} - \frac{1}{-i(\Delta + \omega_m) + \kappa/2} \right) \\ &= -|G|^2 \left( \frac{\Delta - \omega_m}{(\Delta - \omega_m)^2 + (\kappa/2)^2} + \frac{\Delta + \omega_m}{(\Delta + \omega_m)^2 + (\kappa/2)^2} \right) \\ \gamma_{OM} &= 2|G|^2 \text{Re} \left( \frac{1}{i(\Delta - \omega_m) + \kappa/2} - \frac{1}{-i(\Delta + \omega_m) + \kappa/2} \right) \\ &= 2|G|^2 \left( \frac{\kappa/2}{(\Delta - \omega_m)^2 + (\kappa/2)^2} - \frac{\kappa/2}{(\Delta + \omega_m)^2 + (\kappa/2)^2} \right). \end{aligned} \quad (2.14)$$

The first term can be seen as an optical modification to the mechanical resonance frequency, the so-called optical spring. The second term is a modification of the damping of the mechanical mode and this optical damping term can give rise to extra damping (cooling) and anti-damping (amplification) of the resonator. In general, we will be interested in what happens when we have a large coherent drive red-detuned a mechanical frequency from the optical mode, that is  $\Delta = \omega_m$ , and our system is sideband resolved ( $\omega_m \gg \kappa$ ). In this case the optical spring is negligible compared to the frequency of the mechanical resonator and the optomechanical damping simplifies to,

$$\gamma_{OM} = \frac{4|G|^2}{\kappa} \quad \text{for } \Delta \approx \omega_m. \quad (2.15)$$




Figure 2.2: **System model for optomechanical wavelength conversion.** Two optical cavities  $\hat{a}_1$  and  $\hat{a}_2$  are coupled to the same, co-localized mechanical mode  $\hat{b}$  with strength  $G_1$  and  $G_2$  respectively. The intrinsic loss rate for the optical ( $\kappa_{i,1}$  and  $\kappa_{i,2}$ ) and mechanical modes ( $\gamma_i$ ), along with the optical modes' extrinsic coupling ( $\kappa_{e,1}$  and  $\kappa_{e,2}$ ) which provide a means of getting the input signals in ( $\hat{a}_{in,k}$ ) and out ( $\hat{a}_{out,k}$ ).

## 2.3 Theory of Optomechanical Wavelength Conversion

We begin by considering the optomechanical interaction between two distinct optical modes (indexed  $k = 1, 2$ ) coupled to a shared mechanical mode with annihilation operators  $\hat{a}_k$  and  $\hat{b}$  respectively and shown in Figure 2.2. The Hamiltonian describing this system is,

$$\hat{H} = \sum_k \hbar \Delta_k \hat{a}_k^\dagger \hat{a}_k + \hbar \omega_m \hat{b}^\dagger \hat{b} + (\hat{b} + \hat{b}^\dagger) \sum_k \hbar g_k \hat{a}_k^\dagger \hat{a}_k. \quad (2.16)$$

Each optical mode has a frequency  $\omega_{o,k}$ , and is driven by a laser at frequency  $\omega_{\ell,k}$ . The Hamiltonian above is written in the interaction picture with  $\Delta_k = \omega_{o,k} - \omega_{\ell,k}$  as above, and more details in Appendix B.2.2.

As shown above in Section 2.2.1, driving the system with a strong coherent field a mechanical frequency red-detuned,  $\Delta_k \cong \omega_m$  (see Figure 2.3), causes an effective beam splitter interaction to take place between the mechanical mode and each optical mode at an enhanced coupling rate  $G_k = g_k |\alpha_k|$ , where  $\alpha_k$  is the square root of the photon occupation in optical mode  $k$ , that is  $\alpha_k = \sqrt{n_{c,k}}$ . As long as this coupling is weak with respect to the optical linewidths  $\kappa_k$  ( $G_k \ll \kappa_k$ ), an adiabatic elimination of the optical cavities results in new effective mechanical loss rates  $\gamma_{OM,k}$  into each of the  $k$  optical degrees of freedom in the system. This “loss” can provide an effective coupling between the optical modes by allowing the exchange of excitations between the optical resonances through the mechanical motion of the system. We calculate exactly what these conversion rates are by using a scattering matrix formulation to understand the behaviour of the system. Though






Figure 2.3: **Pumping scheme for wavelength conversion.** Schematic showing the two optical modes of the system with the blue optical mode representing the higher optical frequency ( $\omega_{o,1} > \omega_{o,2}$ ). Coherent driving tones with strength  $\alpha_k$  are tuned a mechanical frequency red of both cavities ( $\omega_{\ell,k} = \omega_{o,k} - \omega_m$ ), allowing the interaction to be linearized and giving rise to the beam splitter like interaction Hamiltonian.

completely general expressions can be derived, to best understand the processes involved we focus on the parameter regime relevant to this experiment, i.e., the weak-coupling, sideband-resolved case where  $\omega_m \gg \kappa \gg \gamma_{OM}$ , and follow a scattering matrix derivation of the induced optomechanical coupling between the optical waveguides.

With strong coherent driving tones set a mechanical frequency red-detuned, the interaction Hamiltonian can be linearized as mentioned in Section 2.2.1, giving rise to a beam splitter like interaction between the two optical modes and common co-localized mechanical mode,

$$H_{\text{int}} = \hbar G_1(\hat{a}_1^\dagger \hat{b} + \hat{a}_1 \hat{b}^\dagger) + \hbar G_2(\hat{a}_2^\dagger \hat{b} + \hat{a}_2 \hat{b}^\dagger). \quad (2.17)$$

The essential physics of the process giving rise to the wavelength conversion can be seen already in this equation. Here, this interaction allows photons in cavity mode 1 to be destroyed to coherently create phonons in the mechanical mode at a rate  $G_1$ . We can also see that phonons can be destroyed to coherently create photons in cavity mode 2 at a rate  $G_2$ , which is the basis for the wavelength conversion scheme. This process is completely symmetric and we can see the opposite is also true, allowing photons to be converted from mode 2 to mode 1 in exactly the same manner.

The linearized Heisenberg-Langevin equations of motion are found using Equation (2.16), the displacement  $\hat{a}_k \rightarrow \alpha_k + \hat{a}_k$ , and the inclusion of the optical input signal ( $\hat{a}_{\text{in},k}(t)$ ), optical vacuum

noise ( $\hat{a}_{i,k}(t)$ ), and mechanical thermal fluctuation fields ( $\hat{b}_{\text{in}}(t)$ ):

$$\begin{aligned}\dot{\hat{b}}(t) &= -\left(i\omega_m + \frac{\gamma_i}{2}\right)\hat{b} - i\sum_k G_k(\hat{a}_k + \hat{a}_k^\dagger) - \sqrt{\gamma_i}\hat{b}_{\text{in}}(t) \\ \dot{\hat{a}}_k(t) &= -\left(i\Delta_k + \frac{\kappa_k}{2}\right)\hat{a}_k - iG_k(\hat{b} + \hat{b}^\dagger) \\ &\quad - \sqrt{\kappa_{e,k}/2}\hat{a}_{\text{in},k}(t) - \sqrt{\kappa'_k}\hat{a}_{i,k}(t) \quad (k = 1, 2).\end{aligned}$$

The mechanical loss rate  $\gamma_i = \omega_m/Q_m$  determines the coupling of the system to the thermal bath, and provides the most significant contribution in terms of noise processes relevant to the performance of the optomechanical wavelength converter. The other loss rates,  $\kappa_k$ ,  $\kappa_{e,k}$ , and  $\kappa'_k$ , are the optical loss rates for all photons leaving cavity  $k$ , photons leaving cavity  $k$  through the coupling into a two-sided waveguide, and all photons leaving the cavity through channels that are undetected representing a loss of information. Due to the local evanescent side-coupling scheme studied here,  $\kappa_k = \kappa_{e,k}/2 + \kappa'_k$ , i.e., only a fraction  $\eta_k \equiv \kappa_{e,k}/2\kappa_k$  of the photons leaving the cavity get detected and  $\eta_k \leq 1/2$ . A comparison of single and double sided coupling can be found in Appendix B.2.1.

The equations of motion are linear (we linearized them above) and, as was done in Section 2.2.1, the system can be analyzed more simply in the frequency domain. To begin, we first look at the spectrum of the equations of motion for  $\hat{b}$  and  $\hat{a}_k$  in the frequency domain,

$$\begin{aligned}\hat{b}(\omega) &= \frac{-\sqrt{\gamma_i}\hat{b}_{\text{in}}(\omega)}{i(\omega_m - \omega) + \gamma/2} + \sum_k \frac{-iG_k(\hat{a}_k(\omega) + \hat{a}_k^\dagger(\omega))}{i(\omega_m - \omega) + \gamma_i/2} \\ \hat{a}_k(\omega) &= \frac{-iG_k(\hat{b}(\omega) + \hat{b}^\dagger(\omega))}{i(\Delta_k - \omega) + \kappa_k/2} - \frac{\sqrt{\kappa_{e,k}/2}\hat{a}_{\text{in},k}(\omega) + \sqrt{\kappa'_k}\hat{a}_{i,k}(\omega)}{i(\Delta_k - \omega) + \kappa_k/2}.\end{aligned}\quad (2.18)$$

Solving for the spectrum of the mechanical mode  $\hat{b}(\omega)$  in terms of the input noise operators, we find:

$$\begin{aligned}\hat{b}(\omega) &= \frac{-\sqrt{\gamma_i}\hat{b}_{\text{in}}(\omega)}{i(\omega_m - \omega) + \gamma/2} \\ &\quad + \sum_k \frac{iG_k}{i(\Delta_k - \omega) + \kappa_k/2} \frac{\sqrt{\kappa_{e,k}/2}\hat{a}_{\text{in},k}(\omega) + \sqrt{\kappa'_k}\hat{a}_{i,k}(\omega)}{i(\omega_m - \omega) + \gamma/2} \\ &\quad + \sum_k \frac{iG_k}{-i(\Delta_k + \omega) + \kappa_k/2} \frac{\sqrt{\kappa_{e,k}/2}\hat{a}_{\text{in},k}^\dagger(\omega) + \sqrt{\kappa'_k}\hat{a}_{i,k}^\dagger(\omega)}{i(\omega_m - \omega) + \gamma/2}\end{aligned}$$

where the mechanical frequency  $\omega_m$  is now modified by the optical spring, and the mechanical linewidth is given by  $\gamma \equiv \gamma_i + \gamma_{\text{OM},1} + \gamma_{\text{OM},2}$  as was done in Section 2.2.2. The important term to remember is the optomechanical damping terms  $\gamma_{\text{OM},k}$  which come from coupling of the mechanical

system to the optical mode  $k$ , and are given by the relation

$$\begin{aligned}\gamma_{\text{OM},k} &= 2|G_k|^2 \text{Re} \left[ \frac{1}{i(\Delta_k - \omega_m) + \kappa_k/2} - \frac{1}{-i(\Delta_k + \omega_m) + \kappa_k/2} \right] \\ &= \frac{4|G_k|^2}{\kappa_k} \quad \text{for } \Delta \approx \omega_m.\end{aligned}\tag{2.19}$$

The last expression is a simplification that is often made, and is equivalent to looking at spectral properties at detunings from the mechanical frequency much smaller than  $\kappa$  (so the optical lineshape isn't taken into account). This approximation is quite valid as the mechanical mode will spectrally filter the wavelength conversion process with a bandwidth of  $\gamma$ , which is much smaller than  $\kappa$  here. The full equation for  $\hat{a}_k(\omega)$  is then,

$$\begin{aligned}\hat{a}_k(\omega) &= \frac{iG_k\sqrt{\gamma_i}\hat{b}_{\text{in}}(\omega)}{[i(\Delta_k - \omega) + \kappa_k/2][i(\omega_m - \omega) + \gamma/2]} \\ &+ \sum_j \frac{G_k G_j}{[i(\Delta_k - \omega) + \kappa_k/2][i(\Delta_j - \omega) + \kappa_j/2]} \frac{\sqrt{\kappa_{e,j}/2}\hat{a}_{\text{in},j}(\omega) + \sqrt{\kappa'_j}\hat{a}_{i,j}(\omega)}{i(\omega_m - \omega) + \gamma/2} \\ &+ \sum_j \frac{G_k G_j}{[i(\Delta_k - \omega) + \kappa_k/2][-i(\Delta_j + \omega) + \kappa_j/2]} \frac{\sqrt{\kappa_{e,j}/2}\hat{a}_{\text{in},j}^\dagger(\omega) + \sqrt{\kappa'_j}\hat{a}_{i,j}^\dagger(\omega)}{i(\omega_m - \omega) + \gamma/2} \\ &+ \frac{iG_k\sqrt{\gamma_i}\hat{b}_{\text{in}}^\dagger(\omega)}{[i(\Delta_k - \omega) + \kappa_k/2][-i(\omega_m + \omega) + \gamma/2]} \\ &+ \sum_j \frac{-G_k G_j^*}{[i(\Delta_k - \omega) + \kappa_k/2][i(\Delta_j - \omega) + \kappa_j/2]} \frac{\sqrt{\kappa_{e,j}/2}\hat{a}_{\text{in},j}(\omega) + \sqrt{\kappa'_j}\hat{a}_{i,j}(\omega)}{-i(\omega_m + \omega) + \gamma/2} \\ &+ \sum_j \frac{-G_k G_j^*}{[i(\Delta_k - \omega) + \kappa_k/2][-i(\Delta_j + \omega) + \kappa_j/2]} \frac{\sqrt{\kappa_{e,j}/2}\hat{a}_{\text{in},j}^\dagger(\omega) + \sqrt{\kappa'_j}\hat{a}_{i,j}^\dagger(\omega)}{-i(\omega_m + \omega) + \gamma/2} \\ &- \frac{\sqrt{\kappa_{e,k}/2}\hat{a}_{\text{in},k}(\omega) - \sqrt{\kappa'_k}\hat{a}_{i,k}(\omega)}{[i(\Delta_k - \omega) + \kappa_k/2]}.\end{aligned}\tag{2.20}$$

We will be setting the coherent drive tones at mechanical frequency red-detuned from the optical cavities, that is  $\Delta_k = \Delta_j = \omega_m$ . The fourth, fifth, and sixth terms arise from the coupling to  $\hat{b}^\dagger(\omega)$  in Equation 2.18 and are centered around  $-\omega_m$ , which we can immediately discard as the optical mode spectrally filters out this contribution. That is, the optical mode  $\hat{a}_k$  is centered at  $\omega_m$  (recall we moved to a rotating frame and are looking at the slowly varying equations of motion for  $\hat{a}_k$ ) with a linewidth of  $\kappa_k \ll \omega_m$ —we are in the good cavity (sideband resolved) limit. In the remaining terms, we can see with  $\Delta_k = \omega_m$ , the filtering of cavity  $\hat{a}_k$  is  $i(\Delta_k - \omega) + \kappa_k/2 \rightarrow i(\omega_m - \omega) + \kappa_k/2 \rightarrow \kappa_k/2$  since we will only be looking at detunings  $(\omega_m - \omega)$  on the order of  $\gamma \ll \kappa_k$ . Similarly, the spectral filtering of cavity mode  $\hat{a}_j$  can be approximated by  $i(\Delta_j - \omega) + \kappa_j/2 \rightarrow \kappa_j/2$  for the same reason. The other term becomes  $-i(\Delta_j + \omega) + \kappa_j/2 \rightarrow -2i\omega_m$  as again,  $\omega_m \gg \kappa_j$ . Under this approximation,

the optical spectrum is

$$\begin{aligned}
\frac{\kappa_k}{2}\hat{a}_k(\omega) &= \frac{iG_k\sqrt{\gamma_i}\hat{b}_{\text{in}}(\omega)}{i(\omega_{\text{m}} - \omega) + \gamma/2} \\
&+ \sum_j \frac{2G_jG_k}{\kappa_j} \frac{\sqrt{\kappa_{e,j}/2}\hat{a}_{\text{in},j}(\omega) + \sqrt{\kappa'_j}\hat{a}_{\text{i},j}(\omega)}{i(\omega_{\text{m}} - \omega) + \gamma/2} \\
&+ \sum_j \frac{iG_jG_k}{2\omega_{\text{m}}} \frac{\sqrt{\kappa_{e,j}/2}\hat{a}_{\text{in},j}^\dagger(\omega) + \sqrt{\kappa'_j}\hat{a}_{\text{i},j}^\dagger(\omega)}{i(\omega_{\text{m}} - \omega) + \gamma/2} \\
&- \sqrt{\kappa_{e,k}/2}\hat{a}_{\text{in},k}(\omega) - \sqrt{\kappa'_k}\hat{a}_{\text{i},k}(\omega).
\end{aligned} \tag{2.21}$$

From this expression, we see that there are several noise operators incident on optical mode  $k$ . The thermal fluctuations from the environment,  $\hat{b}_{\text{in}}(\omega)$ , are converted into noise photons over the mechanical bandwidth  $\gamma$ . There is also an induced coupling to the optical mode  $j \neq k$ , and photons originally incident only on mode  $j$ , i.e.,  $\hat{a}_{\text{in},j}(\omega)$ , are now also coupled into  $\hat{a}_k$ . Finally we note that vacuum noise creation operators are also present in this expression. These give rise to quantum noise through spontaneous emission of phonons, though we show below that, as  $\omega_{\text{m}}$  is made very large with respect to  $\kappa_k$ , these terms diminish in importance.

From outside the optomechanical system we only have access to photons sent into the system,  $\hat{a}_{\text{in},j}$ , and those leaving the system,  $\hat{a}_{\text{out},k}$ , on transmission. The relation between these operators is best understood through scattering parameters, and can be derived using the equation for  $\hat{a}_k(\omega)$  (Equation 2.21) and the input-output boundary conditions  $\hat{a}_{\text{out},k}(\omega) = \hat{a}_{\text{in},k}(\omega) + \sqrt{\kappa_{e,k}/2}\hat{a}_k(\omega)$ . After some algebra, the output operator is expressed as

$$\begin{aligned}
\hat{a}_{\text{out},k}(\omega) &= s_{\text{th},k}(\omega)\hat{b}_{\text{in}}(\omega) \\
&+ t_k(\omega)\hat{a}_{\text{in},k}(\omega) + s_{kj}(\omega)\hat{a}_{\text{in},j}(\omega) \\
&+ \sum_m n_{\text{opt},m}(\omega)\hat{a}_{\text{i},m}(\omega) \\
&+ \sum_m s_{\text{adj},\text{in},m}(\omega)\hat{a}_{\text{in},m}^\dagger(\omega) \\
&+ \sum_m s_{\text{adj},\text{i},m}(\omega)\hat{a}_{\text{i},m}^\dagger(\omega).
\end{aligned} \tag{2.22}$$

The scattering coefficient  $s_{\text{th},k}(\omega)$  is the conversion efficiency of mechanical thermal noise to photons, and is given by

$$s_{\text{th},k}(\omega) = i\sqrt{\eta_k} \frac{\sqrt{\gamma_{\text{OM},k}\gamma_i}}{i(\omega_{\text{m}} - \omega) + \gamma/2}. \tag{2.23}$$

We can understand this term as follows: Thermal noise is coupled into the mechanical mode at a

rate  $\gamma_i$  as thermal noise phonons. These phonons are then converted to photons in optical mode  $\hat{a}_k$  at a rate  $\gamma_{\text{OM},k}$ , and these photons are coupled into the waveguide to be collected for detection with an efficiency  $\eta_k$ . The entire process is spectrally filtered by the mechanical mode, as can be seen from the denominator in Equation 2.23. These “thermal” photons represent a source of noise in the conversion process and will be discussed later.

The coefficient  $t_k(\omega)$  in our system is the transmission amplitude, given by

$$t_k(\omega) = (1 - 2\eta_k) + \eta_k \frac{\gamma_{\text{OM},k}}{i(\omega_m - \omega) + \gamma/2}. \quad (2.24)$$

Here we see photons from the waveguide are coupled into optical mode  $\hat{a}_k$  with efficiency  $\eta_k$  and converted into phonons at a rate  $\gamma_{\text{OM},k}$ . These phonons are then converted back into photons at the same rate and coupled back into the waveguide with the same efficiency. The process again has the bandwidth of the mechanical mode. Since this entire conversion process is coherent, these photons interfere with photons that bypassed the optical cavity and photons that coupled into the optical cavity but were not converted to phonons. In principle, this coefficient is the electromagnetically induced transparency transmission coefficient, though here it is written about the mechanical frequency for detunings on the order of  $\gamma$ , and the  $\kappa$ 's are too large ( $\kappa \gg \gamma$ ) for the optical lineshape to be considered in the expression. This is the expected result of the weak coupling approximation [26].

Finally, the most important coefficient is the wavelength conversion coefficient  $s_{kj}(\omega)$ , which is given by

$$s_{kj}(\omega) = \sqrt{\eta_k \eta_j} \frac{\sqrt{\gamma_{\text{OM},k} \gamma_{\text{OM},j}}}{i(\omega_m - \omega) + \gamma/2}. \quad (2.25)$$

This term can be similarly understood. Photons are coupled into optical mode  $\hat{a}_j$  from the waveguide with an efficiency  $\eta_j$  and converted into phonons at a rate  $\gamma_{\text{OM},j}$ . These phonons are then converted into photons in cavity mode  $\hat{a}_k$  at a rate  $\gamma_{\text{OM},k}$ , and coupled into the waveguide to be detected with an efficiency  $\eta_k$ . The entire process is filtered spectrally by the mechanical mode. For the moment we will ignore the quantum noise terms.

### 2.3.1 Efficiency, Bandwidth, and Noise

We calculate the spectral density of the output field on port  $k$  ( $S_{\text{out},k}(\omega)$ ), assuming an input field spectral density on the opposing optical port  $j$  ( $S_{\text{in},j}(\omega)$ ), vacuum inputs on all other optical channels, and thermal noise from a phonon bath with thermal occupation  $n_b$ . The spectral densities here have units of photons/Hz  $\cdot$  s and can be interpreted as the photon flux per unit bandwidth. At first we ignore the field creation operators in the scattering relation (2.22) (which give rise to quantum noise and can be neglected when the mechanical resonator occupation number is greater

than one as will be shown later) and arrive at the following expression:

$$\begin{aligned}
S_{\text{out},k}(\omega) &= \int_{-\infty}^{\infty} d\omega' \langle \hat{a}_{\text{out},k}^{\dagger}(\omega) \hat{a}_{\text{out},k}(\omega') \rangle \\
&= \int_{-\infty}^{\infty} s_{\text{th},k}^*(-\omega) s_{\text{th},k}(\omega') \langle \hat{b}_{\text{in}}^{\dagger}(\omega) \hat{b}_{\text{in}}(\omega') \rangle \\
&\quad + t_k^*(-\omega) t_k(\omega') \langle \hat{a}_{\text{in},k}^{\dagger}(\omega) \hat{a}_{\text{in},k}(\omega') \rangle \\
&\quad + s_{kj}^*(-\omega) s_{kj}(\omega') \langle \hat{a}_{\text{in},j}^{\dagger}(\omega) \hat{a}_{\text{in},j}(\omega') \rangle \\
&\quad + n_{\text{opt},k}^*(-\omega) n_{\text{opt},k}(\omega') \langle \hat{a}_{i,k}^{\dagger}(\omega) \hat{a}_{i,k}(\omega') \rangle \\
&\quad + n_{\text{opt},j}^*(-\omega) n_{\text{opt},j}(\omega') \langle \hat{a}_{i,j}^{\dagger}(\omega) \hat{a}_{i,j}(\omega') \rangle d\omega'.
\end{aligned}$$

Using the definitions for the noise operator correlations (both mechanical and optical, see Appendix A.2), this reduces to

$$\begin{aligned}
S_{\text{out},k}(\omega) &= \eta_k \frac{\gamma_i \gamma_{\text{OM},k}}{(\omega + \omega_{\text{m}})^2 + (\gamma/2)^2} n_{\text{b}} \\
&\quad + \eta_k \eta_j \frac{\gamma_{\text{OM},k} \gamma_{\text{OM},j}}{(\omega + \omega_{\text{m}})^2 + (\gamma/2)^2} S_{\text{in},j}(\omega).
\end{aligned} \tag{2.26}$$

From here, we see that the maximum photon wavelength conversion efficiency is given by

$$\eta_{\text{max}} = \eta_j \eta_k \frac{4\gamma_{\text{OM},j} \gamma_{\text{OM},k}}{(\gamma_i + \gamma_{\text{OM},j} + \gamma_{\text{OM},k})^2}. \tag{2.27}$$

The conversion process has a bandwidth (BW) equal to the optomechanically damped (cooled) mechanical linewidth as it is the most spectrally selective system involved

$$\text{BW} \equiv \gamma_i + \gamma_{\text{OM},j} + \gamma_{\text{OM},k}. \tag{2.28}$$

From the same expression (Equation 2.26), we calculate a value for the optical signal to noise ratio (OSNR), to find

$$\text{OSNR}_{kj}^{\text{classical}} = \eta_j \frac{\gamma_{\text{OM},j}}{\gamma_i n_{\text{b}}} S_{\text{in},j}(\omega). \tag{2.29}$$

This expression only takes into account the thermal noise in the system and therefore holds only when  $\gamma_i n_{\text{b}} / \gamma_{\text{OM},j} \gg 1$ , that is, when the mechanical mode population is much greater than one ( $\langle n \rangle \gg 1$ ). At higher powers, the quantum noise processes which give rise to the back-action limit in cooling [27, 28, 30, 34] also give rise to excess noise for wavelength conversion. To consider these terms, the creation operators in Equation (2.22) must not be neglected, and the full relation is then

found to be

$$\begin{aligned}
S_{\text{out},k}(\omega) = & \eta_k \frac{\gamma_{\text{OM},k}}{(\omega + \omega_m)^2 + (\gamma/2)^2} \gamma_i n_b \\
& + \eta_k \frac{\gamma_{\text{OM},k}}{(\omega + \omega_m)^2 + (\gamma/2)^2} \gamma_{\text{OM},k} \left( \frac{\kappa_k}{4\omega_m} \right)^2 \\
& + \eta_k \frac{\gamma_{\text{OM},k}}{(\omega + \omega_m)^2 + (\gamma/2)^2} \gamma_{\text{OM},j} \left( \frac{\kappa_j}{4\omega_m} \right)^2 \\
& + \eta_k \eta_j \frac{\gamma_{\text{OM},k} \gamma_{\text{OM},j}}{(\omega + \omega_m)^2 + (\gamma/2)^2} S_{\text{in},j}(\omega).
\end{aligned} \tag{2.30}$$

The quantum limited OSNR is then found to be

$$\text{OSNR}_{kj} = \frac{\eta_j \gamma_{\text{OM},j} S_{\text{in},j}(\omega)}{\gamma_i n_b + \gamma_{\text{OM},k} \left( \frac{\kappa_k}{4\omega_m} \right)^2 + \gamma_{\text{OM},j} \left( \frac{\kappa_j}{4\omega_m} \right)^2}. \tag{2.31}$$

For the quantum back-action terms to become significant as compared to the thermal noise on the mechanical system, one needs the terms in the denominator of Equation 2.31 to be of similar magnitude. Specifically,  $\gamma_i n_b / \gamma \sim \left( \frac{\kappa}{4\omega_m} \right)^2$ . This is a regime that is yet to be reached in experiments, and as such, the quantum back-action noise terms can be neglected for experiments to date. Perhaps more importantly, we ask, “What is the amount of noise added to a signal of a single photon?” We can express the signal to noise ratios as  $\text{OSNR}_{kj} = S_{\text{in},j}(\omega) / n_{\text{added}}$ , with

$$n_{\text{added}} = \frac{\gamma_i n_b + \gamma_{\text{OM},k} \left( \frac{\kappa_k}{4\omega_m} \right)^2 + \gamma_{\text{OM},j} \left( \frac{\kappa_j}{4\omega_m} \right)^2}{\eta_j \gamma_{\text{OM},j}}. \tag{2.32}$$

Here we can see the terms in the numerator correspond to thermal phonons in the mechanical mode, as well as photons from both pump beams scattered a mechanical frequency red, generating phonons. This last process is filtered by the cavity as the photons are scattered into the tail of the optical cavity as opposed to the center of the cavity, and is correspondingly much less likely. The scattering rates  $\gamma_i$ ,  $\gamma_{\text{OM},j}$ , and  $\gamma_{\text{OM},k}$  are the rates a single quanta of noise is generated and the terms multiplying those rates represent the source of the population (i.e., thermal bath and scattered photons due to imperfect sideband resolution). The term on the denominator of Equation 2.32 is the input coupling efficiency of photons from the waveguide to the optical mode, and the efficiency to convert those photons to phonons (in this case a single photon). Assuming  $\kappa_j = \kappa_k = \kappa$ ,  $\gamma_i \ll \gamma$ , and  $\gamma_{\text{OM},j} = \gamma_{\text{OM},k} = \gamma_{\text{OM}}$ , this expression becomes

$$n_{\text{added}} \approx 2\eta_j^{-1} \left( \frac{\gamma_i n_b}{\gamma} + \left( \frac{\kappa}{4\omega_m} \right)^2 \right). \tag{2.33}$$

To achieve the threshold where a single photon incoming on port  $j$  is converted, and is more likely to be detected on port  $k$  than an upconverted thermal excitation, we require  $n_{\text{added}} < 1$ . This requirement is equivalent to stating that in the absence of signals on port  $j$  and  $k$ , a cooled phonon occupation of  $\bar{n} = \eta_j/2$  must be achieved. We note that  $\eta_k$  does not make an appearance in this equation, since reduced output coupling efficiency attenuates the thermal noise and signal equivalently. For the same reason,  $1/\eta_j$  is present in the noise quanta expression. The factor of two can be understood to come from the fact that for good conversion, we require the photon to enter and exit the system before a phonon is upconverted. The former two processes happen at  $\gamma_{\text{OM}}$ , while the latter is given by the thermalization rate  $\gamma_i n_{\text{b}}$ .



## Chapter 3

# Fabrication of Nanobeam Optomechanical Crystals

The process of fabrication of micro and nano-scale devices is iterative, starting with design and ending with characterization of the properties of the device. These properties, when compared to the desired results, can initiate improvements at any of the steps along the way: design, electron beam lithography, etching, all the way to the storage and transport of the device. Due to the scales of the devices we wish to make, the methods available to us can be restrictive, blunt (sledgehammer not scalpel), and somewhat mysterious in their specific action. Combined with our inherent difficulties in applying the intuition we have from everyday length scales to these devices, and their sensitivities to all sorts of perturbations we can't see, it can often be difficult to know where a change needs to be made in the fabrication steps. For this reason it is very important to attempt to do everything along the way exactly the same, to control all the parameters available since there are so many that are not, and since there are some things we don't even know are parameters. This also tends to lead to the notion that fabrication is a bit like black magic. When something is determined to work, through either intense, systematic study and trial and error, or serendipity, the procedure is codified and never changed, often with the original reasoning lost (if there ever was any, and if it was even correct). The following chapter is an attempt to codify some of the fabrication process used in the Painter lab to fabricate the optomechanical nanobeams that have been the workhorse for much of the science done in our group over the last several years (including the wavelength conversion described in this thesis), and the device I have spent much time fabricating. Many of the techniques and process steps existed before I came; some had justifications and reasons, while others just were. The focus of this chapter will be an attempt to provide an overview for the entire fabrication process, from design to a ready-to-characterize sample. Furthermore, an emphasis will be placed on the reasons behind the steps and suggestions for improvements (or likely possible improvements which would require systematic study). Nothing can replace being instructed first-hand by someone in a clean room showing, teaching, and then overseeing the fabrication, but the goal here will be to

provide a theoretical, codified introduction to our processes. Table 3.1 provides a brief overview of the steps and a description of each, while Figure 3.1 shows the major steps schematically.

| Fabrication Step | Input                    | Description of Process                                                                                                                                                                                                      | Output              |
|------------------|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| Design           | an idea                  | simulation                                                                                                                                                                                                                  | design (.dxf file)  |
| File fracture    | .dxf file                | fracture a geometry into shapes the EBPG can draw                                                                                                                                                                           | .gpf file           |
| Spin and bake    | chip                     | clean a diced chip and spin on E-Beam resist and bake                                                                                                                                                                       | resist coated chip  |
| Beam write       | .gpf, resist coated chip | use the EBPG to expose the resist on the chip with the pattern you designed                                                                                                                                                 | exposed chip        |
| Develop          | exposed chip             | develop the E-Beam resist removing material exposed to electrons leaving the patterned mask                                                                                                                                 | developed chip      |
| RIE/ICP etch     | developed chip           | <ul style="list-style-type: none"> <li>plasma etch the chip transferring the pattern from the mask into the device layer</li> <li>areas with no mask will be etched away and the mask will also be etched slowly</li> </ul> | etched chip         |
| Remove resist    | etched chip              | remove the bulk of the resist with ZDMAC (or TCE) and the remaining resist with piranha                                                                                                                                     | resist striped chip |
| Undercut         | resist striped chip      | use HF to remove the buried oxide the device layer sits on                                                                                                                                                                  | undercut chip       |
| Clean            | undercut chip            | clean the surface with Piranha/HF(diluted) cycles                                                                                                                                                                           | cleaned chip        |
| Characterize     | cleaned chip             | determine the chips properties and compare to desired                                                                                                                                                                       | data to iterate     |

Table 3.1: **Overview of fabrication process.** A brief overview of the steps involved in fabricating a nanobeam along with the nominal inputs and outputs of each step and a description of the tasks involved. See Figure 3.1 for a pictorial representation.

### 3.1 Design

The zeroth step in fabrication is the coming up with the idea of what is to be designed and fabricated. The first step is to design the device that needs to be fabricated. Much of the design work in our group involves optical and mechanical simulations, the majority of which are finite element method (FEM) simulations done by COMSOL [35]. With respect to the design of the optomechanical structures we use in our group, including the nanobeam used in the work in this thesis, there are many resources available [22, 36–39]. Other tools used in conjunction with COMSOL include MIT




Figure 3.1: **Overview of silicon processing for the nanobeams.** Schematic highlighting the major steps in fabrication of a silicon nanobeam. From left to right in order of operation: deposition of the resist layer; exposing the pattern with electron beam lithography; develop the pattern removing expose materials and rinse the developer from the sample to stop the development process; pseudo-Bosch etch to transfer the pattern from the resist layer into the silicon device layer; piranha clean to remove all the resist (often preceded by a solvent clean to remove the bulk of the resist and let the piranha do the final polish); release the sacrificial layer the devices are sitting on with HF (the sacrificial layer is glass since the chip is silicon-on-insulator (SOI)). This figure is from [22].

Photonic-Bands (MPB) [40, 41] a free program for band structure calculations of photonic crystals (one-dimensional and two-dimensional), MEEP [42, 43] a free program for finite-difference time-domain (FDTD) method simulations, and Lumerical FDTD Solutions [44] a commercial FDTD package.

### 3.1.1 Design for Fabrication

The design of a structure cannot be done in a vacuum with no thought towards fabrication if the device is intended to be made. If a structure is designed and optimized to have 10 nm square holes, or requires a 10 nm gap in a 400 nm device layer of silicon nitride ( $\text{Si}_3\text{N}_4$ ), it may be theoretically interesting but nearly impossible to make. The material required and processing necessary must be taken into account at design time. The etching of squares vs. circles is one example of this. Etching very nice right angles into a material can be challenging due to the nature of the chemical and physical processes in etching (see Section 3.3). The etching process tends to round these corners and, if the design requires perfect rectangles to have low optical or acoustic scattering losses, or perhaps even work at all, this can be problematic. Figure 3.2a shows an SEM of the original nanobeam design with rectangular holes [45], or holes that should be rectangular but have been rounded in the etch. Figure 3.2b shows an SEM of the new nanobeam design which has elliptical holes [39]. It turns out that perfectly rectangular holes were not required (or easily possible to fabricate) and that rounded holes (ellipses) had better design properties and were easier to fabricate faithfully.




Figure 3.2: **Comparison of rectangular vs. elliptical holes.** (a) Wireframe geometry of a pattern to be etched with rectangular holes (from [45]). (b) Etched geometry showing rounded rectangular holes. (c) Wireframe geometry of a pattern to be etched with elliptical holes (5G29SS). (d) Etched geometry showing high shape fidelity.

Another good example is attempting to make a very small slot gap in a structure such as a zipper optomechanical cavity [46]. For these devices the optomechanical coupling is very sensitive to the slot gap, with the coupling increasing exponentially as the gap is reduced. Etching small slot gaps is possible but it is difficult to make them very small, and doing so can lead to slanted sidewalls and surface roughness in the etched gap. However, clever design can often help alleviate this problem. Figure 3.3a shows an SEM the original zipper optomechanical cavity with a slot gap of  $s = 120$  nm (see Figure 3.3b) and an optomechanical coupling of  $g_{\text{OM}}/2\pi = 123 \text{ GHz} \cdot \text{nm}^{-1}$  [46]. The attempt to make a very small slot gap has lead to the beams being asymmetrically etched away, though this can easily be compensated for with appropriate design (this was done in later iterations). Instead of attempting to etch a very small slot gap, if a clever design using tethers with widths that are less than the beam widths and offset from the center, the residual stress of the material can pull the beams together after they have been released, producing demonstrated slots as small as  $s = 40$  nm with  $g_{\text{OM}}/2\pi \approx 200 \text{ GHz} \cdot \text{nm}^{-1}$  [19]. This can be seen in Figure 3.3c,d.

The main point is that the design and fabrication cannot be completely decoupled. If the person designing a structure doesn't know the limitations of the fabrication process of different materials, they may design a very difficult structure to fabricate. Conversely, the person fabricating a difficult structure may not realize what sorts of small changes to the design might be possible to greatly increase the ease of fabrication without sacrificing device performance.

The end result of the design step is to have a parameterized geometry specified in a way that is useful to the next step, generally either simulation and/or E-Beam lithography. Since MATLAB naturally interacts with COMSOL, most of the scripting is done in MATLAB. In our case, the parameters of the device are specified and a MATLAB script can generate the required geometry commands COMSOL understands (see Jasper Chan's thesis [22]).




Figure 3.3: **Comparison of a small etched slot and pulled in beams.** (a) SEM of the entire zipper optomechanical cavity showing the etched slot gap and anchoring to the substrate (from [46]). (b) Zoom-in of the etched slot. The holes can again be seen to be rounded and are off-center due to excess blowout from trying to etch such a small slot gap. (c) Full floppy zipper structure showing the nanobeam now anchored to the substrate with long, thin tethers (from [19]). (d) The left panel shows the anchoring to the substrate. The middle panel shows the tether anchoring to the nanobeams. Note the size difference, and offset of the tether center from the nanobeam center. This offset and the residual stress of the material causes the beams to bend in towards each other when they are released. The right panel shows the final separation of the nanobeams, resulting in a very small slot gap. Both devices are made out of silicon nitride ( $\text{Si}_3\text{N}_4$ ).

## 3.2 Electron Beam Lithography

With the design and simulation of the structure complete, the next step is almost always to create an etch mask to be used in the process of transferring the pattern into the device layer. This is true for the nanobeam, and because of the size of the features required for photonic crystals, electron beam lithography (e-beam) is required (as compared to photo-lithography). The main steps here are preparing a chip with electron beam resist, preparing patterns to be written on the chip in the resist to define the etch pattern, writing those patterns into the resist with the electron beam pattern generator (EBPG), and finally developing the chip to create the etch mask. The specifics of these tasks will be elaborated upon in the following sections.

### 3.2.1 Parameter Sweeps

While the results of a simulation are good, they are not perfect. If a COMSOL simulation says that there will be an optical mode at  $\lambda = 1550$  nm with a high optical quality factor ( $Q_o$ ), there will almost certainly be a mode; however, even if the wavelength of the mode is  $\pm 2\%$ , that is  $\pm 31$  nm and the 1550 nm laser we use has a range of 1520–1570 nm. As such, it is necessary to write many

devices with varying parameters—parameter sweeps—on the chip to ensure we hit the “sweet spot” and have several devices in our testable range. For nanobeams we generally wrote 20 nanobeams with a two-dimensional, in-plane scaling of  $\pm \sim 3\text{--}7\%$  depending on our confidence of the nominal parameters (a larger range makes it more likely to have some devices in range, but not as many). This scaled the entire design of the nanobeam in-plane—every dimension except the thickness (which was set by the device layer of our material) was scaled. This scaling set the photonic and phononic properties of the nanobeam. On top of that, the lattice constant for the crosses was also scaled over three devices by  $\pm 2\%$  independently of the nanobeam scaling, giving a total of 60 unique nanobeams. When beginning new structures it is good to have a large sweep to determine where the nominal value to sweep around is. Once that is determined, the sweep can be decreased in order to increase the number of devices with wavelengths in the appropriate range.

For more complicated/involved designs, other relevant parameters would also be varied. When we began using waveguide side-coupling to the nanobeams, a third parameter to vary was now the gap between the waveguide and nanobeam as discussed in Section 5. This gap changed the coupling strengths ( $\kappa_e$ ) between the waveguide and the device. In general these devices are very small, and other than the time it takes to beamwrite a chip, the time to process a chip doesn’t scale with the number of devices on it. In the example above, with 60 different nanobeams in our parameter space, it would be wise to have this repeated many times on the chip to have redundancy (possibly with slightly different doses as discussed in Section 3.2.5).

### 3.2.2 Design Compensation

The results of the design and simulation step provide a target set of parameters to be fabricated to achieve a specific final device. However, during the lithography step, the etching step, and subsequent cleaning steps, the size of the features is modified. For instance, etching almost certainly enlarges the size of a hole and decrease the width of a beam. The cleaning steps tend to remove a small amount of material from all surfaces of the device and the lithography step can produce larger or smaller features depending on the dosing. To achieve the correct final feature sizes, this “blowout” must be understood and compensated for. Table 3.2 contains a list of design parameters, their nominal values, and values the EBPG writes in the resist, which have been corrected for any blowout. Figure 3.4 shows the unit cell for the mirror region of the nanobeam and the cross structure that makes up the phononic shield with the design parameters highlighted.

The first thing to notice is that the lattice constants don’t change. Over the range we ask the EBPG to draw shapes, we assume it to be very accurate at placing one shape a distance  $0.1\text{--}10\text{ }\mu\text{m}$  away from another. As such there is no compensation to the lattice constant. The difference between nominal and written values is also not that great in the example given. This is because these values are for a silicon structure and we have a very good silicon etch. Shapes we etch out of the material,

| Parameter                        | Description                                       | Nominal Value (nm) | Written Value (nm) |
|----------------------------------|---------------------------------------------------|--------------------|--------------------|
| Lattice constant ( $a$ )         | spacing between holes in the mirror region        | 448                | 448                |
| Thickness ( $t$ )                | thickness of the nanobeam                         | 220                | 220                |
| Beam width ( $w$ )               | width of the nanobeam                             | 600                | 623.8              |
| Hole height ( $h_y$ )            | height of the hole perpendicular to the beam axis | 340.8              | 340.1              |
| Hole width ( $h_x$ )             | width of the hole along the beam axis             | 240.8              | 235.4              |
| Cross lattice constant ( $c_a$ ) | spacing between crosses in the phononic shield    | 650                | 650                |
| Cross height ( $c_h$ )           | height of a cross                                 | 552.5              | 519.4              |
| Cross thickness ( $c_t$ )        | width/thickness of an arm of a cross              | 162.5              | 146.1              |

Table 3.2: **Design compensation for silicon nanobeams.** Table indicating some of the parameters required to define a nanobeam and the surrounding phononic shield. The nominal value of the parameters to produce a device with wavelength  $\lambda = 1550$  nm, along with values compensating for blowout during the lithography and etch steps. Buffer values used are from 5G29SS. Design parameters are identified in Figure 3.4.

holes and crosses, are written smaller as they tend to blow out, and the beam is written larger, as the trench we write next to it tends to blow into the beam. For comparison, Table 3.3 contains the values for a silicon nitride ( $\text{Si}_3\text{N}_4$ ) nanobeam aiming to have its optical mode at the same frequency. All the values are larger due to the smaller refractive index of silicon nitride.




Figure 3.4: **Unit cell of the nanobeam and phononic shield.** (a) Unit cell of the mirror region of a 5G nanobeam. The design parameters are shown. (b) Unit cell of the phononic shield which is referred to as a cross. The parameters are again shown. This figure adapted from [22].

The values for silicon nitride are also not too bad, but worse. The etch for silicon nitride was recently optimized when we began work on  $\text{Si}_3\text{N}_4$  nanobeams; before that these values would have been much worse (larger blowout).

| Parameter                        | Description                                       | Nominal Value (nm) | Written Value (nm) |
|----------------------------------|---------------------------------------------------|--------------------|--------------------|
| Lattice constant ( $a$ )         | spacing between holes in the mirror region        | 615                | 615                |
| Thickness ( $t$ )                | thickness of the nanobeam                         | 400                | 400                |
| Beam width ( $w$ )               | width of the nanobeam                             | 1073               | 1115.3             |
| Hole height ( $h_y$ )            | height of the hole perpendicular to the beam axis | 565                | 550.1              |
| Hole width ( $h_x$ )             | width of the hole along the beam axis             | 533.2              | 498.7              |
| Cross lattice constant ( $c_a$ ) | spacing between crosses in the phononic shield    | 1050               | 1050               |
| Cross height ( $c_h$ )           | height of a cross                                 | 892.5              | 869.4              |
| Cross thickness ( $c_t$ )        | width/thickness of an of a cross                  | 231                | 175.2              |

Table 3.3: **Design compensation for silicon nitride nanobeams.** Table indicating some of the parameters required to define a nanobeam and the surrounding phononic shield. The nominal value of the parameters to produce a device with wavelength  $\lambda = 1550$  nm along with values compensating for blowout during the lithography and etch steps. Buffer values used are from NNB05. Design parameters are identified in Figure 3.4.

### 3.2.3 File Fracturing

The purpose of fracturing is to convert a drawn design into shapes the EBPG can understand and write into the resist. The software we use for this is GenISys' Layout BEAMER [47] and it takes Autocad .dxf files as input and outputs Leica/Vistec .gpf files that our EBPG uses. The software has many options and modules that have evolved over time as the company can be very responsive to customer needs. This can be good and bad as it can lead to a lot of very specific but often poorly documented functionality. A sequence of modules, called a flow, defines the manipulations done on the geometry from import to export. The flow for nanobeams is shown in Figure 3.5 and is fairly straightforward. The options in a standard flow are described in Table 3.4 and below.

**Import:** The flow begins with an IMPORT step where the .dxf file is selected. Tolerances are specified and we generally set them to have a very small physical dimension ( $\sim 1$  pm), much smaller than what the machine can handle. We do this to ensure that it is trying to make everything as finely grained as possible, i.e., we want the machine's physical limitations to be the limiting factor when it comes to writing fine structures, not a software setting. It is very likely setting these to more realistic parameters on the order of  $\sim 1$  nm would be sufficient but as there is a single license for the entire university, experimentation is difficult. Setting the tolerances to extremely small values may not be useful but it does not seem to be harmful.



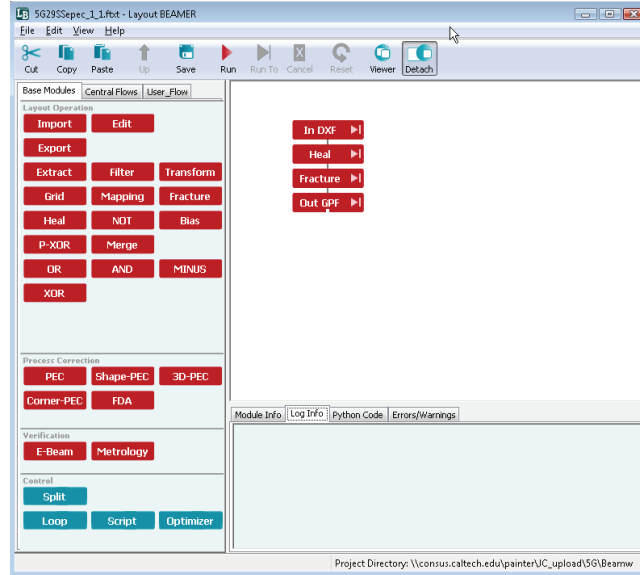


Figure 3.5: **Nanobeam fracture flow.** Screen capture of Layout BEAMER showing the nanobeam fracturing flow: Import (In DXF), Heal (Heal), Prefracture (Fracture), and Export (Out GPF). The prefracture module is a new step we have only recently started trying. It is included here to provide a placeholder for its use.

**Heal:** The next step in the nanobeam flow is a HEAL step. This step is required because of the way we write the cross to be etched in the phononic shield. For simplicity it is written as two overlapping rectangles. To avoid double exposure in the middle a HEAL step is then required in Layout BEAMER. No special parameters need to be set in this module as we are doing a very basic overlap removal so the default values are kept. This step could be easily avoided by writing the cross as a 12 point polyline in Autocad, and in the future is probably a good idea for the sake of simplicity.

**Export:** The final step in the nanobeam flow is the EXPORT step. The first thing required in this step is selecting the appropriate EBPG tool definition; in our case this is “EBPG Caltech 100keV HS” which has all the tool properties preset. The resolution parameter sets the grid size of the pattern to be written and we choose 2.5 nm, the best our tool can do, which defaults the beam step size to 2.5 nm. We want the smallest resolution possible so that curves are as smooth as possible, edges snap to the closest possible position to maintain correct dimensions and periodicity is maintained. Choosing a resolution determines the maximum size of the main field, which is the area the EBPG can write to without having to physically move the stage (when the stage needs to move and continue writing, this is called stitching and in theory is quite good). Our EBPG has a main field DAC size of 16 bits and at 2.5 nm this sets the main field size to 160x160  $\mu\text{m}$ . Each nanobeam is much smaller than this, so there is no reason to draw a pattern larger than the main field size and have to deal with the stage physically moving while trying to draw the pattern. As such, we keep

| Module | Parameter        | Description                                                                                                                                                                              | Selected Option        |
|--------|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| Import | .dxf file        | import the Autocad formatted geometry                                                                                                                                                    | -                      |
|        | Database Grid    | DXF internal - use grid resolution from dxf file                                                                                                                                         | 0.000001 $\mu m$       |
|        | Curve Conversion | converting a curved line to a grid                                                                                                                                                       | 0.000001 $\mu m$       |
|        | Snapping Range   | snapping geometry points to a grid                                                                                                                                                       | 0.0000005 $\mu m$      |
| Heal   |                  | removes overlap in geometry (crosses are two overlapping rectangles - otherwise would double expose)                                                                                     | all default            |
| Export | .gpf file        | export and fracture the geometry for the specific EBPg                                                                                                                                   | EBPg Caltech 100keV HS |
|        | Resolution       | sets the grid size for the final fractured geometry<br>smaller results in smoother curves and snapping of points to desired location                                                     | 0.0025 $\mu m$         |
|        | BSS Fracturing   | Beam Step Size Fracturing: beam step size is by default set to fracturing resolution<br>this is meant to adjust the beam stepping if the number of shots won't fit evenly in the polygon | yes                    |
|        | Main Field Size  | set by the choice of resolution (2.5 nm) and the main field DAC bits of the machine (16)<br>we almost always keep our pattern size within one main field                                 | 160x160 $\mu m$        |
|        | Fracture Mode    | determines how the software breaks up the geometry into polygons the EBPg will then write                                                                                                | conventional           |

Table 3.4: **Layout Beamer fracturing options.** The standard options we use for file fracturing with a description and reason.

our .dxf/.gpf patterns within one main field. There is an option for beam step size (BSS) fracturing; if the beam step size was larger than the resolution, the number of beam steps might not fit within a polygon being draw. Selecting BSS fracturing is meant to deal with this as best as possible. Our beam step size is the same as our resolution so this should not be a problem, but we selected it anyways just in case. The fracture mode parameter determines how the software will break up the pattern into polygons that the EBPg will actually write. There are currently three options in the version we are using: large rectangle fine trapezoid (LRFT), conventional, and curved. We generally only use conventional (and have tried curved) and used to draw our circles and ellipses in Autocad as a 2000 point polyline to force Layout BEAMER to fracture it into very fine lines of trapezoids. When the tool was upgraded from VMS operating system to Linux, with a corresponding change in some of the DAC cards, this started producing terrible results. The machine had been producing sub-par

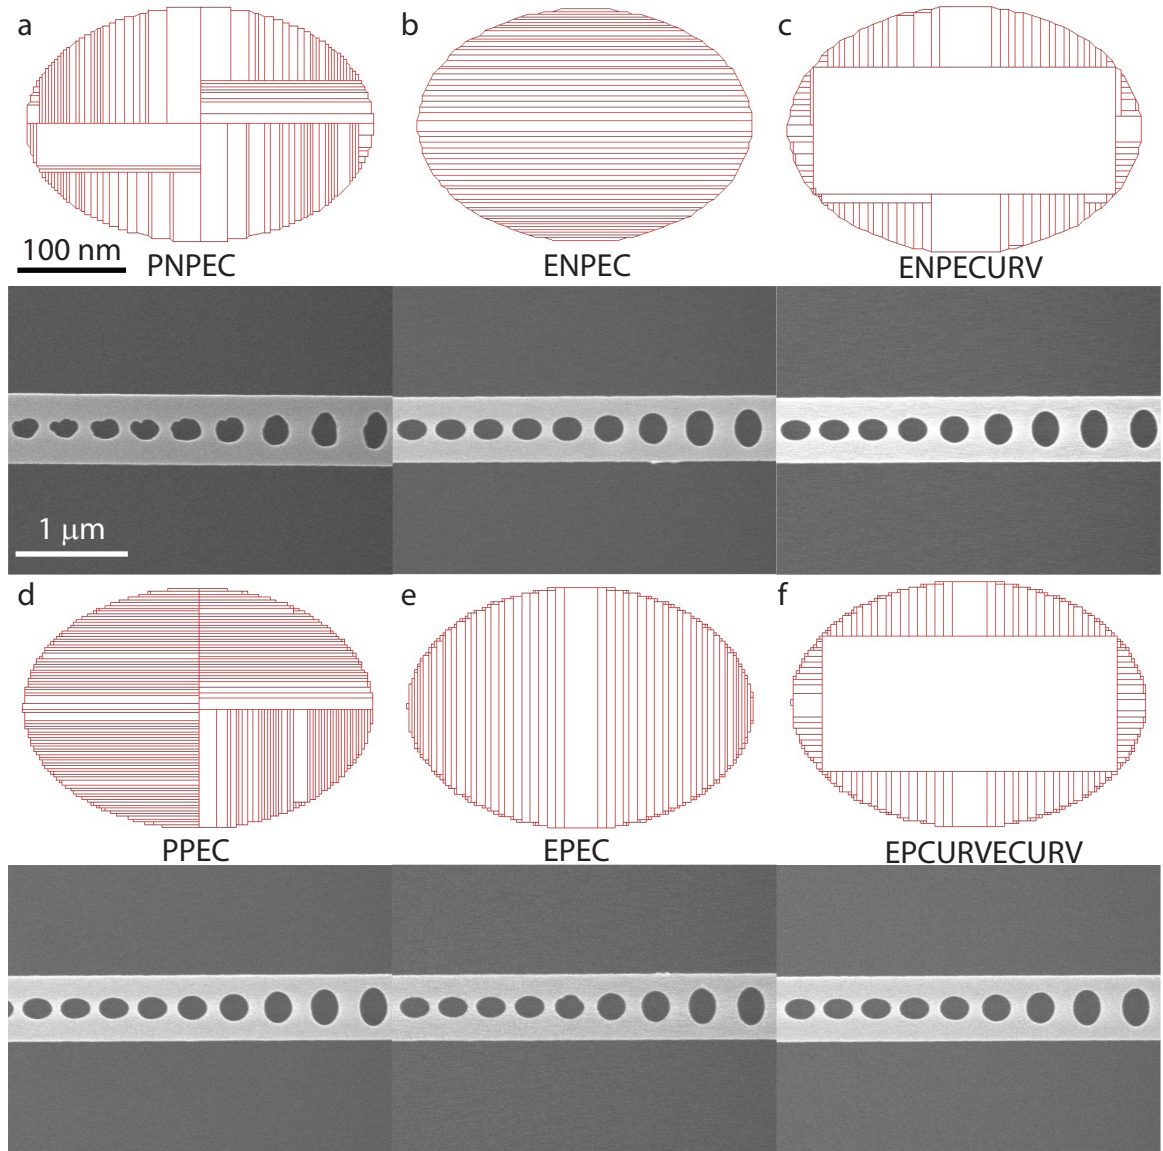


Figure 3.6: **Fracture test patterns and results.** The top is the result of fracturing a mirror period hole. The bottom is a SEM of an etched and undercut device using that fracturing technique. The scales are the same for all fractured images and SEMs, though they are different from each other. Three variables are explored: polyline vs. ellipse, prefabricating or not, curved vs. conventional fracturing. (a) Polyline, no prefabricating, conventional export fracturing. (b) Ellipse, no prefabricating, conventional export fracturing. (c) Ellipse, no prefabricating, curved export fracturing. (d) Polyline, conventional prefabricating, conventional export fracturing. (e) Ellipse, conventional prefabricating, conventional export fracturing. (f) Ellipse, curved prefabricating, curved export fracturing.

results for a while and it wasn't immediately clear that the irregular shapes we were seeing were due to the fracturing. On a hunch we varied the method of writing ellipses (the standard Autocad polyline and Autocad ellipse) and varied the fracturing (conventional and curved along with a new module called PREFRACTURE). The results of the fracturing and SEM of the etched nanobeams

is shown in Figure 3.6. While the fractured holes in Figure 3.6b-f don't look terribly different, the results for the polyline + conventional fracturing in Figure 3.6a is horrible. The results of the others all look very similar. The fracturing of the EPEC and EPCURVECURV have finer trapezoids due to the prefracturing, however the results are indistinguishable from the SEMs. For the time being these are the two fracturing methods we have chosen to use and have seen encouraging results. To properly determine if one method is consistently better than another, a systematic study would need to be taken with the optical and mechanical quality factor determining the best method as they will be more sensitive to small imperfections not visible on the SEM images.

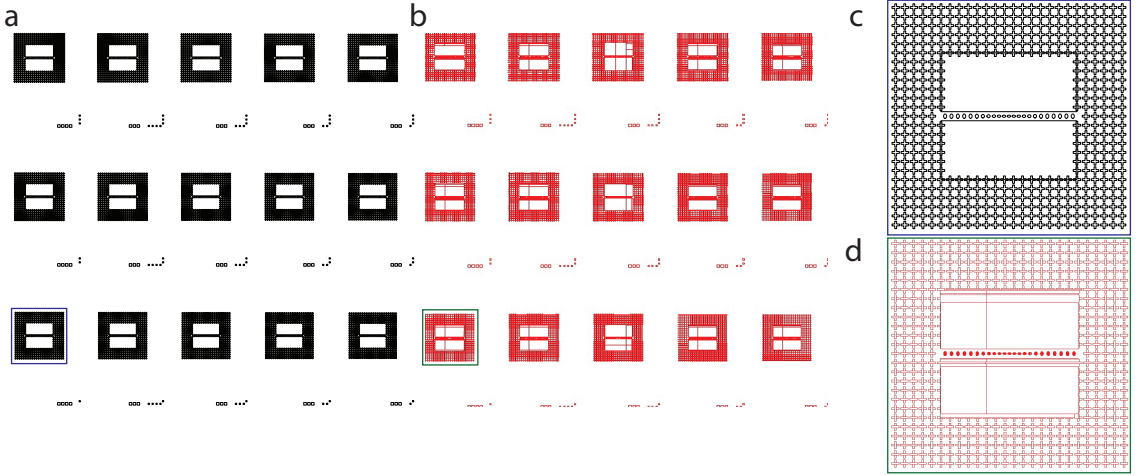


Figure 3.7: **Example .dxf and .gpf files.** (a) An example of one of the Autocad .dxf files generated by the MATLAB script. The parameter sweep can be seen with five different lattice constant nanobeams along each row and each row has a different lattice constant phononic shield. The devices are numbered and this represents one field of view (main field). Here the numbers go from (16,1) in the bottom right to (20,3) in the top left as both lattice constants increase in their respective directions. (b) The fractured version (.gpf) of the same field of view. The quality of the image is much worse as there does not seem to be a good way of exporting a nice image of the fractured file which is unfortunate. (c) Zoomed in view of the (20,1) device from the .dxf file as indicated by the blue box. (d) Zoomed in view of the (20,1) device from the .gpf file as indicated by the green box.

As mentioned, we do not like to have a single pattern span more than one field of view (FoV—same as main field) but at the same time like to write large parameter space sweeps for our nanobeams. The solution we use is to write as many different nanobeams in a FoV as possible and have several different .dxf/.gpf files. The standard nanobeam parameter sweep of 20 optical lattice constants and three cross lattice constants giving 60 devices is split into four FoVs, each with a 5x3 array of nanobeams. Figure 3.7 shows the .dxf file and fractured .gpf file for one of those four patterns. Other solutions to this problem exist and it can be done with one .dxf/.gpf file and either ensuring the nanobeams themselves will not fall along a main field split or providing coordinates/grids for Layout BEAMER to use to split the pattern into main fields.

### 3.2.4 Spinning and Baking of Electron Beam Resist

To define our etch mask we need an etch resist layer that can be patterned by the EBPG. We use ZEP-520A [48], a positive, electron beam resist from Zeon Chemicals. This resist is very good for electron beam lithography of photonic crystals due to the ability to write very small feature sizes, and it stands up well to a silicon etch. The price seems ridiculously high (well over 1000 USD/100 mL) but given the quantity used and other costs associated with processing a chip it is well worth it (not that it should be wasted). The steps we use to spin and bake a chip for nanobeam fabrication are given below.

**Clean:** The first step is to clean a chip. We buy our SOI wafers undiced and then send them out to be diced to a specific size (often 13 mm x 5 mm—this elongated size allows them to be easily clamped in the EBPG for the beamwrite). The diced wafers generally will have photoresist on them to protect them; regardless, it is important to clean the surface of the chip well before spinning and baking. The standard solvent clean is outlined in Section 3.4.1 and generally is sufficient to clean the chips.

**Spin:** We spin the resist on our chip using a vacuum chuck spinner. First, the vacuum is turned on and the chip placed on a chuck of appropriate size for the chip. The chip should be positioned in the center of the chuck so that when it spins it is symmetric about the center. This way it doesn't experience any torque while spinning and the resist will disperse evenly from the center. The recipe we use is ramp: 2500 rpm/s; speed: 5000 rpm/s; time: 60 seconds; and should already be set as the current recipe before beginning. For a 13x5 mm chip like we use for our current silicon processing, this results in a resist thickness of  $\sim 350$  nm. In general, the thickness depends mostly on the spin speed and somewhat on the chip dimensions. It is a good idea to always start the recipe before putting any resist on to ensure that it spins up properly (doesn't have a vacuum error) and is, in fact, the desired recipe (spin it back down once this is verified). Using a fresh pipette (we use glass) that has been blown clean with  $N_2$  and inspected for dirt, attach the pump and insert the pipette into the ZEP520A bottle not touching any of the sides. Ensure the bottom of the pipette is below the top surface (where some particulate may be—hopefully not) but not touching the bottom of the bottle. Draw in the required amount (but not too much more) and remove the pipette while still not touching the bottle. For a 13x5 mm chip, three drops should be sufficient. They should be placed in the center of the chip and on top of each other as this will allow them to spread out evenly and cover the chip while not having different drops merging into each other (which after drying you can see). Immediately start the spinning (usually done with a foot pedal so this is easy). Be over the chip as little as possible to avoid dropping particles into the resist. When the spinning is done, remove the chip to the glass slide for baking and cover it with a upside down aluminum tin (bowl). It is best to have the ZEP bottle uncovered for as little time as possible to reduce the chance of contamination.

**Bake:** When all the chips have finished spinning, the resist needs to bake to remove the solvents.

We use an oven for this and it should set to 180°C (do this at the beginning so it is ready). Place the glass slide with the chips (not the chip cover) in the oven and let bake for 20 minutes. An alternative, which should be better and is certainly faster, is to bake the chips on the glass slide, on a hot plate. Historically this is not what we have done so we have not switched to this method, but it should be better and is worth trying. After the baking is done, remove the chips and place on a surface to cool down; while they are cooling, cover them to avoid particulate falling on them. Let them cool for a reasonable (up to 5 minutes) amount of time and then they are good to go.

### 3.2.5 Electron Beam Writing

The EBPG uses the .gpf file to dose (that is, apply a specific amount of electrons to a unit area and chemically alter the resist) the resist covered chip. The required dose depends on the type of resist, resist thickness, chip material layers, feature size and feature layout. For the nanobeam processing, the optimal dose is  $\sim 235 \mu\text{C}/\text{cm}^2$ . We always write a dose run, which means we write the same set of patterns with many different doses on the same chip. Typically we will write ten different doses with a dose step of  $5 \mu\text{C}/\text{cm}^2$  around this nominal value (these numbers can of course change as required). New processes need a large dose run at first to determine the nominal dose. If a dose isn't high enough the entire feature will not be cleared; that is, a hole in the mask will still have some resist remaining and it will not etch properly. A good example of an underdosed pattern can be seen in Figure 3.8a. Here the underdosing lead to an incomplete removal of the resist in the trench area. The remaining resist is never a nice flat surface and often very wavy showing indications of exactly how the EBPG wrote the area (i.e., you can see shape boundaries, etc.). Upon etching, the resist was finally eaten through and began etching the device layers below, transferring the wavy pattern of the resist resulting in the wavy floor seen in Figure 3.8a. For a comparison of the same etch with a higher dose, see Figure 3.15—the contrast is quite striking and a good example of the effects of dosing. It is not worthwhile to try and optimize an etch until the dose is correct, as the etch times and effects on the sample will clearly be different. If a dose is too high it will cause the feature size (hole) to blow out and be larger than it should be. A good example of an overdosed sample can be seen in Figure 3.8b where all the features are distorted: the elliptical holes are blown out, the nanobeam is severely narrowed and blown in, and the snowflakes are very distorted and almost not recognizable. The snowflakes seen in Figure 3.8b should look similar to the ones in Figure 3.11b; again the contrast is striking and goes a long way towards showing the effects of dosing. These effects can be felt even in the small dose run around the nominal dose as mentioned above. While in the SEM images, the sizes may look very similar, they will in fact be slightly different and the effects of the dose run will be seen in the optical testing of the devices as each dose shifts the frequency of the optical modes very slightly which is not necessarily a bad thing. We have had situations where we have seen a nice device but it was just outside the range of a key piece of equipment (the EDFA

say) and have switched doses and the mode had shifted into the useful range.




Figure 3.8: **Example of under and overdosing the resist.** (a) An example of a very underdosed sample. The wavy floor seen in the trench area is an indication of an incomplete dosing leaving behind an uneven floor of resist which when finally etched through began transferring that pattern into the device. Compare to Figure 3.15. (b) An example of a very overdosed sample. The elliptical holes are blown out, the beam is severely narrowed (blown in) and the snowflakes look like chubby monsters. For comparison see Figure 3.11b.

The choice of a beam size (measured in pA or nA) is partially determined by the resolution of the .gpf file being written and then further by the needs of the user. Smaller resolutions require smaller beams or the electron beam would have to be deflected faster than it was capable of to avoid overdosing. The tool we use can nominally write at a maximum of 25 MHz, but will allow writing up to 50 MHz (probably not advised). Larger beams clearly dose the sample faster and are good for large open areas. Smaller beams dose it more slowly (and thus take much longer) and are good for fine stepping ( $\sim 2.5$  nm resolution) as mentioned above. Smaller beams also have smaller spot sizes, that is, the size (diameter) of the beam of electrons exposing the resist. Clearly for fine features a smaller spot size is required. We generally use a 300 pA beam which has a spot size on the order of  $\sim 4$  nm—about as good as the tool can do.

### 3.2.6 Develop and Rinse

The electrons have exposed the resist and chemically changed it, but the chip must be developed to remove those exposed areas before it can be used as a mask for etching.

**Develop:** The developer is Zeon’s proprietary developer, ZED-N50. We develop for 2:30 (2 minutes, 30 seconds) while holding the chip with tweezers on the sides or a small portion in the corner, gently agitating it in the developer. We use fresh developer each day and change after 4-6 chips.

**Rinse:** The rinse we are currently using is Zeon’s ZMD-D, though we are in the process of



changing this to MIBK (methyl isobutyl ketone), which is the same thing (Zeon is no longer selling ZMD-D). Immediately after developing, the chip is rinsed for 20 seconds and then removed and blown dry with  $N_2$ .

**Note:** The resist will absorb moisture over time changing its properties and bloating—increasing in size. This can ruin fine features and cause unexpected results. As a general rule we develop the chip as soon as possible and etch within several days, never more than a week. This may be overly cautious, but better safe than sorry.

## 3.3 Etching

### 3.3.1 ICP RIE Overview

The goal of this section is not to provide a thorough understanding of ICP RIE etching. That is a much larger and complicated task that is best left to theses, articles, and books with that specific aim (see [49, 50] as good starting points). However, a basic high-level understanding is required. A schematic of a generic ICP RIE system is shown in Figure 3.9.

The goal of etching is to remove material in a controlled way, leaving behind the desired structure. To this end, inductively coupled plasma (ICP) reactive ion etching (RIE) is used. In plasma etching in general, an etch gas is ionized and the resulting ions can either chemically attack the sample or be accelerated towards it to physically mill it or both. For silicon etching the chemical used is sulfur hexafluoride ( $SF_6$ ). When ionized, the fluorine ions (F) chemically attack the silicon and the heavy  $SF_5$ ,  $SF_4$ , etc. ions can be accelerated towards the silicon to physically bombard it, knocking away material. Most new etch tools used for nanofabrication such as the Oxford Instruments PlasmaLab System 100 ICPRIE 180 (the tool we have), decouple the generation of the plasma and the acceleration of the ions towards the sample, allowing the strength of the plasma (ion density) and the momentum/energy these ions have when bombarding the sample to be independently tuned. The ion density of the plasma is set by the ICP power while the strength of the sample bombardment is set by the RF forward power. In this configuration, and using an etch mask as described above, a pattern may be transferred into the silicon device layer below the resist. A problem with the above is while the milling aspect of the etch can be made much larger than the chemical aspect, there will always be some lateral (i.e., not in the direction the ions are accelerated) etching of the material, which is generally unwanted. To compensate for this, an inhibitor can be added to the gas chemistry. For a silicon etch this inhibitor is octafluorocyclobutane ( $C_4F_8$ ) which, when ionized in the plasma, will coat all the surfaces with a protective polymer layer. Then the forward bombardment of the sample will remove the polymer only on the horizontal surface (perpendicular to the direction of the accelerated ions) and produce a very directional etch as desired. This etch is referred to as a **pseudo Bosch** etch [49, 51, 52].



As described above, the ICP power is responsible for setting the number of ions in the plasma, while the RF forward power determines the strength with which those ions bombard the sample surface. The gas chemistry ( $\text{SF}_6/\text{C}_4\text{F}_8$  in this case) determines what materials can be etched and if there is any sidewall passivation. The plasma itself has a characteristic color called the **plasma glow** and can be seen through the view port once the plasma has been struck. This color is dependent on the specific gas chemistry and can be a quick, obvious check if something is wrong and one of the chemicals is not contributing to the plasma. The glow occurs when electrons, ripped from the precursor components, are forced to collide with the heavier gas ions by the oscillating ICP field. This excites the ion and when the electron in the ion de-excites to its lower energy level it emits a photon. Around the edge of the chamber, next to the sample surface, and anywhere else there is something in the chamber is the **dark space**. This region is devoid of electrons, and as such does not glow. Ions that recombine are not impacted by electrons to re-ionize them leading to a space with some density of neutral particles which can deflect other ions on their way to bombard the sample surface. The density of particles in the dark space is set by the chamber pressure and is the largest factor in determining the ion angular distribution function (IADF) which affects the etching in the non-preferred direction. These collisions also change the energy of the ions leading to the ion energy distribution function (IEDF).




Figure 3.9: **Schematic of a generic ICP/RIE etch system.** A schematic of a generic ICP/RIE etch system similar to our own. Many of the key points discussed can be seen. Adapted from [49].

### 3.3.1.1 $\text{SF}_6$ - Sulfur Hexafluoride

Sulfur Hexafluoride ( $\text{SF}_6$ ) is the gas that causes the etching of silicon. When ionized in the plasma, it dissociates into  $\text{SF}_x$  and  $\text{F}_y$  ions which can then be used to bombard the sample surface, knocking out silicon atoms (milling) or chemically attacking the silicon surface. Creating a plasma

with  $\text{SF}_6$  without accelerating towards the sample (i.e., not giving it any preferred directionality) will cause an isotropic etch of silicon due to the extreme reactivity of the fluorine ions (F). At the silicon surface the F ions will react with Si producing the volatile gas  $\text{SiF}_4$  which is pumped away leading to an isotropic etch of silicon. Accelerating this plasma towards the sample surface allows the heavier  $\text{SF}_x$  ions to participate in milling of the silicon surfacing, leading to some degree of directionality of the etch. Clearly a more anisotropic etch is desired which leads to the inclusion of a passivation gas,  $\text{C}_4\text{F}_8$ .

### 3.3.1.2 $\text{C}_4\text{F}_8$ - Octafluorocyclobutane

The inclusion of  $\text{C}_4\text{F}_8$  in the etch chemistry provides means to create a very directional, anisotropic etch. When ionized in the plasma, it dissociates predominantly into  $\text{CF}_2$ ,  $\text{C}_2$ , and  $\text{C}_3$ . The surface passivation comes from the  $\text{CF}_2$  molecules recombining on the surface, producing cross-linked polymer chains protecting the silicon. This polymer will coat all the surface (including etch chamber walls) and prevent etching. The anisotropic aspect of the etch is then regained by the milling of the polymer on the horizontal surfaces only, due to the acceleration of the plasma towards the sample surface. Ideally, the ions will only mill the horizontal surface, removing the protecting polymer and allowing chemical and milling to occur in only this one direction. The introduction of this gas into the mix will also slow the etch in general as the dissociated ions scavenge the F ions to produce harmless  $\text{CF}_4$  molecules.

### 3.3.1.3 ICP Power

The ICP power is used to generate the plasma from the input gases and can be thought of as controlling the number of ions available for etching (either chemical or milling) and passivation. The flow rate of the input gases along with the set pressure will determine the amount of gas in the etch chamber. The ICP power will then determine the amount of gas that will be ionized into the plasma. Too little power and the plasma will continually be bumping into the inert, unionized precursor gas and unable to etch the sample. Increase the power beyond a certain point and all the gas will be ionized and the etch rate will not increase. This is not entirely true for the etch we do where there is more passivation gas ( $\text{C}_4\text{F}_8$ ) than etch gas  $\text{SF}_6$ , and as such increasing the ICP power should produce more passivation ions than etch ions.

### 3.3.1.4 RF Power and DC Bias

The RF power is what we use to accelerate the plasma towards the sample surface for etching. It is primarily responsible for the milling aspect of the etch process. Higher RF powers will tend to mill the surface more, and also improve directionality of the etch as the scattering in the dark space becomes less pronounced. This also means that higher RF power will also consume or damage

the resist faster. The **DC Bias** reading is taken between the etch table and plasma and can be considered a measure of the energy with which the ions are bombarded into the sample surface. It is primarily determined by the RF forward power but also has a smaller dependence on the chamber pressure, ICP power, and to some degree wafer position (which can change the accuracy of the reading of the voltage). On our etch table, part of the coating of the table has been scratched away where the flat of the carrier wafer is, thus exposing the uncoated table to the plasma and improving the ability to read the **DC Bias**. A sudden change in the **DC Bias** of an etch for the same nominal parameters may have more to do with the carrier wafer, or its placement, than with the conditions of the etch. Regardless, it is something that is important to note and attempt to understand as the results may be unpredictable in this situation.

For our etches, we aim for a specific **DC Bias** value and adjust the RF forward power to get that value; that is, we trust the **DC Bias** to be the more representative value of the etch. It is, of course, not a blind trust and the correspondence between **DC Bias** and RF forward power must make sense based on experience and the past history of the etch and tool. Along with the ICP power, the RF power is one of the easier to understand control aspects of the etch.

#### 3.3.1.5 Chamber Pressure

For our etches the chamber pressure ( $P_{\text{chamber}}$ ) is set to 15 mTorr. This pressure is set by the automatic pressure controller (APC), which adjusts the position (angle) of a valve connecting the magnetic levitation (MAGLEV) turbo pump to the etch chamber. The pressure is important for several reasons as it directly affects the plasma. Setting the pressure too low will cause the plasma to be pumped away before it can etch the sample. Setting the pressure too high will cause the molecular density to be too high, increasing the collisions and slowing down the etch, and decreasing the directivity of the ions (decreasing the anisotropy of the etch). Based on the gas chemistry and pressure, the mean free path governing the amount of collisions can be calculated. Most importantly, reasonable values for the pressure are known and have been tested over time (see [50] and reference within). This is not a value we change for the etches we are doing.

#### 3.3.1.6 Helium Backing

The purpose of the helium backing is to provide better thermal contact to the carrier wafer and sample being etched, allowing for stable, repeatable, etch conditions. A substantial amount of power is used to create the plasma required to perform the etch, most of which ends up as heat [50]. Heating of the substrate occurs from both the exothermic chemical reaction and ionic bombardment. Heating of the substrate will provide more energy to overcome the activation energy of the chemical reaction, increasing the chemical aspect of the etch and thus making it more isotropic. More importantly, fluctuations in the temperature will change the etch conditions, making the etch unpredictable and

unrepeatable which is something that needs to be avoided.

For all our etches, we set the **Helium Backing Pressure** ( $P_{\text{He}}$ ) to 10 Torr and record the **Helium Flow** (Flow) required to maintain this pressure. To achieve this backing pressure, helium is flowed up from below the table the wafer sits on, around the wafer lift and presses up against the carrier wafer. Holding the wafer down against the table and flowing helium is the quartz clamp. The flow required to maintain the backing pressure is an indication of how well the wafer is clamped to the table. A low flow ( $\sim 5\text{--}6$  sccm) indicates good clamping and good thermal contact, and is what is desired for repeatable etch conditions. A high flow ( $\sim 7\text{--}10$  sccm) indicates poor clamping, bad wafer position, thinning carrier wafer, etc., and is very undesirable for repeatable etch conditions. Furthermore, a high flow may indicate a problem with the delivery and positioning of the wafer by the robotic arm and wafer lift on the etch table (including poor initial placement of the wafer by the user on the robotic arm in the load lock). If the flow is too high, there is a concern that the wafer is so far out of position that the attempt to lift it back onto the robotic arm during removal will fail and the wafer will become trapped in the etch chamber (requiring it to be vented and opened for removal—to be avoided).

### 3.3.1.7 Etch Table Temperature

The etch table, where the carrier wafer sits while being etched, is cooled by flowing water which is chilled by a chiller in the chase. We set the temperature of the chiller ( $T_{\text{chiller}}$ ) to  $15^{\circ}\text{C}$  for the silicon etch and this gives an initial table temperature of  $19^{\circ}\text{C}$  (the table temperature generally saturates about  $4^{\circ}\text{C}$  warmer than the chiller temperature). As mentioned in Section 3.3.1.6, the temperature changes the etch, and a lower temperature should help with the anisotropy of the etch. During a four minute silicon etch the table temperature will rise to  $20^{\circ}\text{C}$ . The table temperature is not a parameter that we change often. It is possible a lower value would produce nicer etches (e.g., straighter side walls) but would require reworking the etch.

### 3.3.1.8 Matching Capacitors (AMU1 and AMU2)

The etch works by applying RF power to the plasma to both continually generate ions and electrons for the plasma and accelerate it towards the sample surface. The characteristic impedance of the plasma will in general be different from the RF sources being used, and also depends on the specifics of the plasma (e.g., gas chemistry). Due to the large electrical powers (up to  $\sim 1000$  W), an impedance mismatch between the source and load (plasma) could produce substantial amounts of reflected power and damage the machine. It also goes without saying that if the matching is different for different etches, the etch itself will not be the same. The Oxford has an auto match unit (AMU) for both the ICP and RF subsystems which controls two capacitors ( $C_1$  and  $C_2$ ) to attempt to automatically match the impedances, delivering the optimal power to the plasma with the least

reflected power. The ICP is controlled by AMU<sub>1</sub> and the RF by AMU<sub>2</sub>.

Since we always do a conditioning run before the etch, in this run we manually optimize the RF matching (AMU<sub>1</sub>; though in manual not automatic mode) and record those values, as they should be reasonably consistent for the same etch over short periods of time (and perhaps can be an indication of changing chamber conditions). The values for the actual etch should then be set correctly. The goal is to minimize the reflected power, and as a secondary goal to maximize the DC Bias, by adjusting C<sub>2</sub> which is a rough, large adjustment, and then C<sub>1</sub> which is a fine, small adjustment.

The adjustment of the ICP matching (AMU<sub>2</sub>) can be left to be done automatically for set ICP forward powers of  $\leq 600$  W. For higher powers the AMU gets confused, rails the capacitors, and the reflected power is too high. In this case, setting AMU<sub>2</sub> to automatic before starting the etch will set C<sub>1</sub> = C<sub>2</sub> = 5, which is in the middle of the range. Then set it to manual and start the etch. When the plasma is struck, pushing up on C<sub>2</sub> will increase its value and bring the reflected power down quickly to  $\leq 100$  W, at which point AMU<sub>2</sub> can be switched to automatic and the auto-matching will be successful. We leave AMU<sub>2</sub> in automatic mode while etching and AMU<sub>1</sub> in manual mode as mentioned.

### 3.3.1.9 Etch Time

The etch time is set by the conditions of the etch and depth of the device layer to be etched. Longer etches are require for tougher material (Si<sub>3</sub>N<sub>4</sub> over Si) and thicker material (again Si<sub>3</sub>N<sub>4</sub> over Si). Etching for longer may not always be the solution as longer etches tend to heat up the sample and can damage the resist. If an etch isn't getting through, one solution might be to increase the time (the easiest) but if the time becomes excess, re-optimizing the etch may be required.

All of these parameters are recorded for each etch in both the log box for the tool and the user's lab book. An example of a etch log is shown in Figure 3.10 with all the parameters discussed above.

| ICP/RIE Etching Log                             |                                                                                                                       |                                                                  |                                     |
|-------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|-------------------------------------|
| Sample Name: _____                              | Material: [ ] Si [ ] SiN <sub>x</sub> [ ] SiO <sub>x</sub> [ ] SiO <sub>x</sub> +Si+SiO <sub>x</sub> [ ] Other: _____ | Date: _____                                                      |                                     |
| Etching Time = _____                            | AMU <sub>2</sub> : succeed @ _____ Attempt                                                                            | AMU <sub>1</sub> : C <sub>1</sub> = _____ C <sub>2</sub> = _____ | T <sub>chiller</sub> = _____ [ ] °C |
| C <sub>4</sub> F <sub>8</sub> = _____ [ ] sccm, | SF <sub>6</sub> = _____ [ ] sccm,                                                                                     | P <sub>He</sub> = _____ [ ] Torr,                                | Flow = [ ] sccm                     |
| RF Power (F/R) = _____ [ ] W                    | ICP Power (F/R) = _____ [ ] W                                                                                         |                                                                  |                                     |
| DC Bias (V) = _____                             | P <sub>chamber</sub> = _____ [ ] mTorr                                                                                | Thickness = _____ nm                                             |                                     |

Figure 3.10: **Etch log template.** The etch log template filled out every time an etch is done recording relevant parameters about the etch and system conditions. Logging the parameters for each etch provides a way of characterizing the drift of the system, diagnosing problems, and keeping a full record of what has been done.

### 3.3.2 Securing the Chip to the Carrier Wafer for Etching

In order to etch the chip it must first be placed on a 100 mm ( $\sim 4''$ ) diameter carrier wafer. The wafer serves two purposes, perhaps the most important being that our system accepts 100 mm carrier wafers and not chips, so this is a way of loading the samples. The second important purpose is that it mitigates the effects of silicon loading when etching. Loading is the term used to describe the etch conditions changing depending on the area being etched. With the silicon wafer in the etch chamber, the amount of surface area on the chip being etched has a negligible impact on the loading and allows for repeatable etch conditions, no matter the size of the pattern. We secure the chip (in some cases chips) to the carrier wafer using thermal paste (Thermalcote: Thermal Joint Compound). This secures the chip to the wafer and also improves the thermal contact between the chip and the He backed carrier wafer. A few small dabs are all that is needed and it is important not to apply so much that it oozes out the sides. It can be helpful for adhesion to place the chip on the thermal paste and smear it around a small amount to spread out the paste and make the chip stick better. It is important to ensure good contact of the chip with the wafer and using the  $N_2$  gun to spray first down on the chip and then around. Finally, picking up the carrier wafer and giving it a shake test to ensure the chip will not fall off.

### 3.3.3 Silicon Etching and Optimization

#### 3.3.3.1 Original Si Etch




Figure 3.11: **Development of original Si etch.** The original Si etch was developed for a somewhat tricky geometry to etch, the snowflake. (a) A SEM showing how the Si etch performed on the snowflake structures before optimization. (b) The resulting etched geometry after optimization.

The original silicon etch was developed to etch a somewhat tricky geometry; that of the snowflake as seen in Figure 3.11. The geometry is slightly tricky to etch because the long, thin trenches making up the branches of the snowflake etch slower than the larger center. Etching long enough to etch

the trenches nicely leads to blowing out the central region and only long enough to not blowout the central region causes the trenches to not fully etch through and significantly rounds their edges, as can be seen in Figure 3.11a. At the time of beginning to etch the snowflakes, the etch chamber had been oxygen plasma cleaned one too many times, causing the chamber conditions to change significantly and resulting in the silicon etch no longer being optimal.

This was the first time we really had to try to optimize an etch and it was a somewhat long process, but a very valuable learning experience. The results of the optimization can be seen in Figure 3.11b, showing a nearly identical structure under the etch parameters we developed. This structure is much less blown out, through in the trenches with less rounded edges, and is quite faithful to the intended design. This etch was very stable and consistent, producing nice results for the silicon etching the group was doing. It was used for over a year without any substantial changes to the etch parameters.




Figure 3.12: **Original silicon etch.** The original silicon etch showing several non-optimal features. The resist layer, silicon device layer, and buried silicon dioxide layer are labeled and lines guide the eye to the delineation between each layer. The resist layer shows some signs of an angled and rounded top which can be transferred to the device layer and is worse than before. The silicon device layer shows an inward, reentrant tapering that is very undesirable and much worse than before. The silicon dioxide layer is etched quite deeply suggesting over-etching and also contains tapered sidewalls indicating poor etch performance.

Ideally, the etch chamber would remain consistent indefinitely; however this is obviously not the case in reality. Several factors can lead to changes, including having to vent the chamber to air to fix a problem and the build-up of polymers on the chamber walls, to name two. After 14 months of using this etch it was found to have changed enough that reoptimization was required. Figure 3.12 shows a tilted view of the end of a nanobeam where it is anchored to the phononic shield. In this figure the resist, silicon, and silicon dioxide layers are highlighted. The resist itself shows signs of

degradation at the edges, resulting in slanted resist sidewalls. If this slant is too extreme it can result in etching of the top of the silicon producing non-vertical sidewalls of the silicon structure, which is unwanted. The slant in the image does not appear quite that bad but looks to be nearing what would be a problem and is a point of concern. The silicon itself is clearly seen to be not vertical with an inward (reentrant) tapering, which is the biggest concern. The silicon dioxide, while not of any relevance for the final structure, shows a step outward tapering and is etched quite deeply. We need to ensure that we etch fully through the silicon layer, but don't want to etch more than necessary as this only leads to blowout and a concern that the etched silicon dioxide could sputter the sidewalls of the silicon and cause the sidewall surfaces to be unsmooth. These are the issues that we attempted to correct in the silicon etch reoptimization outlined in the next section.

### 3.3.3.2 Si Etch Reoptimization Steps

The goal of this section will be to provide a guide of the thought process used in reoptimizing the silicon etch. At each stage an example SEM showing the current state of the etch will be shown, along with an explanation of the problems seen in the SEM and rationale for the steps to be taken in the next each. As the reoptimization proceeds the steps and way the samples are analyzed changes in accordance with our better understanding of how to optimize an etch. At the end of the section a summary and key points learned will be discussed and this new knowledge is applied to optimizing a silicon nitride etch, which, while similar, has its own unique set of challenges.

It is important when trying to optimize an etch to keep track of as much information about the etch itself and properly and consistently analyze the results. Here all the etch parameters were recorded in an etch log (see Figure 3.10), as well as a power-point presentation to provide etch information coupled with SEMs and thoughts on the results coupled with a reasoned plan on moving forward. It is important to take very good SEMs as the finer details can often be lost with fuzzy images. Consistency is also very important and similar images of each sample and a specified magnification should be used to allow for the easiest comparison. In the images shown here they are not all alike and this is due to our evolving (and we believe improving) understanding of the best method to characterize the quality of the etch.

Silicon is in general, a relatively easy material to etch where the selectivity of etching silicon to the resist is high ( $\sim 4:1$ ) and the thickness we need to etch for our structures is relatively small (220 nm compared to our resist thickness of  $\sim 350$  nm). While this makes the etch much easier to get close to the results you want, it also means that you can be much pickier about the final result. Figure 3.13 shows the evolution of the etch with SEM images taken for each step corresponding to the etch parameters in Table 3.5 in the same order presented. With that said we begin with an analysis of the etch as it stood at the beginning of the etch reoptimization process.






Figure 3.13: **Progression of the Si etch reoptimization.** (a) atnb06x. (b) atnb06y. (c) ET01h. (d) ET01m. (e) ET02e. (f) ET02f. (g) ET01o. (h) ET02h. (i) ET05d.

**original etch:** Etched with the original recipe on line 1 in Table 3.5, results in Figure 3.12.

As mentioned above, the concern with the original etch is primarily the reentrant tapering of the silicon device layer and depth of the etch into the sacrificial silicon dioxide layer below. Of secondary concern is the damage to the resist: rounding of the top corners and the slanting of the edges. To etch less deeply into the buried oxide (BOX) layer the plan is to reduce the DC bias from  $\sim 68$  V to  $\sim 58$  V, a change of almost 15%. Everything else will be kept the same.

**antb06x:** Etched with the corresponding recipe on line 2 in Table 3.5, results in Figure 3.13a.

The results of this etch shown in the tilted SEM image in Figure 3.13a show very little difference from the previous etch. The sidewalls are slightly straighter and the etch into the BOX might be slightly less, but the end result is a negligible change. For the next etch we will increase the  $C_4F_8/SF_6$  ratio, which will slow the etch down and passivate the sidewalls which are both things we need to change about the etch. The increase will be small, going from a ratio of  $2.5 \rightarrow 2.667$ . To see the effects from this change only we will return the DC bias to its original value of  $\sim 68$  V.

**antb06y:** Etched with the corresponding recipe on line 3 in Table 3.5, results in Figure 3.13b.

The results shown in Figure 3.13b are quite similar to the last etch. For the next etch we will

combine the past two changes and see the results of the combination.

**ET01h:** Etched with the corresponding recipe on line 4 in Table 3.5, results in Figure 3.13c.

The results are shown in Figure 3.13c using an etched sample that was then cleaved before being undercut with HF. The patterns were fairly sparse on the chip and it is difficult to cleave exactly where you want given the size of these devices; however, the rewards for cleaving through features you are trying to etch can be very informative. Of course the cleave needs to be done before undercutting. Undercutting the cleaved sample can sometimes produce better images (or at least make it easier to take better images) but information about the depth of the etch into the BOX is lost. The images are of the rectangle cutouts of the phononic shield (the cross structure) and are chosen since they should have a uniform cross-section (unlike a circle for example). The most striking feature of the etch is the reentrant tapering as mentioned before, and it is quite clear in these images. We have been making fairly modest adjustments to the recipe thus far and have seen only small improvements. For the next etch we will make a much larger jump in the  $C_4F_8/SF_6$  ratio going from  $2.667 \rightarrow 3.2$  in hopes of seeing more of a change in results. We keep the DC bias the same at  $\sim 58$  V.

**ET01m:** Etched with the corresponding recipe on line 5 in Table 3.5, results in Figure 3.13d.

The results are shown in Figure 3.13d and are a tilted image of a pre-undercut sample. We would have been better served looking at a cleaved cross-section as in the last etch and a good reason for doing this can't be given. However, the results are starting to look better. The etch into the BOX is less and the sidewalls look better. There is some concern of booting—that is a small lip of the silicon at the silicon-BOX interface. In this image it is quite small to possibly non-existent, but to be sure we increased the time slightly from 4:00 (m:ss) to 4:20 for the next etch.

**ET02e:** Etched with the corresponding recipe on line 6 in Table 3.5, results in Figure 3.13e.

The results are shown in Figure 3.13e with a cleaved and not yet undercut sample giving a good view of the cross section of the cross rectangles. There is a tiny amount of booting right at the silicon-BOX interface, justifying our concern from the previous sample and a small amount of noticeable reentrant tapering. To combat both the tapering and booting we will again increase the  $C_4F_8/SF_6$  from  $3.2 \rightarrow 3.6$  and the time from 4:20  $\rightarrow$  5:00. The time needs to be increased to not only combat the booting of this etch but to counteract the slowdown of the etch that will result from adding more  $C_4F_8$  to the gas chemistry.

**ET02f:** Etched with the corresponding recipe on line 7 in Table 3.5, results in Figure 3.13f.

The results are shown in Figure 3.13f, which is of a cleaved and undercut sample. The reentrant tapering has turned into an inward tapering, suggesting the increased  $C_4F_8$  was more than we compensated for with the increase of the time and either the time must be increased again or the

| Name          | Date<br>dd/mm/yy | Time<br>(mm:ss) | C <sub>4</sub> F <sub>8</sub><br>(sccm) | CF <sub>6</sub><br>(sccm) | C <sub>4</sub> F <sub>8</sub> /SF <sub>6</sub> | RF Forward<br>(W) | DC Bias<br>(V) | ICP<br>(W)  | He Backing<br>(Torr; sccm) | Carrier |
|---------------|------------------|-----------------|-----------------------------------------|---------------------------|------------------------------------------------|-------------------|----------------|-------------|----------------------------|---------|
| original etch | –                | 4:00            | 28 [30]                                 | 12 [12]                   | 2.5                                            | –                 | 68/69          | 600         | 10 [10/-]                  | –       |
| atnb06x       | 16/01/13         | 4:00            | 28 [30]                                 | 12 [12]                   | 2.5                                            | 11 [12/2]         | 58-60          | 600 [582/2] | 10 [10/5.5]                | AJS #2  |
| atnb06y       | 16/01/13         | 4:00            | 30 [32]                                 | 12 [12]                   | 2.667                                          | 12.8 [14/2]       | 68             | 600 [582/2] | 10 [10/5.3]                | AJS #2  |
| ET01h         | 30/01/13         | 4:00            | 30 [32]                                 | 12 [12]                   | 2.667                                          | 10.5-10 [12/2]    | 58-60          | 600 [582/1] | 10 [10/5.5]                |         |
| ET01m         | 11/02/13         | 4:00            | 30 [32]                                 | 10 [10]                   | 3.2                                            | 11.1-10.7 [13/2]  | 58-60          | 600 [582/2] | 10 [10/5.5]                | AJS #1  |
| ET02e         | 14/02/13         | 4:20            | 30 [32]                                 | 10 [10]                   | 3.2                                            | 11 [12/2]         | 58             | 600 [582/2] | 10 [10/5.2]                | AJS #1  |
| ET02f         | 19/02/13         | 5:00            | 34 [36]                                 | 10 [10]                   | 3.6                                            | 11.7 [13/2]       | 57/58          | 600 [582/2] | 10 [10/5.4]                | AJS #2  |
| ET01o         | 19/02/13         | 5:00            | 32 [34]                                 | 10 [10]                   | 3.4                                            | 11.3 [13/2]       | 57/58          | 600 [582/2] | 10 [10/5.5]                | AJS #1  |
| ET02h         | 20/02/13         | 5:30            | 32 [34]                                 | 10 [10]                   | 3.4                                            | 11 [12/2]         | 57/58          | 600 [582/2] | 10 [10/5.7]                | AJS #1  |
| ET05d         | 09/03/13         | 5:30            | 32 [34]                                 | 10 [10]                   | 3.4                                            | 11.0 [12/2]       | 57-60          | 600 [582/2] | 10 [10/5.3]                | AJS #1  |

Table 3.5: **Si etch re-optimization parameters.** The parameters used in re-optimizing the silicon nanobeam etch. The table includes all etches in the order they were done to optimize the etch.

$C_4F_8$  decreased. There is a concern that increasing the time too much will just add to overall blowout of the holes, which, while they may still be vertical, is generally unwanted. The  $C_4F_8/SF_6$  ratio is also quite high, and as such we will decrease this ratio to attempt to remove the inward slanting sidewalls; changing the ratio from  $3.6 \rightarrow 3.4$  and keeping everything else including the time the same.

**ET01o:** Etched with the corresponding recipe on line 8 in Table 3.5, results in Figure 3.13g.

The results are shown in Figure 3.13g, cleaved and not undercut. The sidewalls appear straighter but still have an inward slope and some booting can be seen. To combat this slight inward slope and booting the time of the etch will be increased from 5:00  $\rightarrow$  5:30 with all other parameters kept the same. The rationale is that the etch itself is pretty good at this point and if we just increase the time a little it might be enough to fix the problems we see.

**ET02h:** Etched with the corresponding recipe on line 9 in Table 3.5, results in Figure 3.13h.

The results are shown in Figure 3.13h, cleaved and not undercut. The results of this etch are quite good. There appears to be a very small amount of blowout and a tiny amount of booting on some surfaces, but overall quite good. For the next etch we change none of the etch parameters, however we change the etch pattern. Up until this point we have been etching test chips with the actual devices on them, cleaving them and hoping we get a good cleave showing a nice cross-section. While developing the etch it isn't necessary to etch the final pattern but instead could be useful to etch a test pattern with similar geometries that are easier to analyze.

**ET05d:** Etched with the corresponding recipe on line 10 in Table 3.5, results in Figure 3.13i.

The results are shown in Figure 3.13i, cleaved and not undercut. The trench sidewalls are very straight and the depth into the BOX layer is very minimal. It is important to note that the E-beam written width of these trenches was 110 nm and the roughly measured width of them is 150 nm. It is promising that we can etch through a small trench with such straight sidewalls all the way down. This test pattern has trench widths ranging from 100 nm to 300 nm in increments of 10 nm. Many of these test patterns were placed next to each other on the chip to create a 1.5 mm long set of trench with very little space between them. This then maximized the likelihood that the cleave would go through the trenches. In the direction perpendicular to the cleave we did a dose run of the trenches to further examine the effects of dose on the trenches of different sizes. The larger trenches start showing a small reentrant taper with increasing trench width, which is expected. However, this pattern will let us easily see the effects of the etch over a wide variety of trench widths and finely tailor the etch to the pattern we actually wish to write. With this new pattern in hand we move on to the issue of developing a silicon nitride etch in the next section after first re-iterating the key points learned.



Figure 3.14: **Etch test pattern.** The etch test pattern used for etch optimization of the silicon nitride nanobeams. The pattern consists of groups of four trenches a full FoV long. The trench width varies over each group of four with the first group having a width of 100 nm, the last having a width of 300 nm with the width increasing by 10 nm each group. The squares separating each set of four groups provide a count allowing easier determination when imaging (SEMing).

### 3.3.3.3 Key Points and Lessons Learned

1. Know what you want to image. Over the course of the optimization we switched from tilted images to cross-sections. Both can be valuable, but the cross-sections provided more of the information we wanted and we should have used them the entire time.
2. Image consistently. It is best to image the same part of the structure at the same magnification and make sure the images are high quality. It is a good idea to image multiple areas of the structure consistently.
3. Consider a test pattern. It seems better to use the test pattern of Figure 3.13i and with the entire pattern shown in Figure 3.14. The final design can be etched occasional along the way or at the end to verify the results of the test pattern.
4. Record all information. It is sometimes hard to know what is needed now, or in the future what information you wished you had recorded.

### 3.3.4 Silicon Nitride Etching and Optimization




Figure 3.15: **Original silicon nitride etch.** (a) Tilted image of the original  $\text{Si}_3\text{N}_4$  etch with the resist, 400 nm  $\text{Si}_3\text{N}_4$  device layer, and base Si layer shown. The etch is very deep in the open trench and there is a reasonable amount of resist left. There is clear degradation in the resist with sloping sides. (b) The same chip after undercutting to determine if the holes were etched through and they are clearly not. This etch is very aggressive in open areas but much slower in the small area holes. (c) A look at the crosses of the phononic shield shows significant booting similar to the holes.

The original silicon nitride etch was developed in the group mainly for zipper optomechanical devices [36, 46] and has always been a tough etch due to the nearly 1:1 selectivity and the thickness of the silicon nitride device layer (400 nm). It has, however, served its purpose well producing many nice results in the group [19, 46, 53, 54]. Starting to make the nanobeams [39, 45] in silicon nitride,

coupled with the recent reoptimization of the silicon etch, provided an excellent opportunity to take a look at the silicon nitride etch and attempt to optimize it.

The current silicon nitride etch is shown in Figure 3.15. The resist, device, and substrate silicon layers are shown in Figure 3.15a with a black line showing the delineation of the different layers. The results of this etch were better than we expected, with what looked like quite a bit of resist left when it was generally a given that the nitride etch consumed all or almost all the resist. The depth which the silicon substrate was etched was also quite surprising; a testament to the aggressiveness of the nitride etch when exposed to large open areas of silicon. The resist itself looks like it has been strongly attacked and degraded, showing a lot of sloping in its edges. To see if the holes were through, the resist was stripped and the devices undercut. It is clear from Figure 3.15b that the holes are not through despite the depth of the open trenches. Looking at the crosses from the phononic shield, significant booting can also be seen as in Figure 3.15c. Given these problems, the results of the etch look promising with room for optimization. For the optimization of the etch, the test pattern shown in Figure 3.14 will be used. In the post etch images the columns will correspond to written trench widths of 100 nm, 200 nm, and 300 nm respectively, while each row will be a different test chip/etch as in Figure 3.16.

**nmb01a:** Etched with the original recipe on line 1 in Table 3.6, results in Figure 3.15.

The results are shown in Figure 3.15 and were discussed above. The results looked promising and with the amount of resist left gave hope of room to optimize further. The holes were clearly not through and therefore we will etch slightly longer with the first test pattern and look at the results. The time of the etch will be increased from 4:50  $\rightarrow$  5:20.

**NET01c:** Etched with the corresponding recipe on line 2 in Table 3.6, results in Figure 3.16a-c.

The results are shown in Figure 3.16a-c, and all three trench widths show scalloping in the sidewalls. The 100 nm trench shown in Figure 3.16a shows the most significant scalloping decreasing with increasing trench width. It is possible this is a combination of the angular dependence of the incoming ions, the selective masking by the different trench widths combined with too little passivation on the sidewalls. For the narrowest trench, the filtering of off-normal impact ions is much higher than with the widest trench. As the resist is worn down, the off normal ions can attack the nitride trench lower and lower, leading to the scalloping observed. The widest trench produces very little filtering of the off-normal impact ions, and as such much more of the nitride sidewall is attacked all the time leading to less scalloping. This only works because the sidewall passivation is insufficient given the strength of the plasma etch. A similar idea was recently used in a controlled manner to intentionally scallop pillars [55]. Regardless of the exact reason, there is clearly insufficient passivation and as such we will increase the  $C_4F_8$ , changing the  $C_4F_8/SF_6$  ratio from 2.0  $\rightarrow$  2.5. All other parameters will be left the same.






Figure 3.16: **Progression of the silicon nitride etch optimization.** Results of the test etches in order of the optimization. Refer to Table 3.6 for parameters. From left to right the written trench widths are 100 nm, 200 nm, and 300 nm. (a-c) Results of the NET01c test etch. (d-f) Results of the NET01d test etch. (g-i) Results of the NET01a test etch.

**NET01d:** Etched with the corresponding recipe on line 3 in Table 3.6, results in Figure 3.16d-f.

The results are shown in Figure 3.16d-f are quite promising, showing the increased passivation worked well. There is still a significant amount of resist left and its sidewalls look reasonably straight at the resist-nitride interface. The etch itself seems to work well for the larger trench widths and is almost through for the smallest. Being through for the smallest written width of 100 nm is not necessary for the final etch since this width is much smaller than the smallest feature in the silicon nitride nanobeam design. Some slanting in the sidewalls of the nitride is present. There isn't much to correct here. In an attempt to reduce blowout, and perhaps produce straighter sidewalls, we will increase the  $C_4F_8/SF_6$  ratio from  $2.5 \rightarrow 3.2$  and also increase the etch time to compensate for the increased passivation.

**NET01a:** Etched with the corresponding recipe on line 4 in Table 3.6, results in Figure 3.16g-i.

The results are shown in Figure 3.16g-i show less blowout in general than the last etch with similar etch depths. The sidewalls show much similarity in straightness as the last etch and there is still a lot of resist. Overall this etch is quite good and a reasonable place to stop the optimization



for now and test the etch on an actual nanobeam pattern. A slight improvement could be to etch a little longer or slightly increase the DC bias. Further optimization steps were taken with an ICP power of 600 W and are still being explored. The results so far are shown in Figure 3.17 in a similar manner as before with the associated etch parameters in Table 3.6.

### 3.3.4.1 Key Points and Lessons Learned

1. Don't be afraid to take large steps in parameter space. While small steps are safe, taking big steps can help set bounds on the parameter space and allows you to really see the effect of changing a specific parameter.

## 3.4 Wet Chemistry

### 3.4.1 Standard Solvent Clean

The current recipe is a 10 minute bath in heated TCE with a spin bar at 250 rpm (hot plate set to 80°C). The chips are removed and, before being allowed to dry, thoroughly rinsed with acetone from a spray bottle followed by an IPA rinse from a spray bottle. The chips are set down on a tek-wipe, held with tweezers and blown dry with nitrogen ( $N_2$ ). It is best to never let the dirtier solvents (TCE, acetone) dry on a chip as it can be hard to get the residue off. The solvent clean goes from most powerful and dirtiest to least aggressive and cleanest. We used to end the clean with deionized (DI) water before the  $N_2$ , which may be a good idea. We removed that step as we believed our IPA should be very clean.

### 3.4.2 Resist Strip

Following the etch step there is often a large amount of resist left on the chip. For the silicon etch, the initial resist thickness is  $\sim 350$  nm and after the etch there is a substantial 250 – 300 nm remaining (see Figure 3.13i for a good example). The removal of this resist is done in two steps. The first is the removal of the majority of the resist in either a TCE bath or using Zeon's ZEP remover, ZDMAC. The second step is using a piranha solution to fully clean away any remaining resist on the chip. Two steps are used because it has been found that placing a sample with such a large layer of etched resist into a piranha solution tends to remove it, but leaves a resist slime layer that can then redeposit on the surface. This appears to be due to the thickness of the resist and the way it is burnt in the etch process. While a long enough piranha clean would almost certainly fully dissolve this material, we use two steps and let the piranha do the final clean. The first step alone, either TCE or ZDMAC, does not fully remove the resist. It will remove the majority of the resist but tends to leave a "whispy", smooth, thin layer coating the surface or part of the surface.

| Name   | Date<br>dd/mm/yy | Time<br>(mm:ss) | C <sub>4</sub> F <sub>8</sub><br>(sccm) | SF <sub>6</sub><br>(sccm) | C <sub>4</sub> F <sub>8</sub> /SF <sub>6</sub> | RF Forward<br>(W)     | DC Bias<br>(V) | ICP<br>(W)   | He Backing<br>(Torr; sccm) | Carrier |
|--------|------------------|-----------------|-----------------------------------------|---------------------------|------------------------------------------------|-----------------------|----------------|--------------|----------------------------|---------|
| NNB01a | 25/02/13         | 4:50            | 22 [24]                                 | 12 [12]                   | 2.0                                            | 24.4-24.1 [25/4]      | 81-82          | 1000 [969/2] | 10 [10/5.5]                |         |
| NET01c | 25/02/13         | 5:20            | 22 [24]                                 | 12 [12]                   | 2.0                                            | 26.3-25.4 [27/6]      | 81-82          | 1000 [969/2] | 10 [10/5.4]                | JC #2   |
| NET01d | 28/02/13         | 5:20            | 28 [30]                                 | 12 [12]                   | 2.5                                            | 25.8,26.2,25.5 [27/7] | 81-82          | 1000 [969/2] | 10 [10/5.6]                | JC #2   |
| NET01a | 28/02/13         | 6:00            | 30 [32]                                 | 10 [10]                   | 3.2                                            | 26.5 [27/6]           | 81-82          | 1000 [969/2] | 10 [10/5.5]                | JC #1   |
| NET01b | 28/02/13         | 6:30            | 30 [32]                                 | 10 [10]                   | 3.2                                            | 18.4-17.3 [19/4]      | 81-82          | 600 [582/2]  | 10 [10/5.4]                | JC #2   |
| NET01e | 04/03/13         | 11:00           | 30 [32]                                 | 10 [10]                   | 3.2                                            | 17.5-16.2 [18/3]      | 81-82          | 600 [582/2]  | 10 [10/5.4]                | JC #1   |
| NET01f | 04/03/13         | 15:00           | 34 [36]                                 | 10 [10]                   | 3.6                                            | 18-16.8 [18/3]        | 81-82          | 600 [582/2]  | 10 [10/5.4]                | JC #2   |
| NET01g | 04/03/13         | 17:00           | 38 [40]                                 | 10 [10]                   | 4.0                                            | 19.8-18.1 [21/4]      | 81-82          | 600 [582/2]  | 10 [10/5.4]                | JC #1   |
| NET01h | 06/03/13         | 15:00           | 38 [40]                                 | 10 [10]                   | 4.0                                            | 22-20.6 [22/4]        | 91-92          | 600 [582/2]  | 10 [10/5.4]                | JC #1   |
| NET01i | 06/03/13         | 15:00           | 36 [38]                                 | 10 [10]                   | 3.8                                            | 24.5-22 [25/4]        | 101-102        | 600 [582/2]  | 10 [10/5.4]                | JC #2   |

Table 3.6: **SiN etch optimization parameters.** The parameters used in optimizing the silicon nitride nanobeam etch. The table includes all etches in the order they were done.

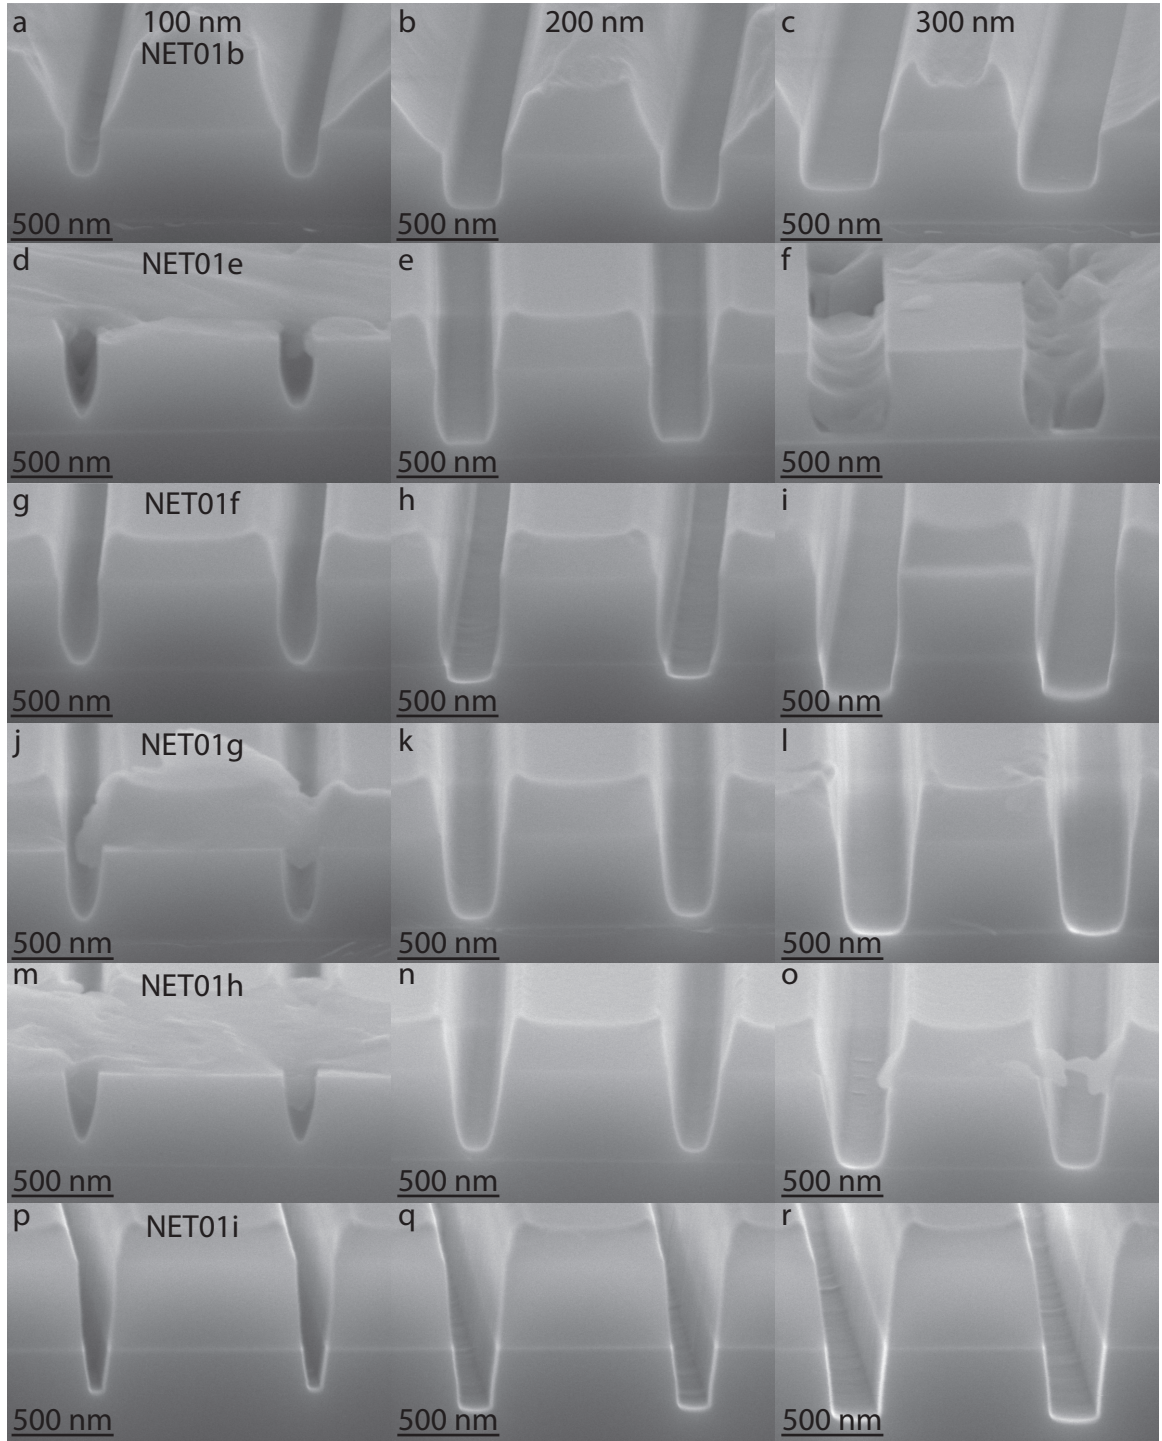


Figure 3.17: **Progression of the silicon nitride etch optimization with 600 W ICP power.** Results of the test etches in order of the optimization. Refer to Table 3.6 for parameters. From left to right the written trench widths are 100 nm, 200 nm, and 300 nm. (a-c) Results of the NET01b test etch. (d-f) NET01e. (g-i) NET01f. (j-l) NET01g. (m-o) NET01h. (p-r) NET01i.

This can be seen in Figure 3.18, where several images of the same chip show the surface with only a TCE bath (**a-c**) and then after a piranha clean, after that TCE bath (**d-f**). With only the TCE, a thin layer of resist can be seen on the chip (note the lines on the surface) and also in some of the crosses. After a subsequent piranha clean the silicon surface looks clean and the holes are free from material. The specifics of the TCE bath and ZDMAC initial resist strip and final piranha polish will be outlined next. First, a note about removing the etched chip from the carrier wafer.




Figure 3.18: **Comparison of resist stripping.** (**a-c**) SEM of sections of the phononic shield after stripping resist only with hot TCE (80° C, 10 minutes, 250 rpm stirring). Some resist can clearly still be seen. (**d-f**) SEM of the same chip, showing devices neighboring the ones above, after a piranha, after the TCE step (10 minutes). The sample is much cleaner.

#### 3.4.2.1 Removing the Etched Chip from the Carrier Wafer

After the etch, the chip will be more firmly secured to the carrier wafer than before due to the heat experienced in the etch chamber and its effects on the thermal paste. The best way to remove the chip is not to attempt to lift it straight off the wafer but to grasp it on the side with carbon fiber tweezers and slide the chip parallel to the sides it is being grasped (gently wiggling it side-to-side a bit). As you are sliding the chip, gently begin applying pressure to lift it off the surface (likely need to gently hold the edge of the carrier wafer). This sliding and lifting motion should allow the chip to come easily free without excess application of force, which could cause the tweezers to slip and scratch the chip or carrier wafer (or even cause the chip to be thrown across the wet bench).

In order to properly image the resist layer of the chip, the thermal past on the back of the chip must be removed without damaging the top resist layer. Importantly, solvents and even their fumes can damage this resist and make learning about the effects of the etch on the resist difficult

or impossible. To remove the thermal paste without damaging the resist, a small amount of TCE in a beaker or tin is required along with 2-3 soft tipped applicators (like q-tips but without the individual threads). While still holding the chip edges with the tweezers, dip the applicator tip into the TCE making it wet (you may need to press the tip against a tek-wipe if it is too wet). With small strokes, sweep the tip across the backside, removing the thermal paste. This should be done in strokes that do not approach the edge of the chip (we do not want TCE wicking over the edge to the front surface) and each of the two sides of the applicator can generally be used twice, once in each direction. Attempt to get most of the thermal paste off this way. A final applicator with very little or no TCE can be used to gently scrub and clean the back surface completely. The chip can now be SEMed or stored without fear of contamination from the thermal paste and the resist layer should be intact.

**Note:** Even if you do not intend to SEM the chip with resist on, it is good practice to carefully remove the resist every time as it does take some practice and skill to do well and easily.

### 3.4.2.2 TCE Bath or ZDMAC Clean

To clean the bulk of the resist off the chip we have historically used a TCE bath. A Petri dish filled with TCE is placed on a hot plate which is set to 80°C; a stir bar is placed inside and the rotational speed set to 250 rpm. It is best to have the stir bar sit off to one side and then place the chips on the other (the currents created by the stir bar can be powerful and if the chips are placed too close, disaster can strike). The TCE is given ~10 minutes to heat up. By grabbing the chip on the edge with carbon fiber tweezers, it is inserted into the TCE, resist side up, and held very near the stir bar for ~15–30 seconds, allowing strong currents of hot TCE to wash over the top and bottom. The chip is then placed far enough from the stir bar that it isn't caught in the current and moved, but close enough to have flow over it. Multiple chips can be done at once. The chips should sit in the bath for 10 minutes and the TCE level should never be allowed to get so low such that the chips may have TCE residue dry on their surface. The chips are then given a standard solvent rinse and blown dry as outlined in Section 3.4.1.

**Note:** TCE is very volatile and toxic. It will evaporate quickly when heated and should be kept far back in a well-vented fume hood.

While TCE is a very powerful solvent, it is also very toxic and we are attempting to replace it in our processing where possible. For the removal of resist, Zeon, the company that makes the resist, sells a removal chemical: ZDMAC. We can replace the TCE bath with a step using ZDMAC and get similar results in less time. To use the ZDMAC, a small beaker is filled with ZDMAC. The chip is then grasped on the edges with carbon fiber tweezers and inserted fully into the ZDMAC and shaken gently for 60 seconds. It is then removed and given a standard solvent rinse and blown dry (see Section 3.4.1).

**Note:** ZDMAC reacts with the tek-wipes differently than other solvents (which tend to just evaporate). It will strongly attack them and turn the blue tex-wipes white. Care should be taken when using it and the area properly cleaned up (including removal of all stained tex-wipes). With the bulk of the resist removed, the next step is a strong clean of the surface with piranha.

### 3.4.2.3 Piranha

Piranha solution is a very powerful cleaning solution that gets its name from its ability to easily remove organic matter due to its properties as both an acid and strong oxidizing agent. Piranha is a mixture of sulfuric acid ( $\text{H}_2\text{SO}_4$ ) and hydrogen peroxide ( $\text{H}_2\text{O}_2$ ) combined, in our case, in a ratio of 3:1 after the sulfuric acid is heated to  $\sim 75^\circ\text{C}$  on a hot plate. When the hydrogen peroxide is added to the sulfuric acid the reaction is very energetic and exothermic easily bringing the temperature of the solution to  $> 120^\circ\text{C}$ , with visible piranha mist coming off the liquid. For these reasons all safety precautions should be taken including using in a well-vented fume hood with the sash down, wearing acid gloves, a mask, and acid apron, and ensuring the temperatures are within reasonable limits and the ratio of the chemicals is correct. Incorrect ratios or temperatures can lead to extremely violent reactions and the solution could explode (I'm told). The specifics of how we use and prepare piranha are outlined next.

To make the piranha we start with a 100 mL Pyrex beaker, placing a 1" stir bar inside and then adding 50 mL of sulfuric acid (resulting in  $\sim 45$  mL of acid in the beaker when accounting for the stir bar's volume). The stir bar is added first so that it doesn't need to be dropped into the acid, which could cause the acid to splash. This beaker is next placed on a hot plate which is set to  $145^\circ\text{C}$  with a rotational speed of 250 rpm. At this stage it is best to ensure that the supporting clamps and metal pieces are in the correct location to hold the sample in the piranha—when the piranha is active we do not want to be fine tuning the location of a clamp or swapping out one piece for another. A thermometer may be used to measure the temperature of the sulfuric acid, however extensive experience has indicated it takes  $\sim 12$  minutes to reach the desired temperature (this of course depends on many things including glassware, hot plate, etc, and must be properly calibrated). While the sulfuric acid is heating up the hydrogen peroxide is measured in a graduated cylinder (specific for this purpose) and we need 15 mL to give a 3:1 ratio. Also, while the acid is heating, the chip is mounted in a Teflon holder which will be used for the piranha and any following steps (such as HF undercut and piranha/HF cleans).

When the sulfuric acid is done heating ( $\sim 12$  minutes or  $\sim 75^\circ\text{C}$ —a small variation in starting temperature is not a problem) the hydrogen peroxide is poured in slowly, instantly starting an exothermic reaction that visibly produces gas bubbles. It can be a good idea to have a Plexiglas shield between you and the piranha, or surrounding as much of the solution as possible, as the mist coming off it is acidic. The sample is next attached to the support structure and inserted into the bubbling

piranha to be cleaned. The standard clean lasts 10 minutes over which time the temperature will drop significantly (usually below  $100^{\circ}\text{C}$ ) and the bubbling will stop. While the sample is in the piranha the rinse beakers can be filled. We use two 150 mL Pyrex beakers for the rinse—a first rinse and a second rinse. In general it is a good idea to have the rinse larger than the acid such that more of the sample holder can be submersed in the rinse to neutralize any acid mist that has gotten onto the sample holder (acid mist tends to get on everything from a piranha). Upon removing the sample holder from the piranha it is immediately inserted into the first rinse where it is lightly shaken and moved around. Confined spaces in the sample holder can trap acid or air bubbles, which protect the acid from being neutralized in the water and it is these spaces we wish to ensure are rinsed properly. After 30 seconds in the first rinse the sample holder is immediately placed into the second rinse with the same result in mind.

**Note:** While nominally the sample holder will be completely acid free after two 30 second water rinses, this is almost never entirely true. Blind holes in the sample holder are especially good at trapping acid behind an air bubble so that it does not see the water rinse. The space under the chip, and other small openings, also tend to trap acid. Care should always be taken when removing the sample holder from the second rinse, transporting it and removing the chip (or whatever the subsequent step is).

**Note:** Plan the location of all the beakers holding chemicals and water before doing any work. Understand where your hands will have to move to reach the chemicals and try to never have to reach over a beaker of acid or water if possible. Careful planning of the location of the beakers and movements of your hands to complete your task, done before starting, will pay off later and keep everything safer.

### 3.4.3 HF Undercut

To release the etched and cleaned structure, a HF undercut is performed. The chips we use are SOI with a device layer of 220 nm, and a  $3\text{ }\mu\text{m}$  oxide on top of a silicon substrate (on the order of  $575\text{ }\mu\text{m}$  thick). By removing the oxide under and around the etched structures they are free to move mechanically, and the optical fields we are interested in see only the designed structure and air—with the next silicon layer  $3\text{ }\mu\text{m}$  below; far enough away to not be relevant. It is this reason that SOI is such a great material for photonics.

In general the HF undercut step comes directly after the first piranha clean which removes the remaining resist. As such, the sample is generally still attached to the sample holder and has just had two water rinses after the piranha clean. The sample holder can then be directly placed in a small plastic (high-density polyethylene) beaker containing the HF. For the undercut we use 48% HF, undiluted from the bottle we buy it in. The time of the undercut is 90 seconds (1:30) and is long enough to remove almost (80-90%) all of the oxide layer under the exposed structure. After the

undercut the sample holder is transferred to two 150 mL plastic water rinse beakers for a 30 second water rinse in each, similar to the water rinse after the piranha step.

### 3.4.4 Surface Preparation

The surface preparation of silicon has been long studied thanks to the interest by the electronics industry, and more recently MEMs and photonics communities. For the sub-field of optomechanics, where we are interested in both the optical loss of the bulk and surface scatter of photons, but also how the bulk and surface properties affect the phonons we wish to confine and manipulate. Currently, after the HF undercut we perform two piranha/HF clean cycles which has been proven to be an effective surface termination from the perspective of optical scattering (see [56]). This involves putting the sample into a fresh piranha after the undercut, again for 10 minutes, and following the same procedure through to the rinse stage. The HF stage is the same as for the HF undercut, however instead of 48% HF, the sample is immersed in a 10:1 solution of  $\text{H}_2\text{O}$ :48% HF we make ourselves. This lasts one minute and then the same rinse steps are followed. A fresh piranha can then be made and the step repeated. The important aspect of this surface termination is the final one minute 10:1 HF dip. The piranha clean tends to grow a thin thermal oxide and leaves the silicon surface terminated with OH. The 10:1 HF dip removes this layer and leaves the surface nicely terminated with H atoms, leading to better optical and hopefully mechanical properties. This is an area open to active study and significant improvements could be imagined.



## Chapter 4

# Experimental Demonstration of Coherent Optomechanical Wavelength Conversion

The ability to coherently convert photons between disparate wavelengths has broad technological implications, not only for classical communication systems, but also future quantum networks [57–59]. For example, hybrid quantum networks require a low loss interface capable of maintaining quantum coherence while connecting spatially separate systems operating at incompatible frequencies [60]. For this reason, photons operating in the low loss telecommunications band are often proposed as a conduit for connecting different physical quantum systems [61]. It has also been realized that a wide variety of quantum systems lend themselves to coupling with mechanical elements. A coherent interface between mechanics and optics, then, could provide the required quantum links of a hybrid quantum network [62].

Until now, most experiments demonstrating both classical and quantum wavelength conversion have utilized intrinsic optical nonlinearities of materials [63–68]. The nonlinear interaction of light with acoustic or molecular mechanical vibrations of materials, for instance, enables a great many optical functions used in high-speed optical communication systems today [64]. With the technological advancements in the fields of nanomechanics and nanophotonics, it is now possible to engineer interactions of light and mechanics. Progress in this area has included enhanced nonlinear optical interactions in structured silica fibers [69], near quantum-limited detection of nanomechanical motion [70], and the radiation pressure cooling of a mesoscopic mechanical resonator to its quantum ground state of motion [30, 70]. Coupling of electromagnetic and mechanical degrees of freedom, in which the coherent interaction rate is larger than the thermal decoherence rate of the system as realized in the ground-state cooling experiments, opens up an array of new applications in classical and quantum optics. This realization, along with the inherently broadband nature of radiation pressure, has spawned a variety of proposals for conversion between photons of disparate frequen-

cies [31, 32, 62, 71] through interaction with a mechanical degree of freedom—proposals which are but one of many possible expressions of hybrid optomechanical systems [60]. Such cavity optomechanical systems are not just limited to the optical frequency domain, but may also find application to the interconversion of microwave and optical photons [31, 72], enabling a quantum-optical interface to superconducting quantum circuits [73].

Here we demonstrate optical wavelength conversion utilizing a simple, hybrid optomechanical system consisting of an acoustic and optical resonator formed from the top silicon device layer of a silicon-on-insulator wafer, typically used in the microelectronics industry. Through nanoscale lithographic patterning, an optomechanical crystal (OMC) resonator [45] is formed which supports a 4 GHz mechanical resonance co-localized with two optical resonances in the S and C telecommunications bands. The extreme localization of both acoustic and optical energy in this OMC resonator results in a strong radiation pressure interaction between both optical modes and the mechanical motion of the resonator, enabling conversion of light between the two optical cavity modes which span a frequency of 11.2 THz. Optical wavelength conversion is demonstrated over the 1.5 MHz bandwidth of the mechanical resonator at a peak internal efficiency exceeding 90% (the end-to-end measured efficiency is limited to 2% by the optical fiber coupling efficiency to the optical cavities), and with a thermal-limited noise of only 6 quanta, well above the quantum-limited noise of  $4 \times 10^{-3}$  quanta.

## 4.1 Optomechanical Wavelength Conversion

The detailed mathematical description of wavelength conversion in an optomechanical system was laid out in the theory section 2.3. In the section that follows, a higher level description is given, stressing a more physical understanding and highlighting the most relevant equations required to understand the wavelength conversion process and results presented. As illustrated in Figure 4.1a, conceptually the proposed system for wavelength conversion may be thought of as consisting of two optical cavities coupled to the same mechanical resonator. Motion of the mechanical resonator induces a shift proportional to the amplitude of motion in the resonance frequency of each of the optical cavities, corresponding to the usual radiation pressure interaction in a cavity-optomechanical system. In this case the two optical cavities are assumed to have different resonance frequencies, between which the desired optical wavelength conversion occurs. For each of the optical cavities, a pump laser beam red-detuned by the mechanical resonance frequency is used to couple the optical cavity modes to the mechanical resonator via the radiation pressure interaction. This results in a coherent mixing of the optical and mechanical degrees of freedom. Under these conditions, the shared mechanical resonator acts as a bridge between the two optical cavities, whereby incident optical radiation resonant with one optical mode is converted into mechanical motion, which is then

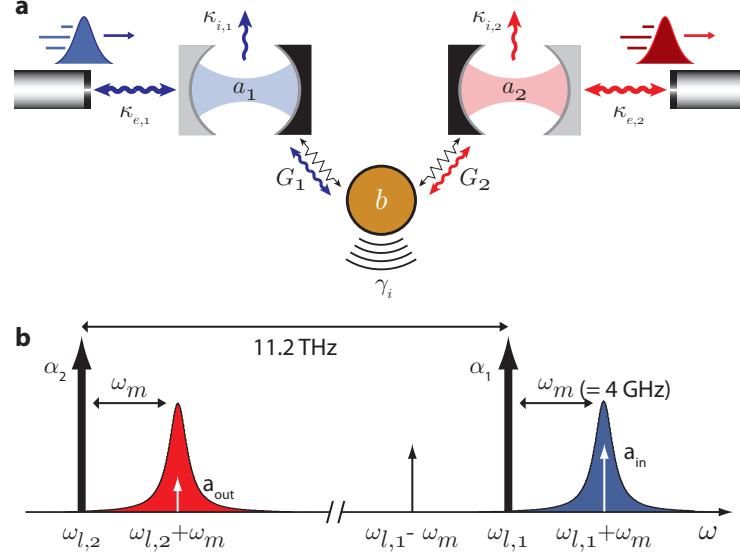


Figure 4.1: **System model.** (a) Diagram of the wavelength conversion process as realized via two separate Fabry-Pérot cavities. The two optical cavity modes,  $a_1$  and  $a_2$ , are coupled to the same mechanical mode,  $b$ , with coupling strengths  $G_1$  and  $G_2$ , respectively. The optical cavity modes are each coupled to an external waveguide (with coupling strengths  $\kappa_{e,1}, \kappa_{e,2}$ ), through which optical input and output signals are sent. The optical cavities also have parasitic (intrinsic) loss channels, labeled  $\kappa_{i,1}$  and  $\kappa_{i,2}$ , whereas the mechanical mode is coupled to its thermal bath at rate  $\gamma_i$ . (b) Schematic indicating the relevant optical frequencies involved in the wavelength conversion process. The cavity control laser beams, labeled  $\alpha_1$  and  $\alpha_2$ , are tuned to a mechanical frequency red of the corresponding optical cavity resonances. An input signal ( $a_{in}$ ) is sent into the input cavity at frequency  $\omega_{l,1} + \Delta_1$ . The input signal is converted into an output signal ( $a_{out}$ ) at frequency  $\omega_{l,2} + \Delta_1$  via the optomechanical interaction.

converted back into optical radiation at the frequency of the other optical mode. This process, akin to the four-wave mixing found in coherent anti-Stokes Raman spectroscopy [74], can result in near-unity conversion efficiency by both matching the pump-induced coupling rates of the mechanical mode to each of the optical cavity modes, and increasing the optomechanical coupling rate above that of the intrinsic loss rate of the mechanical resonator.

The interaction Hamiltonian describing the radiation pressure interaction between the two cavity modes and the mechanical resonator is given by  $H = \sum_k \hbar g_k \hat{a}_k^\dagger \hat{a}_k (\hat{b} + \hat{b}^\dagger)$ , where  $\hat{a}_k$  ( $\hat{b}$ ) are the annihilation operators for the optical cavity modes (common mechanical mode), and  $g_k$  is the optomechanical coupling rate between the mechanical mode and the  $k$ th cavity mode ( $k = 1, 2$ ). Physically,  $g_k$  represents the frequency shift of cavity mode  $k$  due to the zero-point motion of the mechanical resonator. Wavelength conversion is driven by two control laser beams ( $\alpha_k$  in Figure 4.1b), of frequency  $\omega_{l,k}$  and nominal detuning  $\delta_k \equiv \omega_k - \omega_{l,k} = \omega_m$  to the red of cavity resonance at frequency  $\omega_k$ . In the resolved sideband regime, where  $\omega_m \gg \kappa_k$  ( $\kappa_k$  the bandwidth of the  $k$ th cavity mode), the spectral filtering of each cavity preferentially enhances photon-phonon exchange. The

resulting beam-splitter-like Hamiltonian is  $H = \sum_k \hbar G_k (\hat{a}_k^\dagger \hat{b} + \hat{a}_k \hat{b}^\dagger)$  [31, 75], where  $G_k \equiv g_k \sqrt{n_{c,k}}$  is the parametrically enhanced optomechanical coupling rate due to the  $\alpha_k$  control beam ( $n_{c,k}$  is the control-beam induced intracavity photon number). In the weak-coupling limit,  $G_k \ll \kappa_k$ , this interaction effectively leads to an additional mechanical damping rate,  $\gamma_{\text{OM},k} = 4G_k^2/\kappa_k$ . The degree to which this optomechanical loss rate dominates the intrinsic mechanical loss is called the cooperativity,  $C_k = \gamma_{\text{OM},k}/\gamma_i$ . At large cooperativities, the optomechanical damping has been used as a nearly noiseless loss channel to cool the mechanical mode to its ground state [30, 70]. In the case of a single optical cavity system, it has also been used as a coherent channel, allowing inter-conversion of photons and phonons leading to the observation of electromagnetically-induced transparency (EIT) [25, 26]. In the double optical cavity system presented here, coherent wavelength conversion of photons results.

As shown in Figure 4.1a, each optical cavity is coupled not just to a common mechanical mode, but also to an optical bath at rate  $\kappa_{i,k}$  and to an external photonic waveguide at rate  $\kappa_{e,k}$  (the total cavity linewidth is  $\kappa_k = \kappa_{i,k} + \kappa_{e,k}$ ). The external waveguide coupling provides an optical interface to the wavelength converter, and in this work consists of a single-transverse-mode waveguide, bi-directionally coupled to each cavity mode. The efficiency of the input/output coupling is defined as  $\eta_k = \kappa_{e,k}/2\kappa_k$ , half that of the total bi-directional rate. Although the wavelength converter operates symmetrically, here we will designate the higher frequency cavity mode ( $k = 1$ ) as the input cavity and the lower frequency cavity ( $k = 2$ ) as the output cavity. As shown in Figure 4.1b, photons sent into the wavelength converter with detuning  $\Delta_1 \sim \omega_m$  from the control laser  $\alpha_1$ , are converted to photons  $\Delta_1$  detuned from the control laser  $\alpha_2$ , an 11.2 THz frequency span for the device studied here.

The details of the conversion process can be understood by solving the Heisenberg-Langevin equations (see Section 2.3 and [31, 71, 76]). We linearize the system and work in the frequency domain, obtaining through some algebra the scattering matrix element  $s_{21}(\omega)$ , which is the complex, frequency-dependent conversion coefficient between the input field at cavity  $\hat{a}_1$ , and the output field at cavity  $\hat{a}_2$ . This coefficient is given by the expression

$$s_{21}(\omega) = \sqrt{\eta_2 \eta_1} \frac{\sqrt{\gamma_{\text{OM},2} \gamma_{\text{OM},1}}}{i(\omega_m - \omega) + \gamma/2}, \quad (4.1)$$

where  $\gamma = \gamma_i + \gamma_{\text{OM},2} + \gamma_{\text{OM},1}$  is the total mechanical damping rate and equal to the bandwidth of the conversion process. From this expression, the spectral density of a converted signal  $S_{\text{out},2}(\omega)$ , given the input signal spectral density  $S_{\text{in},1}(\omega)$ , may be found and is given by

$$S_{\text{out},2}(\omega) = \eta_2 \eta_1 \frac{\gamma_{\text{OM},2} \gamma_{\text{OM},1}}{(\omega + \omega_m)^2 + (\gamma/2)^2} (n_{\text{added}} + S_{\text{in},1}(\omega)). \quad (4.2)$$

These spectral densities have units of photons/Hz  $\cdot$  s and are proportional to optical power. The

added noise,  $n_{\text{added}}$ , arises from thermal fluctuations of the mechanical system and the quantum back-action noise of light present in each optical mode. If we ignore the input and output coupling efficiencies, or assume a system with ideal cavity-waveguide coupling, ( $\eta_1, \eta_2 = 1$ ), the peak internal photon conversion efficiency is given by

$$\eta_{\text{max,int}} = \frac{4C_1C_2}{(1 + C_1 + C_2)^2}. \quad (4.3)$$

This efficiency only depends on the internal coupling of the optomechanical system, and for both  $C_1 = C_2$  and  $C_1, C_2 \gg 1$ , approaches unity. The latter condition can be understood by requiring the coupling between the optical and mechanical modes to overtake the intrinsic mechanical loss rate, while the first requirement is due to impedance-like matching [31]. The total system efficiency is  $\eta_{\text{max}} = \eta_1\eta_2\eta_{\text{max,int}}$ .

## 4.2 Experimental Setup and Device Characterization

### 4.2.1 Experimental Setup

The full experimental setup required to perform the wavelength conversion and necessary characterization is presented in Figure 4.2. All the electronic equipment including the two Newport tunable external cavity diode lasers (ECDL) is controlled via MATLAB [77] allowing the control of the equipment to be highly automated (setting wavelengths and frequencies, adjusting attenuations, etc.), as well as for data acquisition (recording optical powers, voltages from optical transmission scans, data from electronic spectrum analyzers, etc.). The amount of equipment used for this experiment and the complexity of its interaction necessitate a control system that allows for very repeatable data acquisition, coupled with the modularity required to quickly adjust control and acquisition. The MATLAB framework provides this ability and has been used in several experiments with a growing experimental setup in terms of both equipment used and the demands on the interaction [30, 33, 78–80].

For the wavelength conversion process, each tunable laser is locked to a specific wavelength using a wavemeter ( $\lambda$ -meter) continuously monitoring the lasers' wavelength and feeding back on the piezo control of the laser. The piezo tunability of the lasers is on the order of 300 pm ( $\pm 3$  V), providing a large tuning range and, when combined with the relative sensitivity of the wavemeter ( $\pm 5$  MHz), allows for locking to the desired wavelength within better than 0.1 pm. Each laser is further intensity-stabilized by first sending the light through fiber polarization controllers (FPC) before entering amplitude modulators (a-m)—the modulators are polarization-sensitive and require the correct input polarization to optimize the output power. After the amplitude modulators 10% of the light is picked off via a 90-10 splitter and sent to a lock box feeding back on the modulator. The

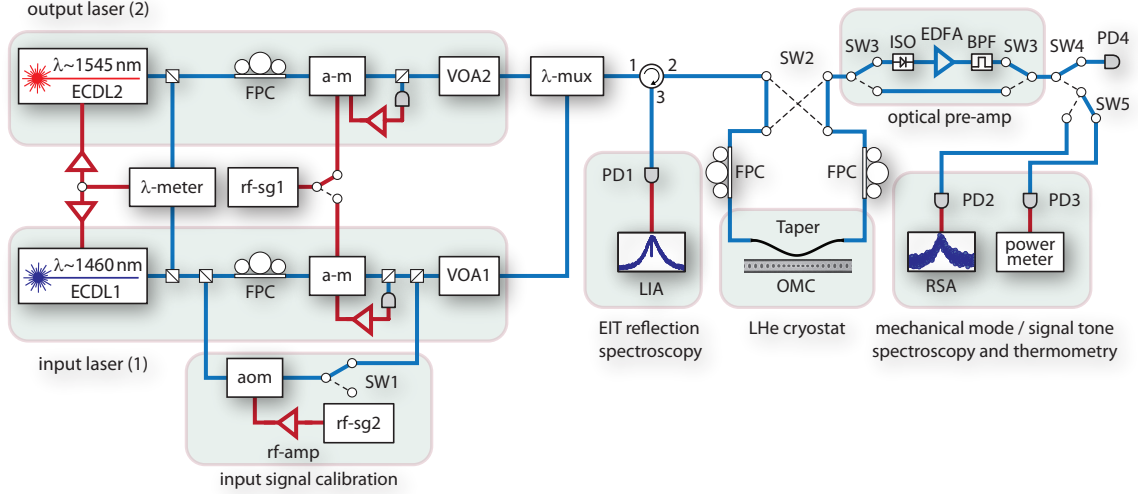


Figure 4.2: **Schematic of the experimental setup.** Two New Focus tunable, external cavity diode lasers (ECDL; ECDL1, ECDL2) are used as intense pump beams locked a mechanical frequency red-detuned from both the first and second order optical mode. The wavelength of both lasers is monitored by a wavelength meter ( $\lambda$ -meter) and locked in a feedback loop. Both beams can be amplitude modulated (a-m) using an RF signal generator (rf-sg1) to perform EIT spectroscopy on reflection using a lock-in amplifier (lock-in), and generate the input signal for the wavelength conversion process. An acousto-optic modulator (aom) is used to calibrate the input signal (switching in the calibration signal with SW1). The light from both lasers can be attenuated (VOA1, VOA2) to the desired power and is then combined using a wavelength multiplexer ( $\lambda$ -mux) and sent into a dimpled optical fibre taper coupled to the optomechanical crystal located in a continuous-flow liquid helium cryostat. Before entering the cryostat, the light passes through fibre polarization controllers (FPC) to align the polarization of the light with the device. The light also passes through an optical switch (SW2), allowing the light to propagate in both directions in the taper to calibrate the losses in the taper before and after the device. The transmitted light from the cavity can be optically pre-amplified (SW3) using an erbium doped fiber amplifier (EDFA) before being sent to a high-speed photodetector (PD2) and then a real-time spectrum analyzer (RSA). Optical switches (SW4, SW5) can also send the light to an optical power meter (PD3, power meter) for power calibration and a photodetector for optical transmission scans (PD4).

modulators have a sinusoidal transfer function and by locking in the middle (linear region) of this transfer function, the output power can be stabilized and small RF modulations about this lock point act to amplitude modulate the optical signal creating tunable sidebands; something required for both EIT spectroscopy and generation of input tones for the wavelength conversion process. After the amplitude modulators, the optical signals pass through variable optical attenuators (VOA), which are used to set the optical power of each signal in the system. It is much better to set the laser power (or laser current), let it stabilize, and never change it while using a VOA to adjust the power in the system. An extra step required for calibration of the input tone for the wavelength conversion process requires 10% of the light (again using a 90:10 splitter) from the input laser (ECDL1) to be sent to an acousto-optic modulator (aom). The aom is driven at very high RF power (+32 dBm; rf-sg2 and rf-amp) at 100 MHz, shifting nearly 100% of the input signal +100 MHz optically

(unshifted light is blocked allowing the aom to be used as an efficient on/off switch). A switch (SW1) determines whether this light is combined (using a 90:10 splitter again, 10% from the aom) for calibration or simply discarded. The aom being fed by the amplified RF source at this power take time to stabilize, so it is better to switch the light in and out rather than switch the RF power on and off. The two optical trains are then combined in a wavelength multiplexer ( $\lambda$ -mux) and sent through an optical circulator ( $1 \rightarrow 2$ ) to the fiber taper in the cryostat. The cryostat is a continuous flow  $^4\text{He}$  cryostat [17] with a cold finger temperature of 9 K for this experiment. Just before the cryostat, the signal is sent through a switch (SW2) and then both arms are sent through FPCs. This allows the direction of the light going into the taper to be switched and for both directions to have their polarizations adjusted such that the light has the same polarization as the optomechanical crystal (OMC) nanobeam. The coupling of the fiber taper to the OMC is touch dependent and can experience more or less optical loss before or after interacting with the OMC. To characterize this loss we send the same amount of power into the device from both directions (left and right determined by SW2) and monitor the effects, thus allowing us to calibrate the losses before and after the taper and determine the power at the device accurately.

The light leaving the device on transmission is sent through a switch (SW3), allowing the signal to be optionally amplified. The light is amplified by an erbium doped fiber amplifier (EDFA) which also creates a lot of excess noise, so the output light is then sent through a bandpass filter (BPF) with a 20 nm optical transmission window (1530 nm - 1550 nm). The EDFA also emits noise into the input channel. Therefore, to avoid having excess noise impinge on the device, an optical isolator (ISO) is placed before the input to the EDFA. After the optional amplification stage a switch (SW4) in the optical train can send the light to a nanosecond photodetector (PD4) for optical transmission scans. This switch can also send the light to the final switch in the optical train (SW5), which sends the light to a power meter (PD3) for calibration of the optical power or to a high-speed photodetector (PD2). The electrical signal output from PD2 is then sent to a real time spectrum analyzer (RSA) for Fourier analysis of the signal. The light reflected from the device passes through the circulator ( $2 \rightarrow 3$ ) and onto a photodetector (PD1). The electrical signal from PD1 is sent to a lock-in amplifier and used for the EIT reflection spectroscopy.

#### 4.2.2 Device Characterization

The optomechanical system used in this work, shown in Figure 4.3, consists not of a separate set of optical cavities, but rather of a single OMC nanobeam cavity. The OMC nanobeam is fabricated from a silicon-on-insulator (SOI) microchip [39, 45], in which the top 220 nm thick silicon device layer is patterned with a quasi-periodic linear array of etched air holes. The resulting silicon beam is 10  $\mu\text{m}$  long and 600 nm wide. The hole radii and positions are designed by numerical optimization and repeated simulations, as discussed in detail in Ref. [39]. The larger air holes




Figure 4.3: **Optomechanical crystal nanobeam.** (a) Scanning electron micrograph (SEM) of the fabricated silicon nanobeam optomechanical cavity. (b, c) Finite element method (FEM) simulation of the electromagnetic energy density, normalized to the maximum modal energy, of the first (b) and (c) second order optical cavity modes of the silicon nanobeam. (d) FEM simulation of the displacement field of the co-localized mechanical mode, normalized to the maximum modal displacement.

on either end of the beam induce Bragg-like reflection of both guided optical and acoustic waves, resulting in strongly confined optical and mechanical resonances at the beam's center. The device used in this work is designed to have both a first-order ( $\hat{a}_1$ ,  $\lambda \approx 1460$  nm) and a second-order ( $\hat{a}_2$ ,  $\lambda \approx 1545$  nm) optical cavity resonance of high quality factor. Both of these optical cavity modes are dispersively coupled to the same gigahertz-frequency mechanical resonance, depicted in Figure 4.3d ( $\hat{b}$ ,  $\omega_m/2\pi = 3.993$  GHz,  $Q_m = 87 \times 10^3$  at  $T \approx 14$  K). A tapered optical fiber waveguide, placed in close proximity ( $\sim 100$  nm) to the OMC nanobeam cavity [81] using precision stages [82], is used to evanescently couple light into and out of the cavity modes.

A schematic of the experimental set-up used to characterize the OMC wavelength converter is shown in Figure 4.2. Measurements are performed in vacuum and at low temperature inside a continuous flow cryostat with a cold finger temperature of 9 K (the corresponding sample temperature, as inferred from the thermal bath temperature of the localized mechanical mode of the OMC cavity, is 14 K [30]). Initial characterization of the optical cavity modes is performed by scanning the tunable control lasers across a wide bandwidth and recording the transmitted optical intensity through the optical fiber taper waveguide. Such a wavelength scan is shown in Figure 4.4a, from which the resonance frequency ( $\omega_1/2\pi = 205.3$  THz,  $\omega_2/2\pi = 194.1$  THz), cavity linewidth ( $\kappa_1/2\pi = 520$  MHz,  $\kappa_2 = 1.73$  GHz), and waveguide coupling efficiency ( $\eta_1 = 0.10$ ,  $\eta_2 = 0.21$ ) of the two nanobeam cavity modes are determined.

Further characterization of the optomechanical cavity is performed by using the control laser beams in conjunction with a weak sideband probe. With control beams  $\alpha_1$  and  $\alpha_2$  detuned, a mechanical frequency to the red of their respective cavity modes ( $\delta_{1,2} = \omega_m$ ), a weak sideband signal generated from  $\alpha_1$  is swept across the first-order cavity mode. The resulting reflected sideband signal



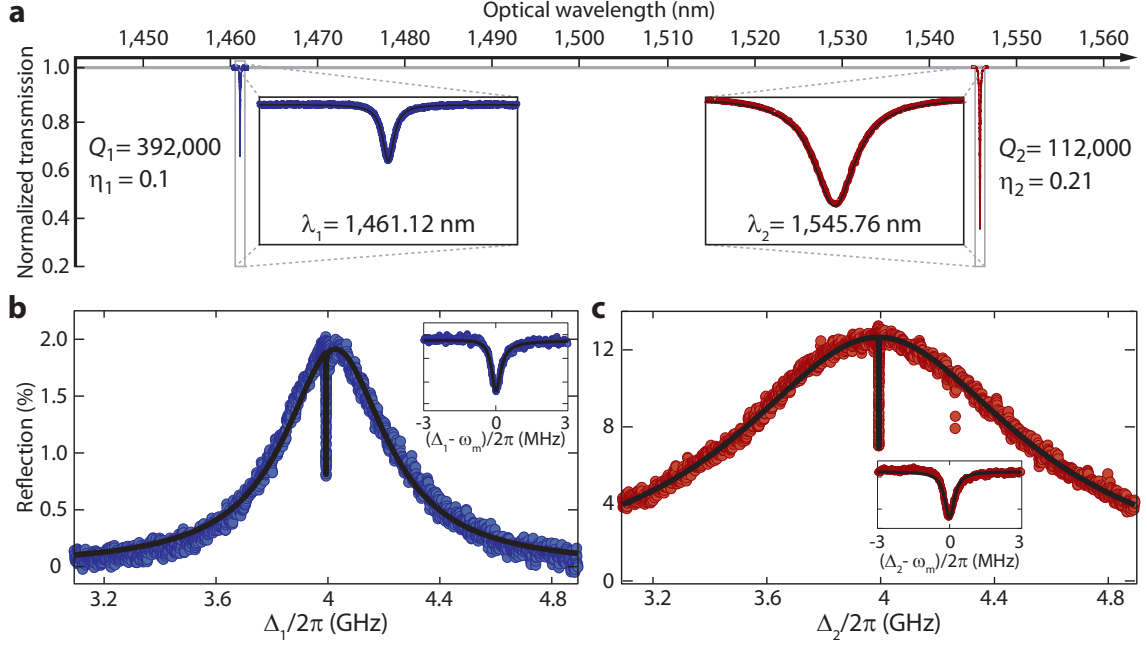


Figure 4.4: **Optical spectroscopy of the device.** (a) Broad wavelength transmission scan showing fiber taper coupling to both the first-order ( $\lambda \approx 1460$  nm) and second-order ( $\lambda \approx 1545$  nm) cavity modes. (b) EIT scan of cavity mode  $\hat{a}_1$ , with the inset showing a zoomed-in scan of the transparency window. The solid black lines correspond to fits to a theoretical model allowing extraction of system parameters ( $\kappa_1$ ,  $\gamma$ ,  $\delta_1$ , and  $\omega_m$ ). (c) Corresponding EIT scan of cavity mode  $\hat{a}_2$ . For measurements in both (b) and (c) the control beam intensities were set such that  $\gamma_{\text{OM},1} \approx \gamma_{\text{OM},2} \gg \gamma_i$ , with detunings  $\delta_1 \gtrsim \omega_m$  and  $\delta_2 \approx \omega_m$ .

versus sideband frequency shift  $\Delta_1$  is plotted in Figure 4.4b, showing the broad cavity resonance along with a narrow central reflection dip (see Figure 4.4b inset). The narrow reflection dip, akin to the EIT transparency window in atomic systems, is due to the interference between light coupled directly into the cavity mode and light coupled indirectly through the mechanical mode [25, 26]. A similar EIT response is shown in Figure 4.4c for the second-order cavity mode. Each of the transparency windows occur at  $\Delta_k = \omega_m$ , with a bandwidth equal to the optically damped mechanical linewidth ( $\gamma = \gamma_i + \gamma_{\text{OM},1} + \gamma_{\text{OM},2}$ ). By fitting the optical resonance lineshape and transparency windows to theory, one can also extract the control beam detunings  $\delta_k$ . In what follows, we use this sort of EIT reflection spectroscopy, time-multiplexed in between wavelength conversion measurements, to set and stabilize the frequency of the control beams to  $\delta_k = \omega_m$ . We note the EIT reflection spectroscopy, though convenient for full characterization and elucidation of the conversion process, is not necessarily performed in a real system application given that both the laser and the cavity resonance have frequency drift, over the measured period of days, that is much smaller than the optical cavity linewidth  $\kappa$  of the devices studied here.

Figure 4.5 plots the inferred optomechanically induced damping of the mechanical resonator as




Figure 4.5: **Characterization of the optomechanical coupling.** Plot of the optically induced mechanical damping versus control beam intensity (intracavity photon number) for each cavity mode, as measured through the mechanical thermal noise spectrum imprinted on the optical output intensity of  $\hat{a}_2$ . The calibration curve for cavity mode  $\hat{a}_2$  is performed with  $\alpha_1$  turned off, whereas the curve for cavity mode  $\hat{a}_1$  is generated with a weak  $\alpha_2$  such that  $\gamma_{\text{OM},2} \ll \gamma_{\text{OM},1}$  for all measured points.

a function of the power (intracavity photon number,  $n_c$ ) of control beams  $\alpha_1$  and  $\alpha_2$ , the slope of which gives the zero-point optomechanical coupling for both cavity modes ( $g_1/2\pi = 960$  kHz and  $g_2/2\pi = 430$  kHz).

### 4.2.3 Generation and Conversion of the Input Signal

We begin by considering the input signal to the first optical cavity mode. This signal is generated by weakly modulating the  $\alpha_1$  control beam (represented by amplitude  $A_1$ ) at a frequency  $\Delta_1$ , using an electro-optic modulator (EOM) producing

$$A_1 \rightarrow A_1 + \frac{\beta_1}{4} A_1 e^{-i\Delta_1 t} + \frac{\beta_1}{4} A_1 e^{i\Delta_1 t} e^{i\phi}, \quad (4.4)$$

where  $\beta_1$  is the modulation index of the input signal and  $\phi$  accounts for any phase difference between the sidebands ( $\phi = 0$  if there is pure amplitude modulation). The input sideband at frequency  $\omega_{1,1} + \Delta_1$  is nearly resonant with the first cavity mode at  $\omega_1$ , whereas the lower frequency sideband at  $\omega_{1,1} - \Delta_1$  is detuned by approximately  $-2\omega_m$ . As such, only the upper frequency sideband is resonant with the transparency window of the first cavity mode, and only this sideband is converted by the optomechanical crystal cavity into a sideband at the second, output cavity mode  $\hat{a}_2$ . The converted sideband is generated from the control laser beam  $\alpha_2$  of the second cavity mode with frequency  $\omega_{1,2} + \Delta_1$ . Thus, the output field amplitude near the frequency of the second control laser beam (amplitude  $A_2$ ) is, ignoring  $\phi$ ,

$$A_2 \rightarrow A_2 + s_{21}(\Delta_1) \frac{\beta_1}{4} e^{i\Delta_1 t}. \quad (4.5)$$

The power measured by a photodetector from the optical signal emanating from the output cavity mode is proportional to

$$P_2 = |A_2|^2 \left( 1 + |s_{21}(\Delta_1)| \frac{\beta_1 A_1}{4A_2} \cos \Delta_1 t \right). \quad (4.6)$$

We thus define  $\beta_2 = |s_{21}(\Delta_1)|\beta_1 A_1/4A_2$  as the modulation index of the output signal,

$$P_2 = |A_2|^2 (1 + \beta_2 \cos \Delta_1 t). \quad (4.7)$$

By careful calibration of  $\beta_1$  and  $\beta_2$ , and measurement of the control beam power ( $P_{\alpha_k}$ ), the conversion scattering matrix element can be determined:

$$s_{21}(\Delta_1) = 2 \frac{\beta_2 A_2}{\beta_1 A_1} = 2 \frac{\beta_2}{\beta_1} \sqrt{\frac{P_{\alpha_2}}{P_{\alpha_1}}}. \quad (4.8)$$

#### 4.2.4 Calibration of the Input Signal




Figure 4.6: **Calibration tones.** Schematic of the four tones appearing in the electronic power spectrum of the photocurrent generated by the optical input signal and the additional calibration signal. The three additional tones proportional to  $A_3$  are generated to calibrate the input signal modulation index,  $\beta_1$ .

The appearance of  $\phi$  in Equation (4.4), along with its sensitivity to effects such as chromatic dispersion in the optical fiber and components, make calibration of  $\beta_1$  more challenging due to interference between the upper and lower frequency sidebands. Direct photodetection of the optical input signal results in a  $\phi$ -dependent measured beat signal between the carrier and the two sidebands. To bypass this problem, we generate an additional single sideband using an acousto-optic modulator (AOM), and beat it against each of the input optical sidebands. This is accomplished by splitting off a portion of the optical input signal prior to creating the EOM sidebands, and frequency shifting it by  $\Delta_{AOM}$  (see Figure 4.6). This AOM-shifted signal is then recombined with the main input signal, giving an overall signal equal to

$$A_{out,1} = A_1 + \frac{\beta_1}{4} A_1 e^{-i\Delta_1 t} + \frac{\beta_1}{4} A_1 e^{i(\Delta_1 t + \phi)} + A_3 e^{i\Delta_{AOM} t}, \quad (4.9)$$

where  $A_3$  is the field amplitude of the signal split off from the input signal prior to the EOM. The total photodetected signal is then given by

$$\begin{aligned}
|A_{\text{out},1}|^2 &= |A_1|^2 + |A_3|^2 \\
&+ |A_1|^2 \beta_1 \cos(\Delta_1 t + \phi/2) \cos(\phi/2) \\
&+ 2|A_1 A_3| \cos(\Delta_{\text{AOM}} t) \\
&+ |A_1 A_3| \frac{\beta_1}{2} \cos[(\Delta_1 - \Delta_{\text{AOM}})t] \\
&+ |A_1 A_3| \frac{\beta_1}{2} \cos[(\Delta_1 + \Delta_{\text{AOM}})t + \phi].
\end{aligned} \tag{4.10}$$

From this equation we see that the optical power at the detector consists of components at zero frequency (DC) as well as four modulated tones (see Figure 4.6). By taking the ratio of the tone at  $\Delta_{\text{AOM}}$  and  $\Delta_1 - \Delta_{\text{AOM}}$ , the value of  $\beta_1$  can be accurately determined independent of  $\phi$ . Note that in order to determine  $\beta_2$ , no such additional calibration signal is required. The filtering properties of the input and output cavities results in a single output sideband, with no phase dependence of the photodetected intensity spectrum. In this case, careful optical and electronic calibration of the measured output cavity ( $\hat{a}_2$ ) optical transmission intensity, similar to that described in the Supplementary Information of ref. [30], provides an accurate estimate of  $\beta_2$  and the wavelength converted signal strength.

#### 4.2.5 Power-Dependent Optical Cavity Loss

To model the wavelength conversion process, all of the optical cavity mode and mechanical resonator parameters entering into Eq. 2.25 must be independently determined. Most of these parameters are measured as described in the main text; however, additional measurements were performed to determine the power dependent optical cavity loss in the silicon optomechanical crystal device. Due to two-photon absorption in silicon, the parasitic optical cavity loss of both cavity modes is not static, but rather depends upon both control beam powers. This effect is particularly acute for the higher  $Q$ -factor first-order cavity mode ( $\hat{a}_1$ ), where in Figure 4.7a we plot the normalized on-resonant optical transmission ( $T_{0,1}$ ) of the first-order cavity mode versus the power (as represented by the intra-cavity photon population,  $n_{c,2}$ ) of control beam  $\alpha_2$  feeding the second-order cavity mode. The steady rise in the on-resonance transmission versus  $n_{c,2}$  is attributable to the increased optical absorption loss stemming from free-carriers generated by two-photon absorption of the control beam  $\alpha_2$ . The corresponding change in coupling efficiency,  $\eta_1 \equiv \kappa_{e,1}/2\kappa_1$ , is shown on the right-side axis of Figure 4.7a. Neglecting the effects of nonlinear optical absorption and free carriers on the intrinsic optical quality factor of  $\hat{a}_1$  leads to a theoretical model of conversion efficiency denoted by the dashed black line in Figure 4.7b. Taking into the account the nonlinear optical loss measured in




Figure 4.7: **Effect of large intracavity photon population.** (a) Recorded coupling depth of the first order optical cavity as a function of the intracavity photon population of the second order mode (green circles). The black curve represents a fit to the extracted coupling depth. (b) Conversion efficiency as a function of  $C_2/C_1$ . The dashed black line represents the theoretical values from the nominal system parameters while the solid red line takes into account the effect of the large intracavity photon population of cavity mode 2 ( $n_{c,2}$ ). The blue circles are the extracted data points and show much better correspondence to theory when this effect is accounted for.

Figure 4.7a leads to the theoretical conversion efficiency denoted by the solid red line, showing much better correspondence with the data (the solid red line is the model fit shown in the main text). The effect of two-photon absorption and free-carrier absorption is much smaller on the second-order cavity mode ( $\hat{a}_2$ ) due to its already lower optical  $Q$ -factor, and we neglect it here. The control-beam-generated free carriers also affect the intrinsic mechanical quality factor ( $\gamma_i$ ), as was shown in ref. [30]. However, this effect is small for the control beam powers studied here (the efficiency curve of Figure 4.7b is also not modified as we plot it versus the ratio of cooperativities for which  $\gamma_i$  cancels out).

### 4.3 Experimental Results

Coherent wavelength conversion can be thought of as occurring between the input transparency window of  $\hat{a}_1$  and the output transparency window of  $\hat{a}_2$ , with the phonon transition mediating the conversion. As shown in Section 2.3.1, the conversion efficiency is theoretically optimized for matched cooperativities of control beam  $\alpha_1$  and  $\alpha_2$  (or equivalently  $\gamma_{\text{OM},1} = \gamma_{\text{OM},2}$ ). To characterize the wavelength conversion process we will want to understand its efficiencies, bandwidth, and noise properties as well as the associated limiting factors. We begin by looking at the efficiency.

### 4.3.1 Conversion Efficiency

To quantify the efficiency of the wavelength conversion process, the input cavity control beam  $\alpha_1$  is held fixed at a detuning  $\delta_1 = \omega_m$  and a power producing an intracavity photon population of  $n_{c,1} = 100$ , corresponding to large cooperativity  $C_1 \approx 16$ . The input signal,  $a_{\text{in}}$ , is generated as an upper sideband of  $\alpha_1$  using an electro-optic intensity modulator (the lower sideband is rejected by the wavelength converter as it is detuned to  $\Delta_1 \approx -\omega_m$ ). Conversion of the  $a_{\text{in}}$  sideband to an output signal emanating from cavity mode  $\hat{a}_2$  is completed by applying control beam  $\alpha_2$  with detuning  $\delta_2 = \omega_m$ . The converted output signal sideband emitted from cavity  $\hat{a}_2$  is beat against the transmitted control beam  $\alpha_2$  on a high-speed photodiode. The amplitude of the input signal tone near resonance with cavity mode  $\hat{a}_1$  is calibrated using a second reference signal generated through acousto-optic modulation of  $\alpha_1$ , whereas the amplitude of the output signal sideband is inferred from calibration of the control beam intracavity number  $n_{c,2}$  and the optical transmission and detection chain (see Figure 4.2). Figure 4.8 plots the resulting power conversion efficiency, given by the magnitude squared of Equation 4.1, versus the ratio of  $C_2/C_1$  as the power of control beam  $\alpha_2$  is varied and for an input signal exactly resonant with cavity mode  $\hat{a}_1$  ( $\Delta_1 = \omega_m$ ). The




Figure 4.8: **Wavelength conversion efficiency.** End-to-end power conversion efficiency of input signal to output signal for  $\Delta_1 = 0$  as a function of the ratio of the cooperativities of the control beams. In this plot the control beam intensity for the first-order cavity mode is held fixed (with  $C_1 \approx 16$ ), while the intensity of the control beam of the second-order cavity mode is swept from  $C_2 \ll C_1$  to  $C_2 > C_1$ . The blue circles correspond to measured data points, whereas the solid red line is a theoretical curve using independently measured system parameters.

power conversion is referred directly to the input and output of the optomechanical cavity (i.e., not including additional losses in the optical link). The solid red line shows the theoretical model for the conversion efficiency using the independently measured system parameters, and taking into account power-dependent effects on the optical losses in the silicon cavity modes (see Section 4.2.5). Good correspondence is seen between the measured data and theory, with a maximum of the conversion occurring at  $C_1 \approx C_2$  as expected from the matching condition.

### 4.3.2 Bandwidth

The bandwidth of the wavelength conversion process is also probed, using matched conditions ( $C_1 = C_2 \approx 16$ ), and, as above, with control beam detunings  $\delta_k = \omega_m$ . The modulation frequency generating  $a_{\text{in}}$  is now varied from  $(\Delta_1 - \omega_m)/2\pi = -10$  MHz to  $+10$  MHz, sweeping the narrow-band input sideband signal across the transparency window of cavity mode  $\hat{a}_1$ . The corresponding measured output signal frequency shift follows that of the input signal ( $\Delta_2 = \Delta_1$ ) as can be seen in Figure 4.9. At each detuning the resulting power conversion efficiency is determined and plotted in




Figure 4.9: **Output tones for input tone detunings.** Measured power spectrum of the output optical channel for  $C_1 = C_2 \approx 16$ , showing a series ( $\Delta_1$  sweeps) of converted input tones (green) sitting on top of a much smaller thermal noise pedestal (blue). Inset shows a zoom-in of the thermal noise pedestal.

Figure 4.10. A peak efficiency, measured from the input to the output port at the optomechanical cavity system, is  $\eta = 2.2\%$ , corresponding to an internal conversion efficiency of  $\eta_{\text{max,int}} = 93\%$ . Fitting the results to a Lorentzian as in Equation (2.25), we find the conversion bandwidth to be 1.55 MHz, equal to the optically damped mechanical linewidth of the mechanical resonance,  $\gamma/2\pi$ .




Figure 4.10: **Bandwidth of the wavelength conversion.** The conversion efficiency versus input signal detuning for matched control beams  $C_1 = C_2 \approx 16$  (corresponding to the conditions for the highlighted data point in Figure 4.8, indicating a conversion bandwidth of  $\sim 1.55$  MHz). The blue circles correspond to the measured conversion efficiency taken at each detuning, while the red curve is the theoretical model.




Figure 4.11: **Converted tone on thermal noise.** Converted tone sitting on top of the mechanical mode driven by thermal fluctuations. The inset shows a zoomed-in mechanical mode with a phonon population of  $\langle n \rangle \approx 3$ . This is the cooled phonon population resulting from the large coherent drive tones a mechanical frequency red detuned from the two optical modes giving rise to cooling.

### 4.3.3 Noise Considerations

The spectrum of the converted output signal is shown in Figure 4.11. As can be seen from this plot and the zoomed-in inset, the narrow-band converted photons sit atop a noise pedestal of bandwidth corresponding to the damped mechanical resonator. This noise is quantified by the added noise quanta number component of Equation 4.2. As shown in Section 2.3.1, under matched conditions the theoretical added noise is,

$$n_{\text{added}} = 2\eta_1^{-1} \left( \frac{\gamma_i n_b}{\gamma} + \frac{1}{2} \left( \frac{\kappa_1}{4\omega_m} \right)^2 + \frac{1}{2} \left( \frac{\kappa_2}{4\omega_m} \right)^2 \right), \quad (4.11)$$

where the first term arises due to thermal noise of the cooled and damped mechanical resonator, and the last two terms are quantum noise resulting from the spontaneous scattering of the control beams (quantum back-action noise [34]). These sources of noise are independent and uncorrelated, and as such, add together incoherently. Figure 4.12a pictorially indicates the various cooling, heating, and spontaneous scattering mechanisms that lead to the output noise of  $\hat{a}_2$ . A plot showing the measured output noise spectral density, calibrated in units of photon number, is given in Figure 4.12b for the optimal conversion efficiency of Figure 4.8. From the peak of the output noise spectral density ( $= \eta_1 \eta_2 n_{\text{added}}$ ) in this plot, the added noise referred to the input of  $\hat{a}_1$  is estimated to be  $n_{\text{added}} \approx 60$  quanta. The corresponding internal added noise (for  $\eta = 1$ ) is only  $n_{\text{added}} \approx 6$ , predominantly due to the cooled phonon occupation of the mechanical resonator ( $\langle n \rangle \approx 3$ ). Due to the strong filtering provided by the high- $Q$  cavities ( $\kappa_k/\omega_m \ll 1$ ), the quantum back-action noise contributes an insignificant  $4 \times 10^{-2}$  ( $4 \times 10^{-3}$ ) added input (internal) noise photons. The effectiveness of noise reduction in the optomechanical cavity system is highlighted in Figure 4.12b by showing the estimated thermal output noise in the absence of sideband cooling by the control beams (dashed red



curve) and the quantum back-action noise in absence of cavity filtering (green dashed curve).

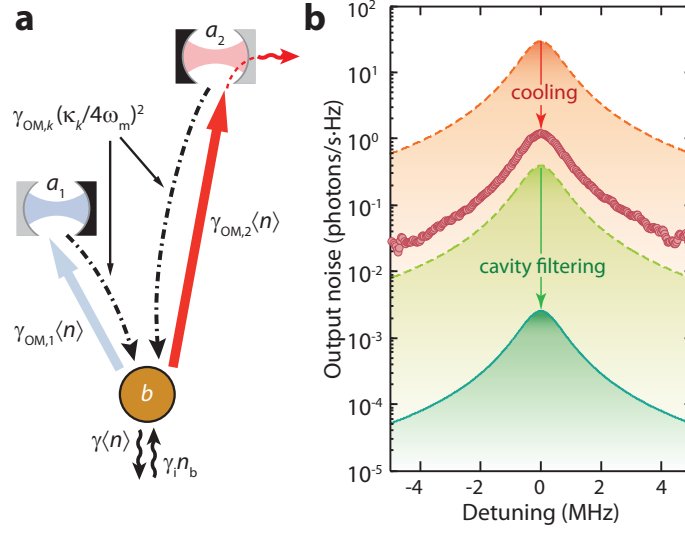


Figure 4.12: **Output noise.** (a) Schematic showing the relevant input noise terms contributing to the total noise at the output. The rates of “noisy” phonon generation are  $\gamma_i n_b$  due to the thermal bath and  $\gamma_{OM,(1,2)}(\kappa_{(1,2)}/4\omega_m)^2$  due to spontaneous Stokes scattering from the two control beams. Cooling of the mechanical mode is also performed by the control beams, resulting in a theoretically cooled phonon occupancy of  $\langle n \rangle = (1/\gamma)(\gamma_i n_b + \sum_j \gamma_{OM,j}(\kappa_j/4\omega_m)^2)$ . The output noise power is proportional to the output cavity cooling rate,  $\gamma_{OM,2}\langle n \rangle$ . (b) Output noise power spectral density. The measured noise data (red circles) is shown along with several theoretical noise curves. The dashed red curve is the corresponding output noise power in the absence of control beam cooling. The solid (dashed) green curve is the output noise power due to spontaneous Stokes scattering in the presence (absence) of cavity filtering.

## 4.4 Towards Single-Photon Wavelength Conversion

Although the demonstrated wavelength conversion in this work represents an important proof of principle, there are still considerable practical challenges to utilizing such cavity optomechanical devices for classical or quantum wavelength translation [58, 65]. With a thermal phonon occupancy  $\langle n \rangle \leq \eta_1/2$  required for single photon conversion at unity signal-to-noise, a primary concern for quantum networking [57] is the thermal energy stored in the mechanical resonator. Although resolved sideband cooling of nanomechanical resonators to occupancies below one has recently been demonstrated [30, 70], this relation emphasizes the additional need for efficient cavity coupling. For chip-scale photonic devices similar to the one presented here, the technology already exists for efficient optical connectivity, with optical fiber-to-chip coupling of  $\sim 95\%$  [83] and waveguide-to-cavity coupling efficiency of  $\eta \gtrsim 0.999$  having been realized [84]. Further reduction in the output noise may also be realized by simply reducing the temperature of operation: for the 4 GHz phonon frequency of the device in this study,  $n_b \lesssim 10^{-3}$  for a bath temperature of  $T \sim 10$  mK realizable in a helium

dilution refrigerator. Integration of cavity optomechanical devices into millikelvin experiments is particularly relevant for superconducting quantum circuits, which have already been strongly coupled to mechanical resonators [85]. Addition of an optical cavity coupled directly, or via an additional acoustic waveguide, to the same mechanical resonator would provide a microwave-to-optical quantum interface [31, 72] similar to the optical-to-optical interface shown here. In the case of solid-state qubits with optical transitions far from the telecommunication bands, such as the nitrogen-vacancy (NV) electron spin in diamond [86], the radiation pressure nonlinearity of cavity optomechanics could be substituted for intrinsic materials nonlinearities in proposed cavity-based schemes for single photon wavelength conversion and pulse shaping [87]. Due to the excellent mechanical properties of diamond and recent advancements in thin-film diamond photonics [88], a cavity-optomechanical wavelength converter could be fabricated around an NV qubit from the same diamond host.

## Chapter 5

# Adiabatically Tapered Waveguide Coupling

### 5.1 Motivation

In the experiment on wavelength conversion using the optomechanical nanobeam, the internal efficiency was demonstrated to be 93%, however, the end-to-end efficiency including the coupling of the light into the device and out of the device was limited to  $\sim 2\%$ . For this experiment the device was bidirectionally coupled to the fiber taper, providing coupling efficiencies of  $\eta_1 = 0.1$  and  $\eta_2 = 0.21$  to the optical modes at 1460 nm and 1545 nm respectively (see Figure 4.4). This bidirectional coupling means half of the photons leaving the cavity couple into the optical fiber in the transmitted direction, while the other half couple into the fiber in the reflected direction limiting our collection efficiency to 50% (for more information see Appendix B.2.1). A similar problem occurs when we try to couple the photons into the cavity with a maximum of 50% coupling in. This problem clearly represents an opportunity to greatly improve the end-to-end wavelength conversion efficiency without compromising any other aspect of the device or measurement.

In the current method of coupling to a device with a fiber taper, we need to bring the taper very close to the device to try and achieve a high coupling efficiency. This poses several problems, the first being that if we touch the device we can ruin the mechanical mode's properties (i.e., it has a large piece of glass now sitting on it). Another problem is that this coupling method is somewhat stochastic. To couple to the device we use precision stages to bring the fiber into the near-field of the device while monitoring the transmission of light through the fiber. When the taper gets close to the device, it loads it—that is, the device now sees the taper as a new loss channel ( $\kappa_e$ ) so that the total loss channel is  $\kappa = \kappa_i + \kappa_e$  and the coupling efficiency is defined as  $\eta = \kappa_e/2\kappa$  (where the 50% maximum efficiency limit is clearly visible for perfect coupling:  $\kappa \rightarrow \kappa_e$ ). From the transmission scan of the optical mode we can quantify the coupling efficiency and position the taper to achieve the required coupling. It can be quite difficult to properly position the taper to get the exact coupling

required as the taper can jump and stick to the surface when it is close. The taper itself can also cause extra loss in the device, which can reduce the coupling efficiency (i.e.,  $\kappa$  increases and the ratio of  $\kappa_e/2\kappa$  decreases).

We want to develop a unidirectional (single-sided) coupling method that can bypass the limitations of our current bidirectional coupling that limits our input and output coupling each to 50%. A single-sided coupling method would allow a theoretical coupling efficiency of 100%, meaning, for the wavelength conversion, we would be able to theoretically get 100% of the photons we wish to convert into the device, and then collect 100% of the converted photons from the device. This would have given a 93% end-to-end conversion efficiency in the experiment, everything else staying the same. Previously, the theoretical maximum conversion efficiency was limited to 25% because of the double-sided coupling. We would like to achieve this single-sided coupling while still using the optical fiber taper as it allows us to quickly characterize many devices on a chip, allowing the best device to be chosen to work with. The coupling to the device should also be more engineerable, repeatable, and be such that we know we are not touching the mechanical mode and possibly ruining the mechanical properties of the device (it can be frustrating to not know if the touch is bad and ruining the mechanics or if the device is simply bad). It is with these goals in mind that we set out to develop a better coupling scheme.

The solution we came up with divided the problem into two parts. First, a method of efficiently coupling the light from the optical fiber into an on-chip waveguide by adiabatically tapering the waveguide was devised. This allowed the taper-waveguide coupling to be physically separated from the device, ensuring the mechanical properties were not affected. It also provide a way to couple all the photons into and out of the on-chip waveguide. With the light now in the on-chip waveguide, the second problem was engineering a single-sided coupling of the waveguide to the nanobeam photonic crystal. In this scheme, the coupling of the fiber to the waveguide, and the waveguide to the nanobeam, are independent and can be optimized and engineered separately. The details and results of this technique will be described in the following sections.

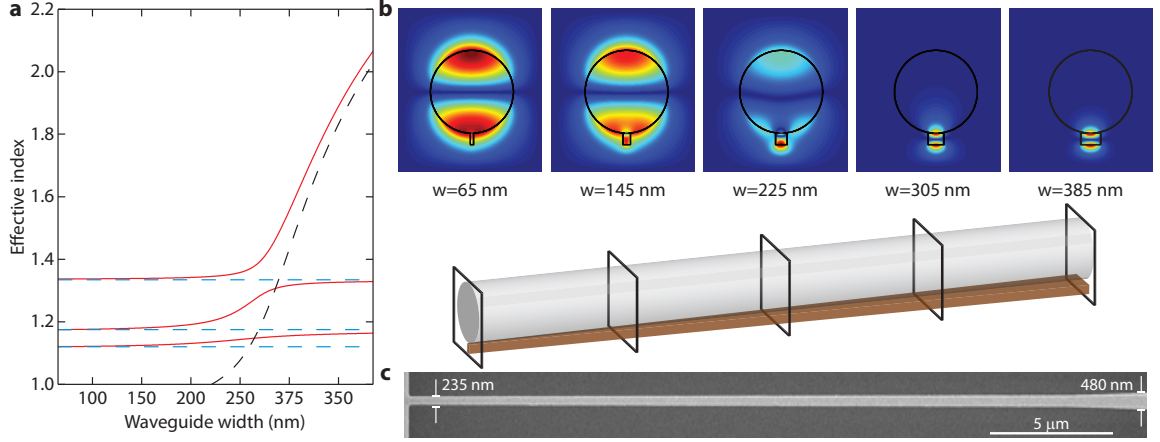
## 5.2 Adiabatically Tapered Waveguide-Fiber Coupling

The problem of efficiently getting light to on-chip devices and waveguides is a fundamental problem for integrated optics [89]. A fundamental issue is the difference in the mode profile of on-chip waveguides and that of an optical fiber. Many techniques involving tapered waveguides [54, 90, 91] and even tapered fibers [92] have been used to inter-convert between the differing spot sizes. Most of these techniques are well suited to the permanent placement of the fiber on-chip but do not provide an easy means to quickly test and switch between multiple devices. Another technique involves phased-matched coupling between an optical fiber taper and on-chip waveguide [93–95], preserving

the benefits of optical fiber testing but imposing bandwidth restrictions due to the phase matching condition required for efficient power transfer. The solution we propose and demonstrate involves aspects of both techniques, utilizing adiabatically tapered on-chip waveguides coupled to optical fiber tapers with a large transfer bandwidth and allowing for the testing of multiple devices.

The concept of adiabatic transfer of energy between coupled modes dates back to the work of Landau and Zener [96] on molecular energy level crossings. These ideas were proposed for microwave systems in a series of papers from Bell Labs in 1955 [97–99]. Work on optical devices in the 1970s [100, 101] with an optical switch using the principle of adiabatic eigenmode conversion being demonstrated a decade later [102]. The technique was even proposed and demonstrated to allow for coupling between off-chip optical fibers and on-chip waveguides [103]. An early demonstration of several different types of adiabatic couplers and a short review of examples of these types of couplers at that time is given in ref. [104]. The basic idea involves coupling two systems, each able to be described independently by normal modes, but when they are coupled it is more accurate to speak of the coupled modes of the system. In the region these system interact, if one of the system's parameters is varied slow enough, such that the energy of its normal mode adiabatically changes from lower to higher (or vice versa) than that of the other normal mode, then an excitation initially in the first system can be perfectly converted into an excitation of the second system. Here adiabatic means that the parameter is varied slowly enough such that the change in the energy level of that mode is small compared to the other energy levels in the system.

Specifically, our two systems are the bare optical fiber taper which has been thinned to a diameter of  $1.6\ \mu\text{m}$  such that it has an evanescent tail outside of the fiber [81]. The other system is an on-chip optical waveguide made in the 220 nm thick device layer of a silicon-on-insulator chip. When these two system are brought into close proximity they can be made to interact strongly, leading to a coupled mode picture of the fiber and waveguide system with even and odd symmetry supermodes. We can change the effective index of the waveguide mode by modifying its width, with a smaller width resulting in a smaller effective index. By adiabatically sweeping the effective index of the waveguide from below to above the index of the optical fiber ( $n = 1.33$  for the fiber mode of interest), all the power originally in the waveguide can be transferred to the fiber, and vice versa. A plot of the simulation showing the effective indices of the normal modes of the uncoupled system (dashed black for the waveguide and dashed blue for the optical fiber) as well as the coupled mode (solid red) is shown in Figure 5.1a. This system is schematically represented in the bottom of Figure 5.1b where the optical fiber is placed directly on top of the adiabatically tapered waveguide, allowing it to interact. The waveguide width changes from  $w = 65\ \text{nm}$  to  $w = 385\ \text{nm}$ . Simulations of the profile of the coupled mode is shown in the top of Figure 5.1b corresponding to five different slices of the system with different waveguide widths. It can be seen that for the smallest width waveguide,  $w = 65\ \text{nm}$ , the coupled mode is localized entirely in the optical fiber. By the time the width of the



**Figure 5.1: Adiabatically Tapered Waveguide-Optical Fiber Coupling.** (a) Simulation of the effective index of the bare optical fiber (dashed blue lines) and bare adiabatically tapered waveguide (dashed black line) along with the coupled modes' (solid red lines) effective index as a function of the waveguide width. The effective index of the bare waveguide increases with increasing width as expected. The bare optical fiber modes' effective indices have no dependence on the width of the waveguide. The fundamental mode's index is  $n=1.33$  and the higher order optical fiber modes have lower indices as they are less confined in the fiber. The waveguide material is silicon, 220 nm thick, and the optical fiber has a diameter of  $1.6\ \mu\text{m}$ . (b) A model (bottom) of the optical fiber sitting on the tapered waveguide with sections showing the mode profile (top) of the coupled system for various waveguide thicknesses. (c) A SEM image of the waveguide made in the 220 nm thick device layer of an SOI chip. The waveguide's width changes from 235 nm to 480 nm over a distance  $25\ \mu\text{m}$ .

waveguide is  $w = 385$  nm the mode is localized entirely in the on-chip waveguide. As long as the effective index of the waveguide is swept across the effective index of the optical fiber slow enough, the process is adiabatic and can result in near perfect conversion. A fabricated waveguide meeting this condition with its width varying from  $w = 235$  nm to  $w = 480$  nm over a distance of  $25\ \mu\text{m}$  is shown in Figure 5.1c.

### 5.3 Single-Sided Coupling of the Waveguide to the Device

After efficiently coupling light from the optical fiber to an on-chip waveguide, we are now left with the task of coupling the light from the waveguide into the device. With the interaction of the waveguide and optical fiber complete, the waveguide is now tapered up to its final width more quickly, in this case 600 nm. Coupling of the waveguide with the device is accomplished by bringing the waveguide into the near-field of the device, allowing for evanescent coupling. The closer the device and waveguide the stronger the coupling  $\kappa_e$ , and also the heavier the device is loaded,  $\kappa = \kappa_i + \kappa_e$ . The simulated strength of the waveguide-device coupling are plotted in Figure 5.2a. The cavity-waveguide coupling  $\kappa_e$  (solid green line) is seen to increase roughly linearly with decreasing separation. The intrinsic loss rate  $\kappa_i$  increases slightly due to the increased scattering resulting from a small gap

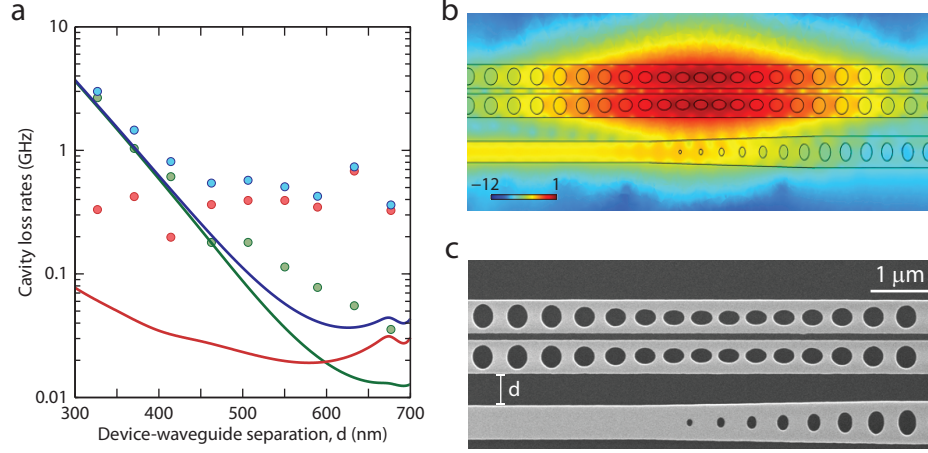


Figure 5.2: **Waveguide-Device Coupling.** (a) Simulation (solid lines) and experimentally extracted values (circles) for the intrinsic loss rate ( $\kappa_i$ ; red), cavity-waveguide coupling ( $\kappa_e$ ; green) and total cavity decay rate ( $\kappa = \kappa_i + \kappa_e$ ; blue). (b) Finite element methods (FEM) simulation of the optical mode, field intensity normalized to its maximum, coupled to the waveguide. The waveguide is designed with a photonic bandgap, at the optical frequencies of interest, on the right of the waveguide, which acts as a perfect mirror only allowing light in and out of the cavity from the left. The tapering of the hole size of the mirror prevents scattering. (c) Scanning electron micrograph (SEM) of the waveguide-device coupling region of a fabricated device. The waveguide-device separation,  $d$ , as plotted in (a) is highlighted.

and the total cavity loss rate  $\kappa$  is asymptotically approached by the cavity-waveguide coupling  $\kappa_e$  as  $\kappa_e \gg \kappa_i$ . To ensure a single-sided coupling between the device and waveguide, the waveguide is terminated by a photonic crystal mirror. The size of the holes of this mirror are adiabatically tapered to avoid scattering losses and the mirror region of the waveguide is the same as for the device. A simulation of the coupled waveguide and fundamental mode of the device is shown in Figure 5.2b. Here the device is a zipper-style double nanobeam used in the squeezing experiment [80]. A SEM of the fabricated device coupled to the waveguide is shown in Figure 5.2c with the waveguide-device spacing  $d$  indicated.

While the devices studied here for the characterization of this coupling scheme are the zipper-style double nanobeams, the method is agnostic to the device used and works equally well with the nanobeams, or other on-chip optical resonators. Figure 5.3 shows a SEM of the full structure with the waveguide coupled to the double nanobeams (Figure 5.3a) and to a single nanobeam (Figure 5.3b). The coupling of the optical fiber to the tapered waveguide is exactly the same for both structures, requiring no redesign. The coupling of the waveguide to the device is fundamentally the same with only a small difference in the coupling rate between the waveguide and device,  $\kappa_e$ , with respect to device-waveguide separation  $d$ .







Figure 5.3: **Waveguide Coupled Zipper Nanobeams and Single Nanobeam.** (a) Waveguide-coupled zipper-style double nanobeams showing both the adiabatically tapered waveguide (left side) and the single-sided waveguide-device coupling (right side) with the photonic crystal mirror terminating the waveguide. (b) Same coupling scheme as in a, this time showing the coupling to a single nanobeam. This coupling scheme can be used with a variety of different on-chip photonic crystal devices. In both a and b the single waveguide couples to two different devices above and below it in the image.

## 5.4 Experimental Results

Full characterization of the system is done in two parts: the characterization of the optical fiber to waveguide coupling and the waveguide to device coupling. The setup used for this characterization is shown in Figure 5.4. One of the features of this coupling scheme is the large bandwidth and high efficiency. As such, a tunable laser is used with a range of 1520–1570 nm allowing the devices to be probed over a 50 nm bandwidth. An electro-optic modulator (EOM) is used with a feedback circuit to ensure a constant power over the tuning range of the laser. Fiber polarization controllers (FPC) are used to set the polarization of the light incident on the device and a circulator allows for the collection of reflected light.





**Figure 5.4: Experimental Setup.** The experimental setup used to characterize the single-sided, adiabatically tapered waveguide couplers. A tunable laser is used to characterize the devices and scan for optical resonances. When an optical resonance is found, the wavelength of the laser is stabilized using a wavemeter ( $\lambda$ -meter) reading in 10% of the light, having been split off using a 90:10 coupler. The light is then amplitude stabilized using an electro-optic modulator as an amplitude modulator (a-m) and then the desired power level is set using an variable optical attenuator (VOA). The polarization of the light is set (using a fiber polarization controller—FPC) before being sent into the 1 port of a circular, exiting the 2 port and coupling to the waveguide. The light reflects off the end mirror of the waveguide, re-entering the circulator at port 2 and exiting port 3. Here the light passes through a switch (SW1) and can be either sent to a slow detector (PD2) to monitor the voltage on reflection or a high-speed detector (PD1). The high-speed detector is connected to a vector network analyzer (VNA) and, along with the a-m, can be used to probe both the amplitude and phase response of the devices allowing for proper determine of under- or over-coupled devices. The final detector (PD3) is used to characterize the transmission if necessary.

An optical fiber taper with a diameter of  $\sim 2 \mu\text{m}$  is positioned on top of the tapered region of the waveguide using precision motorized stages. To provide the largest possible interaction region which maximizes the coupling, the relative angle between the fiber and waveguide is minimized and the fiber is touched down over a  $10\text{--}15 \mu\text{m}$  region of the tapered waveguide. The efficiency of the fiber-waveguide coupling can be determined by comparing the known input power with the power of the light reflected off the waveguide end mirrors (off-resonance from the device) and collected on the reflection port (3) of the circulator. Calibration of the circulator losses from port  $1 \rightarrow 2$  ( $\eta_{12}$ ) and port  $2 \rightarrow 3$  ( $\eta_{23}$ ), optical fiber taper loss ( $\eta_{\text{taper}}$ ), and the loss in optical switch SW1 allows the power measured on the photodetector PD2 to be directly compared to a known input power. The reflected power will also include any scattering loss in the waveguide itself and non-perfect reflection from the end mirrors on the waveguide, both effects we assume to be negligible and ignore. The efficiency measured varies slightly from device to device but is consistently over 90% for single pass coupling and often over 95%. Figure 5.5b shows the reflected signal as the laser is scanned over two different device resonances. The signal level is normalized to the input power so that off-resonance it simply corresponds to the reflected light. In these scans the background approaches a level of  $\sim 0.85$ , giving a single pass coupling efficiency of  $> 90\%$  for both devices.

Characterization of the waveguide-device coupling is more difficult now that the coupling is single-sided. For both a single and double-sided coupling, the normalized on-resonance reflection  $R_0$

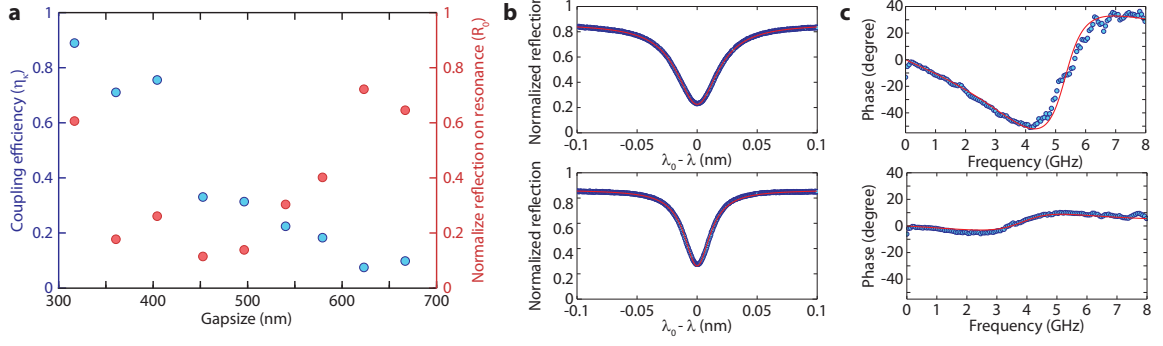


Figure 5.5: **Experimental Results.** (a) Normalized reflection depth on resonance  $R_0$  (red circles; right axis) and coupling efficiency  $\eta_\kappa$  (blue circles; left axis) plotted as a function of the waveguide – device gap size. As the gap size decreases,  $R_0$  decreases as  $\kappa_e/\kappa \rightarrow 1/2$ . For  $\kappa_e/\kappa = 1/2$  where  $\kappa_e = \kappa_i$ ,  $R_0 = 0$  indicating critical coupling. As the gap size continues to decrease,  $\kappa_e$  continues to increase, and  $R_0$  increases back to 1 for  $\kappa_e \approx \kappa$ . (b) Amplitude response of an overcoupled (top) and undercoupled (bottom) cavity. The overcoupled cavity has  $R_0 = 0.26$  and  $Q = 37,500$ , while the undercoupled cavity has  $R_0 = 0.30$  and  $Q = 60,000$ . To determine that these cavities are over- and undercoupled we need to further look at the phase response of the cavities plotted in c. The top plot is the phase response of the overcoupled cavity and the bottom is for the undercoupled cavity. While the amplitude responses are quite similar for the two cases, the phase responses are very different.

(equivalent to  $T_0$  in the double-sided coupling scheme we normally use) is given by  $R_0 = (1 - 2\eta)^2$ , where  $\eta$  is the coupling efficiency and defined as  $\eta_{ss} = \kappa_e/\kappa$  for single-sided coupling and  $\eta_{ds} = \kappa_e/2\kappa$  for double-sided coupling (see Appendix B.2.1 for more details). In the case of double-sided coupling, the  $T_0$  is a single-valued function of the extrinsic coupling  $\kappa_e$  and fitting a single transmission scan to a Lorentzian lineshape is enough to determine the coupling efficiency  $\eta_{ds}$ .

In the case of single-sided coupling, the on-resonance transmission decreases as  $\kappa_e \rightarrow \kappa_i$  and is exactly zero for  $\kappa_e = \kappa_i$ , which is called critical coupling. As the waveguide-device coupling is increased further, the on-resonance transmission increases and  $R_0 \rightarrow 1$  for  $\kappa_e \gg \kappa_i$  (i.e.  $\kappa \sim \kappa_e$ ). Fitting the linewidth of the optical response to determine  $R_0$  and  $\kappa$  is no longer enough to fully determine the coupling efficiency  $\eta$ , as  $R_0$  is not a single-valued function of  $\eta$  for single-sided coupling. This can be seen in Figure 5.5b which shows the reflected signal from the optical scans of two different devices, both of which have similar  $R_0$  values with different linewidths. Not knowing anything else about the devices (i.e., their intrinsic linewidths  $\kappa_i$ ) we can not determine the coupling efficiency. To determine the coupling efficiency for single-sided devices, we need to look at not just the amplitude response, but also the phase response.

We can examine the full, complex response of the optical cavities by using a vector network analyzer (VNA) to sweep a sideband across the cavity while looking at the amplitude and phase of this probe beam. To do this, the laser is locked off-resonance using the wavemeter ( $\lambda$ -meter; see Figure 5.4). The output of the VNA is connected to an electro-optic modulator (EOM) configured

for amplitude modulation (a-m) and generates a optical sideband that can be swept from 0–14 GHz across the optical resonance (this frequency range sets the limit to how far off-resonance the laser can be locked). The beating of the probe tone, as it is swept across the optical resonance, is mixed with the laser, generating a RF signal at the probe tone frequency which is detected on a high-speed photodetector (PD1) connected to the input of the VNA. By calibrating the response with no optical cavity present, the phase and amplitude response of the entire optical train is determined and used to normalize the signal. The amplitude and phase response of two optical cavities is shown in Figure 5.5a and b, respectively. As mentioned above, the amplitude response of the cavities is similar, however, the phase response is quite different due to the fact that the top plots correspond to an over-coupled device while the bottom plots correspond to an under-coupled one.

Using this method, the waveguide-device coupling can be determined for devices of varying gap size as shown in Figure 5.5a. The right-hand axis shows the on-resonance reflection  $R_0$  as a function of gap size, and as expected  $R_0$  at first decreases with decreasing gap size and then begins to increase again with increasing  $\kappa_e$ . The left-hand axis shows the equivalent coupling efficiency when both the amplitude and phase response are used to determine the waveguide-cavity coupling efficiency. Here the coupling efficiency increases roughly linearly with decreasing gap size and a maximum coupling efficiency approaching  $\eta = 0.9$  is demonstrated.

## 5.5 Discussion

By splitting the problem into two independent problems we have managed to engineer a coupling scheme that still allows the use of the fiber taper for rapid optical testing that couples to the device in a single-sided manner. The coupling of the device to the waveguide can be easily modified, while still allowing for very efficient coupling of the optical fiber to the waveguide. We have developed a method for coupling optical fiber tapers to on-chip waveguides with a demonstrated efficiency of 95% and then coupled this waveguide to a photonic crystal with an efficiency of 90%. This means that 85% of the signal photons in the optical cavity are coupled back into the optical fiber for detection. Highly efficient detection is required for measuring the motion of a mechanical oscillator at the standard quantum limit (SQL) [7, 54], performing feedback to cool a mechanical oscillator to its ground state [29, 105], for quantum feedback and control [106], and observing non-classical effects such as squeezing of the optical field [80, 107]. It is this last effect, the squeezing of the optical field due to the interaction of the radiation pressure shot noise with the mechanical oscillator, that was the impetus for the development of this very efficient coupling scheme.

## Appendix A

# Mathematical Tools and Definitions

### A.1 Fourier Transforms and Spectral Densities

$$\begin{aligned}
\hat{A}(t) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} d\omega e^{-i\omega t} \hat{A}(\omega) \\
\hat{A}(\omega) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} dt e^{i\omega t} \hat{A}(t) \\
\hat{A}^\dagger(\omega) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} dt e^{i\omega t} \hat{A}^\dagger(t) \\
\left(\hat{A}(\omega)\right)^\dagger &= \hat{A}^\dagger(-\omega) \\
S_{AA}(\omega) &= \int_{-\infty}^{\infty} d\tau e^{i\omega\tau} \langle \hat{A}^\dagger(t+\tau) \hat{A}(t) \rangle \\
S_{AA}(\omega) &= \int_{-\infty}^{\infty} d\omega' \langle \hat{A}^\dagger(\omega) \hat{A}(\omega') \rangle
\end{aligned} \tag{A.1}$$

### A.2 Noise Correlations

Correlations for noise operators of a thermal bath (i.e.,  $k_b T \gg \hbar\omega$  for temperatures  $T$  and frequencies  $\omega$  of interest).

$$\begin{aligned}
\langle \hat{b}_{\text{in}}^\dagger(t) \hat{b}_{\text{in}}(t') \rangle &= n_b \delta(t - t') \\
\langle \hat{b}_{\text{in}}(t) \hat{b}_{\text{in}}^\dagger(t') \rangle &= (n_b + 1) \delta(t - t') \\
\langle \hat{b}_{\text{in}}^\dagger(\omega) \hat{b}_{\text{in}}(\omega') \rangle &= n_b \delta(\omega + \omega') \\
\langle \hat{b}_{\text{in}}(\omega) \hat{b}_{\text{in}}^\dagger(\omega') \rangle &= (n_b + 1) \delta(\omega + \omega')
\end{aligned} \tag{A.2}$$

Correlations for noise operators of an optical bath (i.e.,  $k_b T \ll \hbar\omega$  for temperatures  $T$  and

frequencies  $\omega$  of interest).

$$\begin{aligned}
\langle \hat{a}_{\text{in}}^\dagger(\omega) \hat{a}_{\text{in}}(\omega') \rangle &= 0 \\
\langle \hat{a}_{\text{in}}(\omega) \hat{a}_{\text{in}}^\dagger(\omega') \rangle &= \delta(\omega + \omega') \\
\langle \hat{a}_{\text{in},i}^\dagger(\omega) \hat{a}_{\text{in},i}(\omega') \rangle &= 0 \\
\langle \hat{a}_{\text{in},i}(\omega) \hat{a}_{\text{in},i}^\dagger(\omega') \rangle &= \delta(\omega + \omega')
\end{aligned} \tag{A.3}$$

### A.3 Some Useful Commutator Relationships

A small collection of useful and important relations related to quantum mechanics:

$$\begin{aligned}
\dot{\hat{A}}(t) &= \frac{i}{\hbar} [\hat{H}, \hat{A}(t)] \\
[\hat{a}, \hat{a}^\dagger] &= 1 \\
[\hat{x}, \hat{p}^\dagger] &= i\hbar \\
[\hat{a}^\dagger \hat{a}, a] &= -\hat{a} \\
[\hat{a} \hat{a}^\dagger, a] &= \hat{a}^\dagger \\
e^{\hat{A}} \hat{B} e^{-\hat{A}} &= \hat{B} + [\hat{A}, \hat{B}] + \frac{1}{2!} [\hat{A}, [\hat{A}, \hat{B}]] + \dots
\end{aligned} \tag{A.4}$$

### A.4 The Mechanical Mode: A Bosonic Harmonic Oscillator

Quadratures of a Bosonic harmonic oscillator:

$$\begin{aligned}
\hat{X}_1 &= \hat{b}^\dagger + \hat{b} \\
\hat{X}_2 &= i(\hat{b}^\dagger - \hat{b})
\end{aligned} \tag{A.5}$$

For a mechanical mode, we can rescale and identify the quadratures with the real, measurable quantities position ( $\hat{x}$ ) and momentum ( $\hat{p}$ ). Keep in mind that they are simply the two quadratures of the complex amplitude of the Bosonic harmonic oscillator in one-dimension.

$$\begin{aligned}
\hat{x} &= \sqrt{\frac{\hbar}{2m\omega}} (\hat{b}^\dagger + \hat{b}) = x_{\text{zpf}} (\hat{b}^\dagger + \hat{b}), & x_{\text{zpf}} &= \sqrt{\frac{\hbar}{2m\omega}} \\
\hat{p} &= \sqrt{\frac{m\omega\hbar}{2}} i(\hat{b}^\dagger - \hat{b}) = p_{\text{zpf}} i(\hat{b}^\dagger - \hat{b}), & p_{\text{zpf}} &= \sqrt{\frac{m\omega\hbar}{2}}
\end{aligned} \tag{A.6}$$

Similarly, the annihilation and creation operators can be found in terms of the position and momentum operators.

$$\begin{aligned}
\hat{b} &= \sqrt{\frac{m\omega}{2\hbar}} \left( \hat{x} + \frac{i}{m\omega} \hat{p} \right) \\
\hat{b}^\dagger &= \sqrt{\frac{m\omega}{2\hbar}} \left( \hat{x} - \frac{i}{m\omega} \hat{p} \right)
\end{aligned} \tag{A.7}$$

The number operator is given by,

$$\begin{aligned}
 \hat{n} = \hat{b}^\dagger \hat{b} &= \frac{m\omega}{2\hbar} \left( \hat{x} - \frac{i}{m\omega} \hat{p} \right) \left( \left( \hat{x} + \frac{i}{m\omega} \hat{p} \right) \right. \\
 &= \frac{m\omega}{2\hbar} \left( \hat{x}^2 + \frac{i}{m\omega} (\hat{x} \hat{p} - \hat{p} \hat{x}) + \frac{\hat{p}^2}{m^2 \omega^2} \right) \\
 &= \frac{m\omega}{2\hbar} \left( \hat{x}^2 - \frac{\hbar}{m\omega} + \frac{\hat{p}^2}{m^2 \omega^2} \right) \\
 &= \frac{m\omega \hat{x}^2}{2\hbar} + \frac{\hat{p}^2}{2m\omega\hbar} - \frac{1}{2}
 \end{aligned} \tag{A.8}$$

where the commutator for position and momentum was used. When substituted into the Hamiltonian for a one dimensional Bosonic mode,

$$\begin{aligned}
 H &= \hbar\omega \left( \hat{b}^\dagger \hat{b} + \frac{1}{2} \right) \\
 &= \hbar\omega \left( \frac{m\omega \hat{x}^2}{2\hbar} + \frac{\hat{p}^2}{2m\omega\hbar} - \frac{1}{2} + \frac{1}{2} \right) \\
 &= \frac{m\omega^2 \hat{x}^2}{2} + \frac{\hat{p}^2}{2m},
 \end{aligned} \tag{A.9}$$

which is the familiar equation for the energy of a one dimensional simple harmonic oscillator.

## Appendix B

# Mathematical Description of an Optical Cavity

Here we provide the necessary mathematical tools and background to model an electromagnetic cavity. We will focus on a single mode of an optical cavity, though the framework used will be easily extensible to multiple modes as other bosonic modes (e.g., microwave modes and phononic modes).

### B.1 A Single Optical Cavity

We begin by considering the Hamiltonian for a single optical mode (e.g., of a photonic crystal)

$$H = \hbar\omega_o \hat{a}^\dagger \hat{a},$$

with  $\hat{a}$  ( $\hat{a}^\dagger$ ) the annihilation (creation) operator for the bosonic optical mode with frequency  $\omega_o$ . Here the mode is modeled as a harmonic oscillator and we can identify the number (mode occupation) operator as  $\hat{n} = \hat{a}^\dagger \hat{a}$ . The Hamiltonian thus represents the energy of the system: the number of photons ( $\hat{n}$ ) multiplied by the energy of each photon ( $\hbar\omega_o$ ), where we are ignoring the contribution from the zero-point fluctuations. In the Heisenberg picture, we can find the equation of motion of an operator (with no explicit time dependence) using  $\hbar\dot{\hat{A}}(t) = i[\hat{H}, \hat{A}(t)]$  and the bosonic commutation relation  $[\hat{a}, \hat{a}^\dagger] = 1$  giving

$$\dot{\hat{a}} = -i\omega_o \hat{a}$$

which can easily be solved to give  $\hat{a}(t) = \hat{a}(0)e^{-i\omega_o t}$  relating the Schrödinger operator,  $\hat{a}(0)$ , to the Heisenberg operator,  $\hat{a}(t)$ . In this picture the Heisenberg annihilation operator can be seen to rotate in the Hilbert space at a frequency  $-\omega_o$  corresponding to a decrease in the mode energy. The creation operator rotates at  $\omega_o$  and corresponds to an increase in the mode energy. This description, however, is not very interesting as it lacks any inputs/outputs, noise, and dissipation. To account for these terms, we need to first determine the type of input/output coupling appropriate for our

device.

## B.2 Input/Output Coupling and Dissipation

An isolated optical cavity is not very interesting, with no means of sending light into it or extracting information from it, nor is it realistic. With this in mind we extend our model to include both methods to inject a signal into our cavity and ways a signal in the cavity can leak out:

$$\dot{\hat{a}} = -i\omega_o\hat{a} - \frac{\kappa}{2}\hat{a} - \sqrt{\Lambda}\hat{a}_{\text{in}}e^{-i\omega_\ell t} - \sqrt{\kappa'}\hat{a}_{\text{in},i}. \quad (\text{B.1})$$

The input term is  $\sqrt{\Lambda}\hat{a}_{\text{in}}e^{-i\omega_\ell t}$ , representing a driving coherent excitation (e.g., a laser) at a frequency of  $\omega_\ell$ . This driving field is coupled to the mode at a rate  $\Lambda$  and the exact form of this coupling rate can take several forms, as will be discussed shortly. In addition to the input term, there is a dissipation term  $(\kappa/2)\hat{a}$  modeling the rate at which a signal in the cavity is lost for any reason (e.g., imperfect design, fabrication, input/output coupling, etc.). As can be seen from the equation of motion of the annihilation operator  $\hat{a}$ , this term represents the field decay rate and gives an energy decay rate of  $\kappa$ . The final term  $\sqrt{\kappa'}\hat{a}_{\text{in},i}$  is a consequence of the fluctuation-dissipation theorem and is the noise input term associated with the cavity loss. The coupling rate of this term,  $\kappa'$ , represents the rate at which information about the cavity is lost (intrinsic dissipation and external coupling where the signal is not detected).

### B.2.1 Single and Double-Sided Coupling

Coupling of the cavity to a waveguide for input and output can take two general forms, single and double-sided coupling as shown in Figure B.1. For single-sided coupling, the cavity mode is coupled to the waveguide unidirectionally at a rate  $\kappa_e$ . Photons in the cavity will now see two loss rates associated with the intrinsic cavity loss due to scattering and imperfections  $\kappa_i$  and the coupling to the waveguide  $\kappa_e$ , giving a total cavity loss rate of  $\kappa = \kappa_i + \kappa_e$ . We can define the coupling efficiency

$$\eta_{\text{ss}} = \frac{\kappa_e}{\kappa}, \quad (\text{B.2})$$

representing the fraction of the photons in the cavity that we can couple out of the cavity to be detected as well as couple into the cavity. For single-sided coupling the maximum coupling efficiency is  $\eta_{\text{ss}} = 1$  for  $\kappa_e \gg \kappa_i$ . This is the best case scenario and gives an undetected photon loss rate of  $\kappa' = \kappa_i$ , only the photons lost through the device's intrinsic loss channel go undetected. This type of coupling occurs in structures requiring a phase matching condition (disks, rings, toroids, etc.) and butt-coupled device (including fiber butt-coupling and taper coupling to waveguides which are



subsequently butt-coupled to the device [93, 95]). We have also developed a very efficient single-sided coupling scheme as outlined in Chapter 5.




Figure B.1: **Schematic of single and double-sided coupling.** (a) Single-sided coupling. (b) Double-sided coupling.

For double-sided coupling, the device couples to two channels of the waveguide. In the case of an optical fiber taper, this is the forward and backward direction along the axis of the taper. The coupling between the device and each of these channels is  $\kappa_e/2$ . The device now sees its intrinsic loss channel as well as two external couplings to the waveguide giving a total loss rate of  $\kappa = \kappa_i + \kappa_e$ , the same as the case of single-sided coupling. However, if we assume that the input light field is incident on the device from one direction only, a very good assumption, it only couples into the device at a rate  $\kappa_e/2$ ; at best only half of the light is coupled into the device. Similarly, light coupled from the device to the waveguide sees two channels with equal coupling rates,  $\kappa_e/2$  and couples into each equally, sending half the light down the fiber in one direction and half of the light the other direction. In general detection is only done at one end of the optical fiber meaning at best half the light will be detected. We can again define a coupling efficiency,

$$\eta_{ds} = \frac{\kappa_e}{2\kappa}, \quad (\text{B.3})$$

with a maximum efficiency of  $\eta_{ds} = 1/2$  for  $\kappa_e \gg \kappa_i$ . In the best case, 50% of the light can be coupled into the device and 50% coupled out into our detector (single-sided detection). The loss rate for photons that go undetected is now  $\kappa' = \kappa_i + \kappa_e/2$ . This is the type of coupling encountered when using an optical fiber taper to couple directly to photonic crystal cavities (both one-dimensional and two-dimensional).

### B.2.2 Solution to the Cavity Equation of Motion

For both the single and double-sided coupling cases we can, without loss of generality, define the input/output coupling rate as  $\Lambda = \eta\kappa$  and leave the information loss rate as  $\kappa'$ . The equation of motion for the cavity we wish to solve is then

$$\dot{\hat{a}} = -i\omega_o \hat{a} - \frac{\kappa}{2} \hat{a} - \sqrt{\eta\kappa} \hat{a}_{in} e^{-i\omega_\ell t} - \sqrt{\kappa'} \hat{a}_{in,i} \quad (\text{B.4})$$

with total energy loss (decay) rate  $\kappa$ , input/output coupling  $\eta\kappa$  to an external source (waveguide)  $\hat{a}_{\text{in}}$ , and coupling  $\kappa'$  which denotes all optical loss channels that go undetected (information lost) and associated noise input  $\hat{a}_{\text{in},i}$ . It should be kept in mind that  $a$  represents an electric field, so that  $\hat{a}^\dagger\hat{a}$  represents the photon number of mode  $\hat{a}$  while  $\hat{a}_{\text{in}}^\dagger\hat{a}_{\text{in}}$  represents the number of photons/s incident on the device. This difference explains the power of the coupling rates in Equation (B.4). To solve this equation of motion it is common to move into a frame rotating at the frequency of the laser drive,  $\omega_\ell$ , simplifying the dynamics of the problem. Thus, defining,

$$\tilde{\hat{a}} = \hat{a}e^{i\omega_\ell t} \quad (\text{B.5})$$

as the slowly varying component of  $\hat{a}$  we find

$$a = \tilde{a}e^{-i\omega_\ell t} \quad (\text{B.6a})$$

$$\dot{a} = \dot{\tilde{a}}e^{-i\omega_\ell t} - i\omega_\ell\tilde{a}e^{-i\omega_\ell t}. \quad (\text{B.6b})$$

Substituting Equation (B.6a) and (B.6b) into Equation (B.4) while making the substitution  $\tilde{a} \rightarrow a$  (for ease and clarity) and multiplying both sides by  $e^{i\omega_\ell t}$  gives

$$\dot{a} = \left(-i\Delta - \frac{\kappa}{2}\right)a - \sqrt{\eta\kappa}a_{\text{in}} \quad (\text{B.7})$$

as the equation of motion for the annihilation operator with  $\Delta = \omega_o - \omega_\ell$  the detuning of the laser from the cavity;  $\Delta > 0$  implies a red detuning and  $\Delta < 0$  a blue detuning with  $\Delta = 0$  being on resonance. For now, we have ignored the (vacuum) noise input term  $\hat{a}_{\text{in},i}$ .

Solving for the steady state solution  $\dot{a} = 0$  gives

$$a = \frac{-\sqrt{\eta\kappa}a_{\text{in}}}{i\Delta + \kappa/2}. \quad (\text{B.8})$$

We can see that the cavity field,  $a$ , is  $180^\circ$  out of phase with the driving field  $a_{\text{in}}$  (recall  $e^{i\pi} = -1$ ). For the transmitted signal in the double-sided case and the reflected signal in the single-sided case we can write the output field as

$$a_{\text{out}} = a_{\text{in}} + \sqrt{\eta\kappa}a, \quad (\text{B.9})$$

which gives the optical power at the detector as  $|a_{\text{out}}|^2$ . The transmission coefficient can then be defined as  $T = \left|\frac{a_{\text{out}}}{a_{\text{in}}}\right|^2$  giving

$$T = \left|1 - \frac{\eta\kappa}{i\Delta + \frac{\kappa}{2}}\right|^2 \quad (\text{B.10})$$

from which we can define the on-resonance ( $\Delta = 0$ ) transmission as  $T_0 = (1 - 2\eta)^2$  for both single and double-sided coupling. Thus, from a transmission scan of the cavity, knowing the linewidth

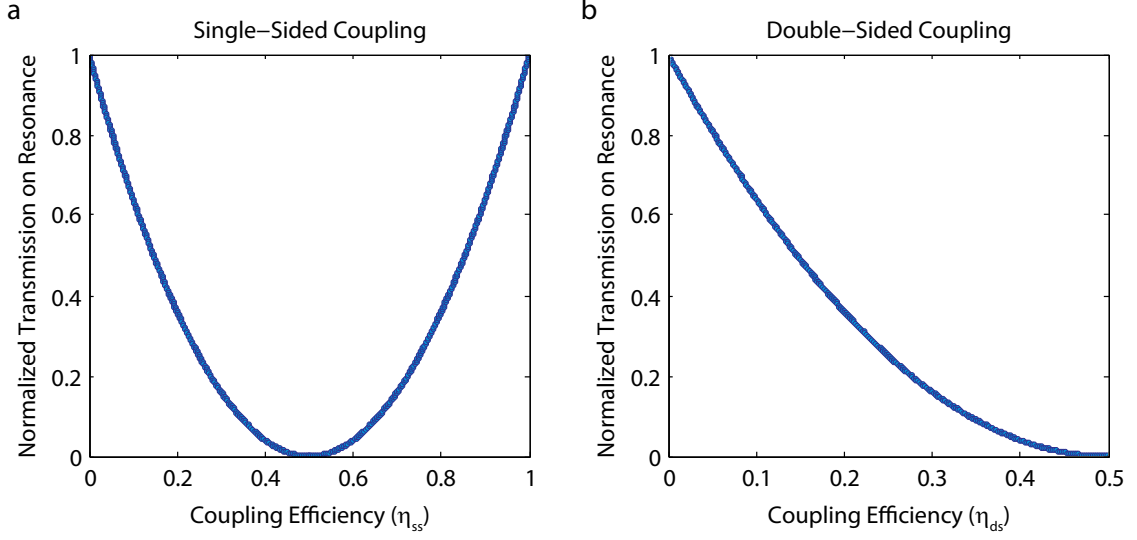


Figure B.2: **Transmission depth on resonance for single and double-sided coupling.** (a) Single-sided coupling. (b) Double-sided coupling.

(proportional to  $\kappa$ ) and the normalized transmission on-resonance ( $T_0$ ), the intrinsic loss rate  $\kappa_i$ , and the external coupling rate  $\kappa_e$  can be determined. For a double-sided cavity, we can solve for  $\kappa_e$  and  $\kappa_i$  knowing  $T_0$  and  $\kappa$  with,

$$\begin{aligned} \frac{\kappa_e}{\kappa} &= 1 - \sqrt{T_0} \\ \frac{\kappa_i}{\kappa} &= \sqrt{T_0}. \end{aligned} \quad (\text{B.11})$$

The situation for a single-sided cavity is more complicated, as more than one value for the coupling efficiency can lead to the same on-resonance normalized transmission as seen in Figure B.2. Unlike the double-sided case, here we can not solve for unique values of  $\kappa_e$  and  $\kappa_i$  given only  $T_0$  and  $\kappa$  - the values we can extract from a simple, calibrated transmission scan:

$$\begin{aligned} \frac{\kappa_e}{\kappa} &= \frac{1 - \sqrt{T_0}}{2} \\ \frac{|\kappa_i - \kappa_e|}{\kappa} &= \sqrt{T_0}. \end{aligned} \quad (\text{B.12})$$

Given a transmission depth, with no other information it is impossible to know the coupling efficiency. Knowing, or guessing, the intrinsic loss rate  $\kappa_i$  is enough to solve this problem. However, looking at the phase response as well as the amplitude response (full complex response) fully determines the coupling efficiency for both the single and double-sided coupling cases.

For a given optical power at the device ( $P_d$ ) the intracavity photon number,  $n_c$  can be calculated from

$$n_c = a^\dagger a = \frac{\eta \kappa |a_{\text{in}}|^2}{\Delta^2 + \left(\frac{\kappa}{2}\right)^2} \quad (\text{B.13})$$

with  $|a_{\text{in}}|^2 = P_d/(\hbar\omega_o)$ . It is easy to see that for the single-sided case,  $\eta\kappa = \kappa_e$ , the intracavity photon number is twice what it is for double-sided coupling when  $\eta\kappa = \kappa_e/2$ , as would be expected for the same parameters since the coupling efficiency for the double-sided case would be double that of the single-sided case.

## Appendix C

# Power Spectral Density of the Mechanical Mode

Here we will evaluate the spectral density of the mechanical mode first without any optomechanical coupling (i.e.,  $G = 0$ ) to see what a mechanical mode driven only by a thermal bath looks like. Next, the coupling will be reinstated, increasing the complexity but providing insight on how the optomechanical coupling affects the mechanical mode. The full equation for the annihilation operator in the Fourier domain is:

$$\hat{b}(\omega) = \frac{-\sqrt{\gamma_i}\hat{b}_{\text{in}}(\omega) - iG(\hat{a}^\dagger(\omega) + \hat{a}(\omega))}{i(\omega_m - \omega) + \gamma/2}. \quad (\text{C.1})$$

### C.1 Mechanical Mode Coupled to Thermal Bath

In this section we will calculate the power spectral density of a mechanical mode coupled only to a thermal bath. The starting point for this is the equation of motion of the mechanical mode in the Fourier domain (Equation C.1), with the optomechanical coupling set to zero ( $G = 0$ ) and the definition of the power spectral density from Appendix A.1. The power spectral density of the annihilation operator is then

$$\begin{aligned} S_{\text{bb}}(\omega) &= \int_{-\infty}^{\infty} d\omega' \langle \hat{b}^\dagger(\omega) \hat{b}(\omega') \rangle \\ &= \frac{\gamma_i}{-i(\omega_m + \omega) + \gamma_i/2} \int_{-\infty}^{\infty} d\omega' \frac{\langle \hat{b}_{\text{in}}^\dagger(\omega) \hat{b}_{\text{in}}(\omega') \rangle}{i(\omega_m - \omega') + \gamma_i/2} \\ &= \frac{\gamma_i \bar{n}}{(\omega_m + \omega)^2 + (\gamma_i/2)^2}, \end{aligned} \quad (\text{C.2})$$

where the correlation for the mechanical noise operators from Appendix A.2 has been used. The result is a Lorentzian centered around  $-\omega_m$ , with width  $\gamma_i$  and an area proportional to  $\bar{n}$ , the average phonon population of the mechanical mode which in this case is in thermal equilibrium with the bath and thus  $\bar{n} = n_{\text{b}}$ . This PSD represents the mechanical mode emitting energy (in this case

into the bath as it is coupled to nothing else), which is why it is proportional to  $\bar{n}$ ; if the mode has no phonons it can not emit energy. Further, we know that we can interpret Lorentzians centered around  $-\omega_m$  as emitting energy (and as such it makes sense since this is the PSD of the annihilation operator).

We look next at the PSD for the creation operator,

$$\begin{aligned} S_{b^\dagger b^\dagger}(\omega) &= \int_{-\infty}^{\infty} d\omega' \langle \hat{b}(\omega) \hat{b}^\dagger(\omega') \rangle \\ &= \frac{\gamma_i(\bar{n} + 1)}{(\omega_m - \omega)^2 + (\gamma_i/2)^2}. \end{aligned} \quad (\text{C.3})$$

The result is again a Lorentzian with width  $\gamma_i$ , however it is now centered around  $+\omega_m$  and has an area proportional to  $(\bar{n} + 1)$ . This PSD represents the ability of the mechanical mode to absorb energy (it is the PSD of the creation operator) and we can see even if  $\bar{n} = 0$ , the case of a quantum ground state mechanical oscillator, it can still absorb energy.

At this point we can make a quick comment regarding the PSDs of the annihilation and creation operators. With the annihilation operator's PSD being proportional to  $\bar{n}$  and the creation operator's PSD proportional to  $(\bar{n} + 1)$ , at high average phonon occupancies the positive and negative frequency PSDs are of equal magnitude. At low phonon occupancy, however, we can see an asymmetry starts to develop, becoming increasingly pronounced as  $\bar{n} \rightarrow 0$  and the resonator becomes “more quantum”—in this case we are referring to the relative magnitudes of the thermal noise contribution to the phonon population and the contribution from the Heisenberg uncertainty (the  $1/2$ ). This asymmetry does not exist in classical mechanics where the positive and negative frequency PSDs have equal amplitudes [7, 78].

The annihilation and creation operators, however, are not quantum observables of the mechanical mode, and as such it is interesting to also look at the PSD of the position operator; the in-phase/real quadrature of the mode. We can do this by using the quadrature definition  $\hat{x} = x_{\text{zpf}}(\hat{b}^\dagger + \hat{b})$ . Then,

$$\begin{aligned} S_{xx}(\omega) &= \int_{-\infty}^{\infty} d\omega' \langle \hat{x}^\dagger(\omega) \hat{x}(\omega') \rangle \\ &= x_{\text{zpf}}^2 \int_{-\infty}^{\infty} d\omega' \langle (\hat{b}(\omega) + \hat{b}^\dagger(\omega)) (\hat{b}^\dagger(\omega') + \hat{b}(\omega')) \rangle \\ &= x_{\text{zpf}}^2 (S_{b^\dagger b^\dagger}(\omega) + S_{bb}(\omega)). \end{aligned} \quad (\text{C.4})$$

Here we note that the real observable is composed of the Lorentzians centered around  $+\omega_m$  and  $-\omega_m$  and normally they are folded back on themselves and we see a symmetrized signal at  $\omega_m$ . If we could somehow look at the positive and negative frequency components of this real observable separately we would be able to see the asymmetry associated with the  $\hat{b}^\dagger$  and  $\hat{b}$  terms, which is exactly what Safavi-Naeini, et al. did [78]. An asymmetric PSD can only occur in quantum mechanics and is a

direct result of the the position and momentum operators do not commute; it is always symmetric in classical mechanics.

## C.2 Mechanical Mode with Optomechanical Coupling

We now wish to consider the PSD of the annihilation operator, including the optomechanical coupling ( $G \neq 0$ ). To do this we need to substitute the equation for the annihilation and creation operators for the optical mode into Equation C.1 given by

$$\hat{a}(\omega) = \frac{-\sqrt{\eta\kappa}\hat{a}_{\text{in}}(\omega) - \sqrt{\kappa'}\hat{a}_{\text{in},i}(\omega) - iG(\hat{b}^\dagger(\omega) + \hat{b}(\omega))}{i(\Delta - \omega) + \kappa/2}. \quad (\text{C.5})$$

We then have for the annihilation and creation operator of the mechanical mode,

$$\begin{aligned} \hat{b}(\omega) &= \frac{-\sqrt{\gamma_i}\hat{b}_{\text{in}}(\omega)}{i(\omega_{\text{m}} - \omega) + \gamma/2} \\ &\quad + \frac{iG(\sqrt{\eta\kappa}\hat{a}_{\text{in}}^\dagger(\omega) + \sqrt{\kappa'}\hat{a}_{\text{in},i}^\dagger(\omega))}{[i(\omega_{\text{m}} - \omega) + \gamma/2][-i(\Delta + \omega) + \kappa/2]} \\ &\quad + \frac{iG(\sqrt{\eta\kappa}\hat{a}_{\text{in}}(\omega) + \sqrt{\kappa'}\hat{a}_{\text{in},i}(\omega))}{[i(\omega_{\text{m}} - \omega) + \gamma/2][i(\Delta - \omega) + \kappa/2]} \\ \hat{b}^\dagger(\omega) &= \frac{-\sqrt{\gamma_i}\hat{b}_{\text{in}}(\omega)}{-i(\omega_{\text{m}} + \omega) + \gamma/2} \\ &\quad + \frac{-iG^*(\sqrt{\eta\kappa}\hat{a}_{\text{in}}^\dagger(\omega) + \sqrt{\kappa'}\hat{a}_{\text{in},i}^\dagger(\omega))}{[-i(\omega_{\text{m}} + \omega) + \gamma/2][-i(\Delta + \omega) + \kappa/2]} \\ &\quad + \frac{-iG^*(\sqrt{\eta\kappa}\hat{a}_{\text{in}}(\omega) + \sqrt{\kappa'}\hat{a}_{\text{in},i}(\omega))}{[-i(\omega_{\text{m}} + \omega) + \gamma/2][i(\Delta - \omega) + \kappa/2]} \end{aligned} \quad (\text{C.6})$$

Defining,

$$\chi_{\text{m}}(\omega) = \frac{1}{i(\omega_{\text{m}} - \omega) + \gamma/2} \quad \chi_{\text{o}}(\omega) = \frac{1}{i(\Delta - \omega) + \kappa/2} \quad (\text{C.7})$$

we can simplify the equations giving a neater

$$\begin{aligned} \hat{b}(\omega) &= -\chi_{\text{m}}(\omega)\sqrt{\gamma_i}\hat{b}_{\text{in}}(\omega) \\ &\quad + iG\chi_{\text{m}}(\omega)\chi_{\text{o}}^*(-\omega)(\sqrt{\eta\kappa}\hat{a}_{\text{in}}^\dagger(\omega) + \sqrt{\kappa'}\hat{a}_{\text{in},i}^\dagger(\omega)) \\ &\quad + iG\chi_{\text{m}}(\omega)\chi_{\text{o}}(\omega)(\sqrt{\eta\kappa}\hat{a}_{\text{in}}(\omega) + \sqrt{\kappa'}\hat{a}_{\text{in},i}(\omega)) \\ \hat{b}^\dagger(\omega) &= -\chi_{\text{m}}^*(-\omega)\sqrt{\gamma_i}\hat{b}_{\text{in}}(\omega) \\ &\quad - iG^*\chi_{\text{m}}(-\omega)^*\chi_{\text{o}}^*(-\omega)(\sqrt{\eta\kappa}\hat{a}_{\text{in}}^\dagger(\omega) + \sqrt{\kappa'}\hat{a}_{\text{in},i}^\dagger(\omega)) \\ &\quad - iG^*\chi_{\text{m}}(-\omega)^*\chi_{\text{o}}(\omega)(\sqrt{\eta\kappa}\hat{a}_{\text{in}}(\omega) + \sqrt{\kappa'}\hat{a}_{\text{in},i}(\omega)). \end{aligned} \quad (\text{C.8})$$

The PSD of the annihilation operator is then,

$$\begin{aligned}
S_{bb}(\omega) &= |\chi_m(-\omega)|^2 \gamma_i n_b + |G|^2 |\chi_m(-\omega)|^2 |\chi_o(\omega)|^2 (\eta\kappa + \kappa') \\
&= |\chi_m(-\omega)|^2 \gamma_i n_b + |G|^2 |\chi_m(-\omega)|^2 |\chi_o(\omega)|^2 (\kappa) \\
&= \frac{\gamma n_f}{(\omega_m + \omega)^2 + (\gamma/2)^2},
\end{aligned} \tag{C.9}$$

where we have used the fact that  $\eta\kappa + \kappa' = \kappa$  for both single and double-sided coupling, the definitions for the correlation between  $\hat{a}_{in}(\omega)$ ,  $\hat{a}_{in}^\dagger(\omega)$ ,  $\hat{a}_{in,i}(\omega)$ ,  $\hat{a}_{in,i}^\dagger(\omega)$  (Equation A.3), and defined the final mechanical mode occupation,

$$n_f = \frac{\gamma_i n_b}{\gamma} + \frac{|G|^2 \kappa}{\gamma} \frac{1}{(\Delta - \omega)^2 + (\kappa/2)^2}. \tag{C.10}$$

We can see from Equation C.10 that the first term is the phonon population due to the coupling to the thermal bath with a factor due to any modification of  $\gamma$  from its intrinsic value  $\gamma_i$ . The second term is due to the optomechanical coupling and is the quantum backaction on the mechanical mode from the optical mode.

### C.2.1 Red-Detuned Pump ( $\Delta = \omega_m$ )

For the case of a red-detuned optical pump ( $\Delta = +\omega_m$ ) we see

$$n_f|_{\Delta=+\omega_m} = \frac{\gamma_i n_b}{\gamma} + \frac{G^2 \kappa}{\gamma} \frac{1}{(2\omega_m)^2} = \frac{\gamma_i n_b}{\gamma} + \frac{\gamma_{OM}}{\gamma} \left( \frac{\kappa}{4\omega_m} \right)^2. \tag{C.11}$$

Here we have made the assumption our system is side-band resolved,  $\omega_m \gg \kappa$ , and  $\omega = -\omega_m$  ( $S_{bb}$  is peaked here and has support only over a small frequency range,  $\gamma$ , centered around  $\omega = -\omega_m$ ).

### C.2.2 On-Resonance Pump ( $\Delta = 0$ )

With the pump beam tuned to the resonance frequency of the optical cavity we have  $\Delta = 0$  and  $\gamma_{OM} = 0$  ( $\gamma = \gamma_i$ ). Again we look only around  $\omega = -\omega_m$  and assume  $\omega_m \gg \kappa$  giving,

$$n_f|_{\Delta=0} = n_b + \frac{G^2 \kappa}{\gamma_i} \frac{1}{\omega_m^2} = n_b + \frac{4G^2}{\kappa \gamma_i} \left( \frac{\kappa}{2\omega_m} \right)^2. \tag{C.12}$$

The first term is clearly the unperturbed phonon occupation due to the coupling to the thermal bath. The second term scales with input power (recall  $G = \sqrt{n_c}$ ) and with no modification to the mechanical mode due to the optomechanical coupling (i.e., no dynamical optomechanical backaction) is the quantum backaction term. This term means that the harder we attempt to measure the mechanical mode (i.e., with more and more optical power to beat down the optical shot noise which



scales with  $N$ , the photon number) we modify the phonon number and measure it less and less accurately.

### C.2.3 Blue-Detuned Pump ( $\Delta = -\omega_{\text{m}}$ )

The last case we look at is with the optical pump blue-detuned ( $\Delta = -\omega_{\text{m}}$ ). In this case the optical spring shift is opposite compared to the red-detuned case, and the optomechanically induced damping ( $\gamma_{\text{OM}}$ ) subtracts from the intrinsic damping  $\gamma_{\text{i}}$ . This gives

$$n_{\text{f}}|_{\Delta=-\omega_{\text{m}}} = \frac{\gamma_{\text{i}} n_{\text{b}}}{\gamma} + \frac{G^2 \kappa}{\gamma} \left( \frac{2}{\kappa} \right)^2 = \frac{\gamma_{\text{i}} n_{\text{b}}}{\gamma} + \frac{|\gamma_{\text{OM}}|}{\gamma}, \quad (\text{C.13})$$

which is valid only for  $\gamma_{\text{OM}} < \gamma_{\text{i}}$  after which the system becomes unstable as it starts to oscillate (self-sustained oscillations). The behavior in this regime becomes non-linear and non-linear effects clamp the amplitude of oscillation to some finite value.

# Appendix D

## Publications

1. A. H. Safavi-Naeini, S. Gröblacher, J. T. Hill, J. Chan, M. Aspelmeyer, and O. Painter, “Squeezing of light via reflection from a silicon micromechanical resonator,” *arXiv:1302.6179* (2013).
2. 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.* **15**, 035007– (2013).
3. J. T. Hill, A. H. Safavi-Naeini, J. Chan, and O. Painter, “Coherent optical wavelength conversion via cavity optomechanics,” *Nat. Commun.* **3**, 1196– (2012).
4. 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.* **101**, 081115– (2012).
5. 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.* **108**, 033602– (2012).
6. J. Chan, T. P. M. Alegre, A. H. Safavi-Naeini, J. T. Hill, A. Krause, S. Gröblacher, M. Aspelmeyer, and O. Painter, “Laser cooling of a nanomechanical oscillator into its quantum ground state,” *Nature* **478**, 89–92 (2011).
7. A. H. Safavi-Naeini, J. Chan, M. Eichenfield, Q. Lin, J. T. Hill, D. Chang, and O. Painter, “Electromagnetically induced transparency and slow light with optomechanics,” *Nature* **472**,

69–73 (2011).

# Bibliography

- [1] V. B. Braginsky, Y. I. Vorontsov, and K. S. Thorne, “Quantum Nondemolition Measurements,” *Science* **209**, 547–557 (1980).
- [2] J. B. Hertzberg, T. Rocheleau, T. Ndukum, M. Savva, A. A. Clerk, and K. C. Schwab, “Back-action-evading measurements of nanomechanical motion,” *Nature Phys.* **6**, 213–217 (2010).
- [3] 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* **463**, 72–75 (2010).
- [4] J. D. Thompson, B. M. Zwickl, A. M. Jayich, F. Marquardt, S. M. Girvin, and J. G. E. Harris, “Strong dispersive coupling of a high-finesse cavity to a micromechanical membrane,” *Nature* **452**, 72–75 (2008).
- [5] A. M. Jayich, J. C. Sankey, B. M. Zwickl, C. Yang, J. D. Thompson, S. M. Girvin, A. A. Clerk, F. Marquardt, and J. G. E. Harris, “Dispersive optomechanics: a membrane inside a cavity,” *New J. Phys.* **10**, 095008 (2008).
- [6] H. Miao, S. Danilishin, T. Corbitt, and Y. Chen, “Standard Quantum Limit for Probing Mechanical Energy Quantization,” *Phys. Rev. Lett.* **103**, 100402 (2009).
- [7] A. A. Clerk, F. Marquardt, and J. G. E. Harris, “Quantum Measurement of Phonon Shot Noise,” *Phys. Rev. Lett.* **104**, 213603 (2010).
- [8] A. Nunnenkamp, K. Brkje, J. G. E. Harris, and S. M. Girvin, “Cooling and squeezing via quadratic optomechanical coupling,” *Phys. Rev. A* **82**, 021806 (2010).
- [9] M. Vanner, “Selective Linear or Quadratic Optomechanical Coupling via Measurement,” *Phys. Rev. X* **1**, 021011 (2011).
- [10] A. Xuereb and M. Paternostro, “Selectable linear or quadratic coupling in an optomechanical system,” *Phys. Rev. A* **87**, 023830 (2012).
- [11] M. Alayo, I. Pereyra, and M. Carreo, “Thick SiOxNy and SiO2 films obtained by PECVD technique at low temperatures,” *Thin Solid Films* **332**, 40–45 (1998).

- [12] Y. Kim, S. Cho, Y. Seo, H. Yoon, Y. Im, and D. Yoon, “Influence of hydrogen on SiO<sub>2</sub> thick film deposited by PECVD and FHD for silica optical waveguide,” *Cryst. Res. Technol.* **37**, 1257–1263 (2002).
- [13] Q. Lin, J. Rosenberg, X. Jiang, K. J. Vahala, and O. Painter, “Mechanical Oscillation and Cooling Actuated by the Optical Gradient Force,” *Phys. Rev. Lett.* **103**, 103601 (2009).
- [14] J. Rosenberg, Q. Lin, and O. Painter, “Static and dynamic wavelength routing via the gradient optical force,” *Nat Photon.* **3**, 478–483 (2009).
- [15] Q. Lin, J. Rosenberg, D. Chang, R. Camacho, M. Eichenfield, K. J. Vahala, and O. Painter, “Coherent mixing of mechanical excitations in nano-optomechanical structures,” *Nature Photon.* **4**, 236–242 (2010).
- [16] N. E. Flowers-Jacobs, S. W. Hoch, J. C. Sankey, A. Kashkanova, A. M. Jayich, C. Deutsch, J. Reichel, and J. G. E. Harris, “Fiber-cavity-based optomechanical device,” *Appl. Phys. Lett.* **101**, 221109 (2012).
- [17] L. Janis Research Company, “Janis ST-500 Continuous Flow Cryostat.”  
<http://www.janis.com/products/productsoverview/SuperTranContinuousFlowCryostatSystems/ST-500MicroscopyCryostat.aspx>
- [18] I. Custom Thermoelectric, “03111-9J30-20CA.”  
[http://www.customthermoelectric.com/tecs/pdf/03111-9J30-20CA\\_spec\\_sht.pdf](http://www.customthermoelectric.com/tecs/pdf/03111-9J30-20CA_spec_sht.pdf)
- [19] R. M. Camacho, J. Chan, M. Eichenfield, and O. Painter, “Characterization of radiation pressure and thermal effects in a nanoscale optomechanical cavity,” *Opt. Express* **17**, 15726–15735 (2009).
- [20] E. D. Black, “An introduction to Pound–Drever–Hall laser frequency stabilization,” *Am. J. Phys.* **69**, 79–87 (2001).
- [21] T. J. Kippenberg and K. J. Vahala, “Cavity Opto-Mechanics,” *Opt. Express* **15**, 17172–17205 (2007).
- [22] J. Chan, *Laser cooling of an optomechanical crystal resonator to its quantum ground state of motion.*, Ph.D. thesis, California Institute of Technology (2012).
- [23] J. T. Hill, Q. Lin, J. Rosenberg, and O. Painter, “Mechanical Trapping in a Quadratically Coupled Optomechanical Double Disk,” in *CLEO:2011 - Laser Applications to Photonic Applications* (Optical Society of America, 2011), p. CThJ3.

- [24] H. Shi and M. Bhattacharya, “Quantum mechanical study of a generic quadratically coupled optomechanical system,” *Phys. Rev. A* **87**, 043829 (2013).
- [25] S. Weis, R. Riviere, S. Deleglise, E. Gavartin, O. Arcizet, A. Schliesser, and T. J. Kippenberg, “Optomechanically Induced Transparency,” *Science* **330**, 1520–1523 (2010).
- [26] A. H. Safavi-Naeini, T. P. M. Alegre, J. Chan, M. Eichenfield, M. Winger, Q. Lin, J. T. Hill, D. E. Chang, and O. Painter, “Electromagnetically induced transparency and slow light with optomechanics,” *Nature* **472**, 69–73 (2011).
- [27] I. Wilson-Rae, N. Nooshi, W. Zwerger, and T. J. Kippenberg, “Theory of Ground State Cooling of a Mechanical Oscillator Using Dynamical Backaction,” *Phys. Rev. Lett.* **99** (2007).
- [28] M. Bhattacharya and P. Meystre, “Trapping and Cooling a Mirror to Its Quantum Mechanical Ground State,” *Phys. Rev. Lett.* **99**, 073601 (2007).
- [29] 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* **77** (2008).
- [30] J. Chan, T. P. M. Alegre, A. H. Safavi-Naeini, J. T. Hill, A. Krause, S. Groeblacher, M. Aspelmeyer, and O. Painter, “Laser cooling of a nanomechanical oscillator into its quantum ground state,” *Nature* **478**, 89–92 (2011).
- [31] A. H. Safavi-Naeini and O. Painter, “Proposal for an optomechanical traveling wave phonon-photon translator,” *New J. Phys.* **13**, 013017 (2011).
- [32] L. Tian and H. Wang, “Optical wavelength conversion of quantum states with optomechanics,” *Phys. Rev. A* **82**, 053806 (2010).
- [33] J. T. Hill, A. H. Safavi-Naeini, J. Chan, and O. Painter, “Coherent optical wavelength conversion via cavity optomechanics,” *Nat. Commun.* **3**, 1196 (2012).
- [34] F. Marquardt, J. P. Chen, A. A. Clerk, and S. M. Girvin, “Quantum Theory of Cavity-Assisted Sideband Cooling of Mechanical Motion,” *Phys. Rev. Lett.* **99**, 093902 (2007).
- [35] COMSOL, “COMSOL Multiphysics.”  
<http://www.comsol.com/products/multiphysics/>
- [36] J. Chan, M. Eichenfield, R. Camacho, and O. Painter, “Optical and mechanical design of a ‘zipper’ photonic crystal optomechanical cavity,” *Opt. Express* **17**, 3802–3817 (2009).

- [37] 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* **17**, 20078–20098 (2009).
- [38] A. H. Safavi-Naeini and O. Painter, “Design of optomechanical cavities and waveguides on a simultaneous bandgap phononic-photonic crystal slab,” *Opt. Express* **18**, 14926–14943 (2010).
- [39] 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.* **101**, 081115 (2012).
- [40] S. G. Johnson and J. D. Joannopoulos, “MIT Photonic-Bands (MPB).”  
[http://ab-initio.mit.edu/wiki/index.php/MIT\\_Photonic\\_Bands](http://ab-initio.mit.edu/wiki/index.php/MIT_Photonic_Bands)
- [41] S. G. Johnson and J. D. Joannopoulos, “Block-iterative frequency-domain methods for Maxwell’s equations in a planewave basis,” *Opt. Express* **8**, 173–190 (2001).
- [42] S. G. Johnson and J. D. Joannopoulos, “MEEP.”  
<http://ab-initio.mit.edu/wiki/index.php/Meep>
- [43] A. F. Oskooi, D. Roundy, M. Ibanescu, P. Bermel, J. D. Joannopoulos, and S. G. Johnson, “MEEP: A flexible free-software package for electromagnetic simulations by the FDTD method,” *Comput. Phys. Commun.* **181**, 687–702 (2010).
- [44] I. Lumerical Solutions, “FDTD Solutions.”  
<http://www.lumerical.com/tcad-products/fdtd/>
- [45] M. Eichenfield, J. Chan, R. M. Camacho, K. J. Vahala, and O. Painter, “Optomechanical crystals,” *Nature* **462**, 78–82 (2009).
- [46] M. Eichenfield, R. Camacho, J. Chan, K. J. Vahala, and O. Painter, “A picogram- and nanometre-scale photonic-crystal optomechanical cavity,” *Nature* **459**, 550–555 (2009).
- [47] GenISys, “Layour BEAMER.”  
<http://www.genisys-gmbh.com/web/products/beamer.html>
- [48] L. Zeon Chemicals, “ZEP520A.”  
<http://www.zeonchemicals.com/ElectronicMaterials/>
- [49] M. D. Henry, *ICP etching of silicon for micro and nanoscale devices*, Ph.D. thesis, California Institute of Technology (2010).
- [50] H. V. Jansen, M. J. de Boer, S. Unnikrishnan, M. C. Louwerse, and M. C. Elwenspoek, “Black silicon method: X. A review on high speed and selective plasma etching of silicon with profile

control: an in-depth comparison between Bosch and cryostat DRIE processes as a roadmap to next generation equipment,” *J. Micromech. Microeng.* **19**, 033001 (2009).

- [51] M. D. Henry, S. Walavalkar, A. Homyk, and A. Scherer, “Alumina etch masks for fabrication of high-aspect-ratio silicon micropillars and nanopillars,” *Nanotechnology* **20**, 255305 (2009).
- [52] C. Welch, A. Goodyear, T. Wahlbrink, M. Lemme, and T. Mollenhauer, “Silicon etch process options for micro- and nanotechnology using inductively coupled plasmas,” *Microelectron. Eng.* **83**, 1170–1173 (2006).
- [53] A. G. Krause, M. Winger, T. D. Blasius, Q. Lin, and O. Painter, “A high-resolution microchip optomechanical accelerometer,” *Nature Photon.* **6**, 768–772 (2012).
- [54] J. D. Cohen, S. Meenehan, and O. Painter, “Optical coupling to nanoscale optomechanical cavities for near quantum-limited motion transduction,” *Opt. Express* **21**, 11227–11236 (2013).
- [55] S. S. Walavalkar, A. P. Homyk, M. D. Henry, and A. Scherer, “Three-dimensional etching of silicon for the fabrication of low-dimensional and suspended devices,” *Nanoscale* **5**, 927–931 (2013).
- [56] M. Borselli, T. J. Johnson, and O. Painter, “Measuring the role of surface chemistry in silicon microphotronics,” *Appl. Phys. Lett.* **88**, 131114–131114-3 (2006).
- [57] H. J. Kimble, “The quantum internet,” *Nature* **453**, 1023–1030 (2008).
- [58] D. Kielpinski, J. F. Corney, and H. M. Wiseman, “Quantum Optical Waveform Conversion,” *Phys. Rev. Lett.* **106**, 130501–130501 (2011).
- [59] S. Ritter, C. Nolleke, C. Hahn, A. Reiserer, A. Neuzner, M. Uphoff, M. Mücke, E. Figueroa, J. Bochmann, and G. Rempe, “An elementary quantum network of single atoms in optical cavities,” *Nature* **484**, 195–200 (2012).
- [60] M. Wallquist, K. Hammerer, P. Rabl, M. Lukin, and P. Zoller, “Hybrid quantum devices and quantum engineering,” *Phys. Scr.* **2009**, 014001 (2009).
- [61] H. de Riedmatten, M. Afzelius, M. U. Staudt, C. Simon, and N. Gisin, “A solid-state light-matter interface at the single-photon level,” *Nature* **456**, 773–777 (2008).
- [62] K. Stannigel, P. Rabl, A. S. Sørensen, P. Zoller, and M. D. Lukin, “Optomechanical Transducers for Long-Distance Quantum Communication,” *Phys. Rev. Lett.* **105**, 220501 (2010).
- [63] J. Huang and P. Kumar, “Observation of Quantum Frequency Conversion,” *Phys. Rev. Lett.* **68**, 2153–2157 (1992).



- [64] J. Toulouse, “Optical Nonlinearities in Fibers: Review, Recent, Examples, and Systems Applications,” *J. Lightwave Technol.* **23**, 3625–3641 (2005).
- [65] S. Tanzilli, W. Tittel, M. Halder, O. Alibart, P. Baldi, N. Gisin, and H. Zbinden, “A photonic quantum information interface,” *Nature* **437**, 116–120 (2005).
- [66] M. T. Rakher, L. Ma, O. Slattery, X. Tang, and K. Srinivasan, “Quantum transduction of telecommunications-band single photons from a quantum dot by frequency upconversion,” *Nature Photon.* **4**, 786–791 (2010).
- [67] M. T. Rakher, L. Ma, M. Davanco, O. Slattery, X. Tang, and K. Srinivasan, “Simultaneous Wavelength Translation and Amplitude Modulation of Single Photons from a Quantum Dot,” *Phys. Rev. Lett.* **107**, 083602 (2011).
- [68] H. J. McGuinness, M. G. Raymer, C. J. McKinstrie, and S. Radic, “Quantum Frequency Translation of Single-Photon States in a Photonic Crystal Fiber,” *Phys. Rev. Lett.* **105**, 093604 (2010).
- [69] M. S. Kang, A. Nazarkin, A. Brenn, and P. S. J. Russell, “Tightly trapped acoustic phonons in photonic crystal fibres as highly nonlinear artificial Raman oscillators,” *Nature Phys.* **5**, 276–280 (2009).
- [70] J. D. Teufel, T. Donner, D. Li, J. W. Harlow, M. S. Allman, K. Cicak, A. J. Sirois, J. D. Whittaker, K. W. Lehnert, and R. W. Simmonds, “Sideband cooling of micromechanical motion to the quantum ground state,” *Nature* **475**, 359–363 (2011).
- [71] Y. D. Wang and A. A. Clerk, “Using Interference for High Fidelity Quantum State Transfer in Optomechanics,” *Phys. Rev. Lett.* **108**, 153603 (2012).
- [72] C. A. Regal and K. W. Lehnert, “From cavity electromechanics to cavity optomechanics,” *J. Phys.: Conf. Ser.* **264**, 012025 (2011).
- [73] R. J. Schoelkopf and S. M. Girvin, “Wiring up quantum systems,” *Nature* **451**, 664–669 (2008).
- [74] R. F. Begley, A. B. Harvey, and R. L. Byer, “Coherent anti-Stokes Raman spectroscopy,” *Appl. Phys. Lett.* **25**, 387–390 (1974).
- [75] M. Aspelmeyer, S. Gröblacher, K. Hammerer, and N. Kiesel, “Quantum optomechanics – throwing a glance,” *J. Opt. Soc. Am. B* **27**, A189–A197 (2010).
- [76] L. Tian, “Adiabatic State Conversion and Pulse Transmission in Optomechanical Systems,” *Phys. Rev. Lett.* **108**, 153604 (2012).

- [77] MathWorks, “MATLAB 2010b,” (2010).  
<http://www.mathworks.com/products/matlab/>
- [78] 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.* **108**, 033602 (2012).
- [79] A. H. Safavi-Naeini, J. Chan, J. T. Hill, S. Grblacher, H. Miao, Y. Chen, M. Aspelmeyer, and O. Painter, “Laser noise in cavity-optomechanical cooling and thermometry,” *New J. Phys.* **15** (2013).
- [80] A. H. Safavi-Naeini, S. Groeblacher, J. T. Hill, J. Chan, M. Aspelmeyer, and O. Painter, “Squeezing of light via reflection from a silicon micromechanical resonator,” *arXiv:1302.6179* (2013).
- [81] 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* **15**, 4745–4752 (2007).
- [82] A. attocube systems, “Attocube.”  
<http://www.attocube.com/nanoPOSITIONING/linear.html>
- [83] B. B. Bakir, A. V. de Gyves, R. Orobtcouk, P. Lyan, C. Porzier, A. Roman, and J.-M. Fedeli, “Low-Loss (<1 dB) and Polarization-Insensitive Edge Fiber Couplers Fabricated on 200-mm Silicon-on-Insulator Wafers,” *IEEE Photonic Tech. L.* **22**, 739–741 (2010).
- [84] M. Notomi, E. Kuramochi, and T. Tanabe, “Large-scale arrays of ultrahigh-Q coupled nanocavities,” *Nature Photon.* **2**, 741–747 (2008).
- [85] 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* **464**, 697–703 (2010).
- [86] L. Childress, M. V. G. Dutt, J. M. Taylor, A. S. Zibrov, F. Jelezko, J. Wrachtrup, P. R. Hemmer, and M. D. Lukin, “Coherent Dynamics of Coupled Electron and Nuclear Spin Qubits in Diamond,” *Science* **314**, 281–285 (2006).
- [87] M. W. McCutcheon, D. E. Chang, Y. Zhang, M. D. Lukin, and M. Loncar, “Broadband frequency conversion and shaping of single photons emitted from a nonlinear cavity,” *Opt. Express* **17**, 22689–22703 (2009).

- [88] A. Faraon, P. E. Barclay, C. Santori, K.-M. C. Fu, and R. G. Beausoleil, “Resonant enhancement of the zero-phonon emission from a colour centre in a diamond cavity,” *Nature Photon.* **5**, 301–305 (2011).
- [89] W. Bogaerts, D. Taillaert, B. Luyssaert, P. Dumon, J. Van Campenhout, P. Bienstman, D. Van Thourhout, R. Baets, V. Wiaux, and S. Beckx, “Basic structures for photonic integrated circuits in Silicon-on-insulator,” *Opt. Express* **12**, 1583–1591 (2004).
- [90] G. Roelkens, P. Dumon, W. Bogaerts, D. Van Thourhout, and R. Baets, “Efficient silicon-on-insulator fiber coupler fabricated using 248-nm-deep UV lithography,” *IEEE Photonic Tech. L.* **17**, 2613–2615 (2005).
- [91] V. Nguyen, T. Montalbo, C. Manolatou, A. Agarwal, C.-y. Hong, J. Yasaitis, L. C. Kimerling, and J. Michel, “Silicon-based highly-efficient fiber-to-waveguide coupler for high index contrast systems,” *Appl. Phys. Lett.* **88**, 081112 (2006).
- [92] T. Alder, A. Stohr, R. Heinzelmann, and D. Jager, “High-efficiency fiber-to-chip coupling using low-loss tapered single-mode fiber,” *IEEE Photonic Tech. L.* **12**, 1016–1018 (2000).
- [93] 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* **20**, 2274–2284 (2003).
- [94] P. Barclay, K. Srinivasan, M. Borselli, and O. Painter, “Experimental demonstration of evanescent coupling from optical fibre tapers to photonic crystal waveguides,” *Electron. Lett.* **39**, 842–844 (2003).
- [95] P. E. Barclay, K. Srinivasan, M. Borselli, and O. Painter, “Efficient input and output fiber coupling to a photoniccrystal waveguide,” *Opt. Lett.* **29**, 697–699 (2004).
- [96] C. Zener, “Non-Adiabatic Crossing of Energy Levels,” *Proc. R. Soc. A* **137**, 696–702 (1932).
- [97] J. S. Cook, “Tapered velocity couplers,” *Bell Syst. Tech. J.* **34**, 807–822 (1955).
- [98] A. G. Fox, “Wave coupling by warped normal modes,” *Bell Syst. Tech. J.* **34**, 823–852 (1955).
- [99] W. H. Louisell, “Analysis of the single tapered mode coupler,” *Bell Syst. Tech. J.* **34**, 853–870 (1955).
- [100] H. Yajima, “Dielectric thinfilm optical branching waveguide,” *Appl. Phys. Lett.* **22**, 647–649 (1973).
- [101] W. Burns and A. Milton, “Mode conversion in planar-dielectric separating waveguides,” *IEEE J. Quantum Elect.* **11**, 32–39 (1975).

- [102] Y. Silberberg, P. Perlmutter, and J. E. Baran, “Digital optical switch,” *Appl. Phys. Lett.* **51**, 1230–1232 (1987).
- [103] D. G. Dalgoutte, R. B. Smith, G. Achutaramayya, and J. H. Harris, “Externally mounted fibers for integrated optics interconnections,” *Appl. Opt.* **14**, 1860–1865 (1975).
- [104] Y. Shani, C. Henry, R. Kistler, R. Kazarinov, and K. Orlowsky, “Integrated optic adiabatic devices on silicon,” *IEEE J. Quantum Elect.* **27**, 556–566 (1991).
- [105] D. Kleckner and D. Bouwmeester, “Sub-kelvin optical cooling of a micromechanical resonator,” *Nature* **444**, 75–78 (2006).
- [106] A. C. Doherty, S. Habib, K. Jacobs, H. Mabuchi, and S. M. Tan, “Quantum feedback control and classical control theory,” *Phys. Rev. A* **62** (2000).
- [107] D. W. C. Brooks, T. Botter, S. Schreppler, T. P. Purdy, N. Brahms, and D. M. Stamper-Kurn, “Non-classical light generated by quantum-noise-driven cavity optomechanics,” *Nature* pp. 476–480 (2012).