Photonic Crystal Waveguides for Integration into an Atomic Physics Experiment

Thesis by Andrew Corby McClung

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

Caltech

CALIFORNIA INSTITUTE OF TECHNOLOGY Pasadena, California

> 2017 Defended June 7, 2017

© 2017

Andrew Corby McClung ORCID: 0000-0001-6995-3289

All rights reserved.

ACKNOWLEDGEMENTS

I would first like to thank my adviser, Jeff Kimble, for his guidance and support throughout my time in his group. Jeff's passion, insight and tenacity drive our work forward, and his vision of the scientific endeavor has been inspirational. I am glad to have had the opportunity to participate in this project from both sides of the gowning room, and I am grateful to Jeff for the breadth of this experience. Thanks are also due to Oskar Painter, for generously allowing me to use his facilities, inviting me to his group meetings, and offering helpful advice with regard to both fabrication and my career.

I am grateful for the capable colleagues I have had in Lab 2. Every one of them contributed significantly to the work presented in this thesis, and to my quality of life both in the basement of East Bridge and outside of it. I had the great fortune to work closely with and learn from three postdocs during my time in Lab 2. Though the vacuum, optical and computer systems have gone through many iterations, they are at their core the result of Jae Hoon Lee's careful planning, and each element (excepting "the tumor") has served us well. Michael Martin was, in his tenure, an exceptional leader, and even after leaving Caltech has continued to be a peerless mentor. He and his wife, Carly Donahue, are each hard-working and talented scientists, and I am glad to be friends with both of them. Alex Burgers, the most recent addition to the Lab 2 crew, has been a joy to work with. His ability to maintain both concentration and levity over two grueling days of fiber coupling is enviable. Heroes are hard to find, which is why I am so glad we found Alex.

I also worked closely with two graduate students in Lab 2, Lucas Peng and Juan Muniz. I like to imagine myself as reasonably computer savvy, but Lucas has shattered that illusion, both with his intricate trajectory simulation work and his ability to interpret the LabVIEW atrocity in which I am complicit. Lucas is decidedly the most culturally in-tune member of the lab, and makes the rest of us feel cool by association. Juan is selfless, sometimes to a fault. I thank him for his generosity, his willingness to indulge my idiosyncrasies, and his patience, which I have rigorously tested. I could not have asked for a better companion with whom to share both the bitter disappointments and moments of triumph in this journey.

The are many group members to thank beyond just Lab 2. Su-Peng Yu and Jon Hood laid the groundwork for the device geometry, fabrication process, and characteriza-

tion facilities, and willingly trained me when I joined the fabrication team. In the cleanroom there is nothing more intimate than shared glassware, and I am honored by their trust. I'd also like to acknowledge the Lab 1 team—Clément Lacroûte, Pol Forn-Díaz, Martin Pototschnig, and especially Daniel Alton—for their patience and guidance in the early years of my PhD, and the Lab 11 team, particularly Akihisa Goban and Chen-Lung Hung, for enlightening interactions. To our theorist in residence, Ana Asenjo Garcia, for helping clarify for me much of the motivation for this work, her husband Héctor, for his delicious frittata, and both of them for their friendship: muchas gracias.

I am also glad to have worked alongside all the fabricators the Painter cleanroom, especially Justin Cohen, Paul Dieterle, Mahmoud Kalaee, and Greg MacCabe. Their willingness to generously share facilities and tricks, and indulge superstitions made the long hours in the yellow glow bearable. Thanks are due also to Barry Baker, who has kept things running, and received calls for help at odd hours without complaint. The KNI staff have also supported this project significantly, particularly Matthew Sullivan Hunt and Melissa Melendes, who expressed not only willingness but interest in solving arcane problems I encountered.

I would like to express appreciation for the love and support of friends and family during my time in graduate school. Close friends and fellow Caltech graduate students Kevin Fiedler, Sam Johnson, Max Jones and Kevin Shan deserve special mention. I also thank my mother and father, who have unquestioningly supported and encouraged me my entire life. And finally, I express my deepest gratitude to Hannah Goodwin, who tolerated a years-long period of weekend commutes, frequent midnight beam writes, and the horrible noisy keyboard on which this document was written, and decided to marry me anyway.

ABSTRACT

Strongly interacting systems of atoms and photons are an important resource in many active areas of research, including quantum information science, quantum simulation, and metrology. Frequently, the strength of these interactions is enhanced by using an optical resonator to confine light to a small volume. In recent years, there have been efforts to replace traditional Fabry-Pérot resonators, formed from macroscopic mirrors, with micro- and nano-fabricated systems, leveraging techniques and infrastructure from semiconductor manufacture to scalably produce high-quality, small mode volume waveguides and resonators. Of particular interest are nano-fabricated photonic crystals, in which very fine control over modal and dispersion properties is possible. Here I describe our efforts to reliably produce photonic crystal waveguides with guided modes designed to trap and interrogate an array of ultracold cesium atoms. Specifically, I present models capturing band placement, modal structure, finite photonic crystal effects, and waveguide input and output coupling; I discuss the techniques we use to fabricate our photonic crystal waveguides; and I describe our characterization capabilities and the packaging and installation of the waveguides into the atomic physics system.

PUBLISHED CONTENT AND CONTRIBUTIONS

 Goban, A., Hung, C.-L., Yu, S.-P., Hood, J., Muniz, J., Lee, J., Martin, M., McClung, A., Choi, K., Chang, D. E., Painter, O. & Kimble, H. J. Atom–light interactions in photonic crystals. *Nat. Comm.* 5, 3808 (May 2014). DOI:10.1038/ncomms4808 ACM contributed to the design and construction of the experiment.

TABLE OF CONTENTS

| Acknowledgements | iii |
|--|---|
| Abstract | · · · · · · · · · · · · · · · · · · · |
| Published Content and Contributions | · · · · · · · · · · · · · · · · · · · |
| Table of Contents | · · · · · · · · · · · · · · · · · · · |
| List of Illustrations | ix |
| List of Tables | xii |
| Chapter 1: Introduction | |
| 1.1 Outline | |
| 1.2 My History in Quantum Optics Group . | |
| Chapter 2: Design and Modeling | |
| 2.1 Waveguide Symmetries and Dispersion | |
| 2.2 Bloch Modes and Band Structure of AF | PCW |
| 2.3 Monogators | |
| 2.4 Dielectric Stack Model | |
| 2.5 Mode Overlap Calculations | |
| Chapter 3: Fabrication | |
| 3.1 Substrate Preparation | |
| 3.2 Photolithography | |
| 3.3 Electron Beam Lithography | |
| 3.4 Reactive Ion Etch | |
| 3.5 Potassium Hydroxide Etch and Wet Che | emistry 61 |
| 3.6 Atomic Layer Deposition | |
| Chapter 4: Characterization and Installation | |
| 4.1 Characterization Setup and Spectra . | |
| 4.2 Atomic Layer Deposition | |
| 4.3 Image Data | |
| 4.4 Science Chips | |
| 4.5 Packaging and Installation | 91 |
| 4.6 Alkali Corrosion | 97 |
| Chapter 5: Atomic Physics Experiment | |
| 5.1 Vacuum System and Magneto-Optical | Fraps |
| 5.2 Transporting Atoms from Science MO | $\Gamma \text{ to } APCW \qquad \qquad 103$ |
| 5.3 Computer Control and Data Acquisition | 107 |
| 5.4 Search for Phase Sensitivity | 110 |
| Chapter 6: Outlook | 113 |
| Bibliography | 114 |
| Appendix A: Timeline of Chins in Lab 2 | 123 |
| Appendix B: MPB Scripts | 123 |
| B 1 APCW Unit Cell | 124 |
| | |

| B.2 Mode Overlap Calculation | | | | | | | | • | | • | • | | 127 |
|------------------------------|--|--|--|--|--|--|--|---|--|---|---|--|-----|
| Appendix C: Film Thicknesses | | | | | | | | | | | | | 131 |
| Appendix D: Cold Develop | | | | | | | | | | | | | 132 |

LIST OF ILLUSTRATIONS

| Number | | P | lage |
|--------|---|---|------|
| 1.1 | Depiction of APCW | • | 4 |
| 2.1 | Waveguide geometry | • | 10 |
| 2.2 | APCW geometry | • | 10 |
| 2.3 | PNW mode symmetries | • | 12 |
| 2.4 | PNW band diagram | • | 12 |
| 2.5 | APCW band diagram | • | 13 |
| 2.6 | <i>z</i> -even mode profiles at $\beta_x = \pi/a$ | • | 15 |
| 2.7 | <i>z</i> -odd mode profiles at $\beta_x = \pi/a$ | • | 15 |
| 2.8 | Mode visibility and group index | • | 16 |
| 2.9 | Band gap placement: width and amplitude | • | 17 |
| 2.10 | Geometric distortions | • | 18 |
| 2.11 | Phase shift | • | 19 |
| 2.12 | Width asymmetry | • | 20 |
| 2.13 | APCW band edges vs. nanobeam separation | • | 21 |
| 2.14 | MPCW | • | 22 |
| 2.15 | Purcell factor, APCW (TE) | • | 23 |
| 2.16 | Purcell factor, MPCW (TE) | • | 24 |
| 2.17 | Purcell factor, MPCW (TM) | • | 25 |
| 2.18 | MPCW external illumination | • | 26 |
| 2.19 | TMM conventions | • | 27 |
| 2.20 | Group index calculations using TMM | • | 29 |
| 2.21 | Taper length: reflectivity and group index | • | 31 |
| 2.22 | Supermodes (TMM) | • | 32 |
| 2.23 | Long-range band gap deformation and losses | • | 34 |
| 2.24 | Comparison of TMM to MPB | • | 35 |
| 2.25 | Coupler mode profiles | • | 37 |
| 2.26 | Mode overlap: geometry and frequency | • | 38 |
| 2.27 | Mode overlap: displacement | • | 38 |
| 3.1 | Fabrication overview | • | 39 |
| 3.2 | Illustration of UV mask | • | 41 |
| 3.3 | Software flow for <i>e</i> -beam patterns | | 43 |

| 3.4 | .param dimensions 47 |
|------|---|
| 3.5 | Deflection in EBPG5000+ |
| 3.6 | Errors introduced by fracturing |
| 3.7 | Loading chips in EBPG |
| 3.8 | Determining substrate rotation |
| 3.9 | Dose-to-clear |
| 3.10 | Electron beam PSF |
| 3.11 | Process blur for 1:1 line–space pattern (simulated) |
| 3.12 | Line–space dose array |
| 3.13 | Simulated process blur |
| 3.14 | Measured process blur |
| 3.15 | APCW dose array |
| 3.16 | Magic factors |
| 3.17 | Etch test SEM |
| 3.18 | ARDE in crystal |
| 3.19 | KOH etch simulation |
| 3.20 | Siderail failure |
| 3.21 | Importance of BHF |
| 3.22 | Illustration of ALD |
| 4.1 | Diagram of polarization control setup |
| 4.2 | Spectral polarization variation of SLED |
| 4.3 | Characterization setup |
| 4.4 | OSA spectra |
| 4.5 | PSD of reflection and transmission spectra |
| 4.6 | TM and TE transmission spectra |
| 4.7 | Scatterer coupling to other symmetries |
| 4.8 | TE coupling efficiency vs. coupler width |
| 4.9 | Simulated and measured band placement |
| 4.10 | Comparing nominally identical waveguides |
| 4.11 | APCW parameter scan |
| 4.12 | Comparing nominally identical prints |
| 4.13 | ALD shift |
| 4.14 | Intensity profiles from image data |
| 4.15 | SEM settings |
| 4.16 | Geometric uniformity across crystal |
| 4.17 | Reflection spectra for Hilde B, dev. 1 |
| | |

Х

| 4.18 | Lil F reflection spectra |
|------|---|
| 4.19 | Lil H transmission spectra |
| 4.20 | Dil D transmission spectra |
| 4.21 | Harald C transmission spectra |
| 4.22 | Hilde E reflection spectra |
| 4.23 | Gluing |
| 4.24 | Signs of cesium exposure |
| 4.25 | In situ cesium shift |
| 4.26 | Simulating adlayers |
| 5.1 | Diagram of vacuum system |
| 5.2 | Mini-MOT kick |
| 5.3 | Enhanced decay rate in mini-MOT |
| 5.4 | Optical lattice simulation |
| 5.5 | Illustration of Lab 2 vacuum chamber |
| 5.6 | Sequence composer GUI |
| 5.7 | GM absorption spectrum |
| 5.8 | Cross-correlation signal |
| A.1 | Timeline of chips in Lab 2 vacuum chamber |
| C.1 | Silicon nitride film color |
| D.1 | Cold develop test |

xi

LIST OF TABLES

| Numbe | Pr Pa | ge |
|-------|--|----|
| 2.1 | Tuning rates of band gap with geometry (simulated) | 18 |
| 3.1 | CAD parameters, part 1 | 45 |
| 3.2 | CAD parameters, part 2 | 46 |
| 3.3 | Magic factors | 59 |
| 3.4 | RIE parameters | 50 |
| 3.5 | Empirically determined DRIE etch rates | 51 |
| 3.6 | KOH etch rates | 52 |
| 4.1 | Tuning rates of band gap with geometry (empirical) | 74 |
| 4.2 | Science chips | 39 |

INTRODUCTION

Laser-cooled and trapped neutral atoms form the basis of many modern physics experiments, admitting the study of a wide variety of phenomena. Atom interferometers and optical clocks allow the measurement of time and fundamental constants with unprecedented precision [1, 2]. Ultracold quantum gases enable controllable studies of interacting quantum many-body systems [3]. High-*Q* cavities are used to controllably exchange quantum information between atomic and photonic degrees of freedom, constituting nodes in a quantum network [4].

Their broad applicability results from the fact that ultracold atoms are "close to ideal realizations of isolated quantum systems" [5]—ultra-high vacuum keeps atoms well-isolated from the environment, and ultra-low temperature virtually eliminates inhomogeneity, allowing them to be treated as identical. In the lab, their primary control interface is near-resonant light, which is used to trap and cool atoms, and can also be used to manipulate and read out an atom's electronic degrees of freedom. The efficiency of this interaction is largely irrelevant in laser trapping and cooling—photons are cheap—but is crucial for implementing quantum information protocols [4].

In free space, the strength of of atom-photon interaction is limited, even for tightlyfocused laser beams [6]. Interaction strength can be significantly enhanced by modifying the local density of optical states (LDOS) [7]. A very successful strategy [8] for increasing the LDOS is to place an atom in a resonant cavity, as first suggested by Purcell [9]. The resulting enhancement can be captured by a Purcell factor, defined here as the ratio of the decay rate into the resonator mode to the decay rate into all other modes. For an atom in a cavity on resonance, the Purcell factor P_c is given by

$$P_{\rm c} \equiv \frac{\Gamma_{\rm c}}{\Gamma'} \approx \frac{3Q}{4\pi^2} \left(\frac{\lambda_0^3}{V_{\rm m}}\right),\tag{1.1}$$

where Γ_c is the decay rate into the cavity mode, Γ' is the decay rate into all other modes, Q is the quality factor of the resonator, λ_0 is the free-space wavelength of the atomic transition and V_m is the volume of the cavity mode. It is apparent from this expression that larger Purcell factors are obtained in higher-Q cavities, and those with more tightly confined modes (smaller $V_{\rm m}$).

Cavities are not the only class of structures which can significantly enhance the LDOS. In the late 1980s it was suggested by Yablonovitch [10] and John [11] that media with spatially periodic dielectric constant could be used to modify the LDOS. These media are called *photonic crystals*, as photons in them behave analogously to electrons in a solid-state crystal [12]. Interesting optical properties only occur when the spatial period is on the same order as excitation wavelength, which for visible or near-infrared operation means patterns on the hundreds-of-nanometers scale. An analytical expression for the Purcell factor P_{wg} can be determined for an atom interacting with a one-dimensional photonic crystal, called a photonic crystal waveguide (PCW) [13]:

$$P_{\rm wg} \equiv \frac{\Gamma_{\rm 1D}}{\Gamma'} \approx \frac{3n_{\rm g}}{4\pi} \left(\frac{\lambda_0^2}{A_{\rm m}}\right),\tag{1.2}$$

where Γ_{1D} is the decay rate into the waveguide, $n_g \equiv c \left(\frac{\partial \omega}{\partial \beta}\right)^{-1}$ is the group index, β is the wave vector, $A_m \equiv V_m/a$ is the mode area, and *a* is the lattice constant (spatial period) of the PCW. Again, tight modal confinement boosts the Purcell factor. In the structures discussed in this thesis, $\Gamma' \approx \Gamma_0$, the decay rate in free space. A waveguide does not rely on resonant enhancement; instead, the LDOS is modified by tailoring dispersion to increase n_g .

The defining feature of photonic crystals is the existence of photonic band gaps: spectral regions for which no guided modes exist. Modes can be localized at defects in a photonic crystal: point defects create photonic crystal cavities and line defects photonic crystal waveguides. At frequencies near but outside the band gap, guided mode dispersion is modified, resulting in large group index, i.e., "slow light." While complete band gaps exist only in 3D, 1D and 2D photonic crystals can also significantly modify the LDOS [14].

The form of equations 1.1 and 1.2 have apparent similarities. The connection between Q and n_g can be made explicit by considering the group index in a linear resonator: in a high-Q Fabry–Pérot cavity of length L, the phase in transmission ϕ increases by π across the cavity linewidth, giving

$$n_{\rm g} \approx c \left(\frac{2\kappa}{\pi/L}\right)^{-1} = \frac{Q\lambda_0}{2L},$$

where we have used $\phi = \beta L$, $Q = \omega_0/2\kappa$, where ω_0 is the angular resonance frequency, κ the half width of the cavity at half maximum, and $\lambda_0 = 2\pi c/\omega_0$. In this context n_g is often identified as the finesse.

While Fabry–Pérot cavities constructed of macroscopic substrate mirrors have been a very successful platform [15–17], micro- and nano-scale resonators and waveguides are attractive, due to their inherently small mode volumes. In recent years ultracold neutral atoms have been integrated successfully with microscopic whispering gallery mode resonators [18–20] and optical nanofiber waveguides [21–23].

Lithographically-defined photonic crystal structures, particularly those manufactured in CMOS-compatible materials, benefit enormously from the mature techniques for and infrastructural investment in producing silicon integrated circuits [24], promising expanded configurability, scalable manufacture, and the ability to engineer dispersion and modal properties. Such structures are already in wide use with solid-state emitters, e.g., quantum dots [25–27] and color centers in crystals [28, 29].

Platforms to integrate photonic crystals with neutral atoms have been recently demonstrated, using photonic crystal cavities [30], and, in our group at Caltech, photonic crystal waveguides [31–33]. Figure 1.1 illustrates the 1D PCW used in our experiment: it consists of two parallel silicon nitride nanobeams with external corrugation. This corrugated geometry is referred to as the 'alligator' photonic crystal waveguide (APCW), and is designed to create stable guided mode optical dipole traps commensurate with a probe mode aligned to the D-line transitions in cesium [34].

Beyond simply creating an improved platform for waveguide-QED [21–23, 35], photonic crystals designed for use with neutral atoms promise to enable the study of qualitatively new regimes of interacting quantum many-body systems. Particularly, light-matter interaction for atoms with transition frequencies inside a band gap exhibit interesting phenomena. As photons cannot propagate in the surrounding medium, an atom behaves like dielectric "defect" around which photons can localize, creating an atom–photon bound state [36]. The length-scale of this localization is determined by the relative detuning of the atomic transition and the band edge. This effect could be leveraged to create tunable-range, photon-mediated atom–atom interactions [37].

Photonic crystals also enable independent control of optical frequency and spa-



Figure 1.1: **Depiction of APCW. (a)** Rendering of atoms (green) loaded in an 'alligator' photonic crystal waveguide (gray) interacting with a guided mode (red). **(b)** SEM micrograph of APCW. Scale bar is $1 \mu m$.

tial periodicity. In our experiment this allows us to make light of very different frequencies spatially commensurate, which aids the construction of guided-mode optical traps [34]. More generally, nano-structured dielectrics could be used to create sub-wavelength optical lattices, opening new regimes for 2D many-body simulations [38].

1.1 Outline

There are many prerequisites to the study of atoms interacting photons in a photonic crystal, the most fundamental of which is the engineering and manufacture of photonic crystal structures with properties that allow atoms to be trapped and interrogated by guided modes. In this thesis I discuss the modeling, fabrication, and characterization of the 'alligator' photonic crystal waveguide, and also some details of its integration into an atomic physics experiment. In chapter 2, I describe important design elements, including dispersion and mode profiles, and present models capturing geometrical effects on band placement; features of finite photonic crystals including tapering, long length-scale distortions and supermode structure; and mode overlap calculations informing coupler geometry. In chapter 3, I outline our fabrication procedure, including lithography, dry and wet etching, and atomic layer deposition, as well as some useful techniques for process characterization. Finally, in chapter 4, I describe the apparatus to characterize our devices; present empirical spectral features, tuning and coupling rates; describe the process to package chips for the vacuum chamber; and discuss waveguide degradation resulting from alkali exposure.

1.2 My History in Quantum Optics Group

It is my hope that this thesis can be used as a guide for future efforts to fabricate nanophotonic structures for use in cold atom experiments. While a chronological account does not elucidate the shortest route to success in this endeavor, I provide it here to contextualize the studies performed and recommendations made later in the document.

I arrived at Caltech in September 2011, and joined Daniel Alton, a senior graduate student, and postdocs Pol Forn-Díaz and Clément Lacroûte in Lab 1. They were working on a cavity QED experiment using microtoroidal whispering gallery mode (WGM) resonators fabricated by the Vahala group. Before my arrival, experiments in the Quantum Optics Group had already demonstrated strong coupling [18], photon routing conditioned on the presence of a single atom [39, 40], and signatures of Casimir-Polder forces between the atom and dielectric resonator [41] for atoms in transit. However, trapped atoms are necessary for implementing more sophisticated quantum information protocols, and the goal when I arrived was to use a nanofiber, already present in the system to couple light into and out of the microtoroid, to create an optical potential to trap atoms in a region with good coupling with the resonator mode.

That fall the Lab 1 team was in the process of constructing 686 nm and 937 nm laser systems to duplicate the "magic" wavelength nanofiber trap demonstrated in Lab 2 [22]. Subsequent numerical simulations showed that simultaneity of stable atom traps, strong atom–photon coupling and good fiber–toroid coupling was difficult [42], and this effort was abandoned in October 2012. Still, my first year with the group was formative: I familiarized myself with trapping and cooling techniques for neutral atoms, the home-built diode laser and tapered amplifier systems, and computer control systems used in the experiment.

The summer and fall of 2012 witnessed a paradigm shift in the Kimble group: the cavity QED system in Lab 11 was disassembled to make room for the new photonic crystal experiment; Lab 11a, previously an optomechanics experiment, was reconfigured into a photonic crystal characterization lab. Soon after that, the new vacuum system being built in Lab 2, originally conceived as an updated nanofiber setup, was converted to a photonic crystal experiment. That October the Lab 1 team followed suit, modifying our apparatus to conform to the proven atom delivery system pioneered by the Lab 11 team.

In spring of 2013, Daniel retreated to write his thesis, and the remaining Lab 1

postdocs left. At this time teams were reconfigured, and I was invited to join graduate student Juan Muniz and postdoc Jae Hoon Lee in Lab 2. Around this time postdoc Mike Martin arrived, who after a short tenure in the characterization lab also joined Lab 2. I began working on modernizing the computer control systems in Lab 2 to simplify the creation of sophisticated experimental sequences, adapting MATLAB and LabVIEW code Jae brought with him from Arizona, and also contributed to the image and photon counting data acquisition systems. Early guided-mode absorption spectra in Lab 2 were seen in May of 2013, and by September the Lab 11 team had seen significant decay rate enhancement for localized (but untrapped) atoms, reported in ref. [31].

These initial demonstrations were encouraging, and each lab proceeded to attempt to load atoms into GM traps. We had initially hoped that these traps could be loaded using techniques established by the preceding nanofiber experiments, but unfortunately loading traps in a photonic crystal is not as straightforward. At this time Labs 2 and 11 developed different techniques to increase the density of atoms near the photonic crystal. The Lab 11 technique uses a single incident beam and its reflection off the photonic crystal to create optical potentials near the waveguide, similar to techniques described in ref. [30].

Lab 2 developed a different approach based on the interference of two counterpropagating beams in an effort to position atoms in between the nanobeams. Detuning the frequency of one of these beams results in a slowly moving optical lattice potential, which, under the appropriate conditions, can position atoms in the target region between the nanobeams. This relative detuning is achieved using acoustooptical modulators driven by phase-stable radio-frequency sources. In summer 2014 I adapted Arduino code Mike brought with him from CU to generate, via direct digital synthesis, the time-varying frequencies necessary to controllably move atoms in the optical lattice.

Up to this point the Lab 2 team had been working exclusively on the atomic physics experiment, while Lab 11 had two graduate students, Su-Peng Yu and Jon Hood, working on device fabrication. The Lab 2 team decided it was important to have some skin in the game, and I asked to join the fabrication effort in August 2014. That fall I trained with Jon and Su-Peng. We worked together to troubleshoot newly-observed contaminants deposited in the wet chemistry process using linear resonators, and Jon and I explored post-fabrication dry etching techniques to obtain finer control over the photonic crystal band placement. My vacancy in Lab 2 was

filled by a younger graduate student, Lucas Peng.

By April 2015 I completed my fabrication training and became relatively independent. That spring we investigated using atomic layer deposition (ALD) to chemically passivate and tune band structure of our APCWs, starting out on an ALD tool in the Atwater group, and in June using a system installed in the Painter group cleanroom. By that summer we had gone through all the wafers from our first silicon nitride deposition run, and switched to a second batch. Unfortunately, the dice angle of this second wafer run had slight rotation with respect to the crystal planes in silicon, resulting in poor device yield and the introduction of additional alignment steps in the electron beam lithography. It was around this time that the Lab 2 team first began to see large absorption signals using lattice transport, and also device degradation resulting from cesium exposure. As the usable device lifetimes in Lab 2 significantly shortened, I worked on fabrication process characterization, in the hopes of both increasing the yield of usable chips and narrowing the spectral distribution of fabrication output. Postdoc Alex Burgers arrived from Michigan in August, overlapping briefly with Mike, who left in the fall. That fall and winter I worked on dose studies aimed at improving pattern fidelity, ultimately choosing to reduce the dose at which patterns were exposed.

In January of 2016 the chip in Lab 2 was replaced. The device we removed allowed us to characterize the geometrical effects of cesium corrosion using SEM images. That spring was a difficult time for fabrication—a lot of the fabricated devices exhibited unexpected features in the band gap. Prof. Oskar Painter suggested asymmetry in the transition region between unmodulated waveguide and nominal crystal could lead to coupling to unintended mode symmetries, potentially explaining these features. To address this possibility I made significant modifications to the fracturing to enforce symmetries as best I could. After some difficulties butt-coupling in June, Lab 2 successfully installed a chip in mid-July—the first install of a chip I fabricated.

Soon afterward the Lab 2 team saw the first hints of a absorption signal in phase with the lattice transport, which is a signature of atoms transiting through the the target region between the nanobeams. In August and September, anticipating another chip transfer, I wrote a very large number of chips, and worked to develop a butt-coupling procedure that avoided the difficulties witnessed earlier that summer. Finally, in late January 2017, we installed another "science" chip, this time with eight coupled devices (previously the maximum number of coupled devices was four). Quite recently we have seen very encouraging phase-dependent signals with signatures

similar to those produced by simulation.

Chapter 2

DESIGN AND MODELING

The chips used in our experiment consist of many optical and structural elements necessary to facilitate the delivery, trapping and coupling of atoms to guided modes (GMs). Our fabrication process (see ch. 3) allows us to create arbitrary twodimensional patterns, but all elements necessarily have a thickness equal to the silicon nitride film thickness, ca. 200 nm. An overview of our waveguide geometry is depicted in fig. 2.1. Light is coupled into and out of the waveguide by placing an optical fiber in a silicon vee groove near a small (ca. 150×200 nm) rectangular waveguide coupler. The coupler is held in place with a 200 nm-wide tether. The width of this waveguide is adiabatically increased to a nominal width of 500 nm over a distance of 200 µm, improving confinement of the mode. This technique is referred to as butt coupling, and is described briefly in sec. 2.5, and in detail in refs. [43–45].

Structures called node arrays mechanically support the waveguide and provide thermal anchoring to side rails on both sides of the waveguide. The first node array is angled to maintain tension in the waveguide and reduce the force on the (comparatively delicate) coupler tether. After approximately 500 μ m of propagation, the waveguide enters the window, a 1.8×6 mm region under which the silicon substrate has been completely removed. The size of the window is sufficient to pass free-space beams through the region near the photonic crystal, necessary for laser trapping and cooling schemes for atoms. Near the center of the window the single-nanobeam waveguide is split in two at the *y* junction, and the exterior sidewall of each nanobeam is sinusoidally modulated. The amplitude of this modulation is increased over ca. 6 μ m (15 unit cells) in the photonic crystal taper, and then held constant for ca. 60 μ m (150 unit cells) in the nominal crystal. All elements are repeated in the opposite order, terminating in a second coupler ca. 2.6 mm away from the first. Many waveguides can be printed on a single chip; the standard number is 16.

Models of critical optical elements in this structure are necessary to understand and optimize their behavior. In this chapter I describe models that fall into two categories: band structure and eigenmode simulations, here performed using MPB [46];



Figure 2.1: Waveguide geometry. (a) Layout of chip (to scale). Substrate is $10 \times 10 \times 0.2$ mm, with a 2×6 mm window over which 16 waveguides are suspended. Fibers (not pictured) are aligned to the waveguides in grooves left and right of the window. (b) Detailed view of waveguide (*not* to scale). $L \approx 2.6$ cm. (i) Light is coupled into and out of the waveguide by butt-coupling (sec. 2.5). Light propagates several millimeters in a single-nanobeam waveguide suspended by (ii) node arrays. Near the center of the window the waveguide splits in two, and light enters the (iii) impedance-matching taper. In the center of the window is (iv) the nominal crystal.



Figure 2.2: **APCW geometry.** Waveguide consists of two parallel nanobeams made of silicon nitride (gray) separated by a gap g. (a) xy-plane. Each nanobeam has an average width w, and is sinusoidally modulated with amplitude A and period a. (b) yz-plane. The transverse cross-section at each x value consists of parallel rectangular waveguides with thickness t. GMs propagate in the x direction.

and studies of termination and finite length effects in the photonic crystal, here treated using a transfer matrix model (TMM). More details about the design and performance of structures not discussed here may be found in refs. [47, 48].

2.1 Waveguide Symmetries and Dispersion

The central element in our waveguide is the APCW, a 1D photonic crystal waveguide consisting of two parallel silicon nitride nanobeams with corrugated exterior sidewalls. Parameterization conventions used in this thesis are presented in fig. 2.2. Before discussing the properties of the APCW, it is informative to consider modal properties of an uncorrugated (i.e., A=0 nm) parallel nanobeam waveguide (PNW).

Waveguides typically support many GMs, which are categorized by their symmetry properties. In a plane perpendicular to the *x* direction, the dielectric function of a PNW has mirror symmetry with respect to the *y* and *z* axes. This symmetry allows us to categorize eigenmodes into even and odd types: if in the *yz*-plane the electric field vector transforms as $\mathbf{E}(x_0, y, z) = -\mathbf{E}(x_0, -y, z)$, i.e., the *y* axis behaves like a mirror for the electric field vector, the mode is said to have even *y* symmetry, whereas if $\mathbf{E}(x_0, y, z) = \mathbf{E}(x_0, -y, z)$, the mode is said to have odd *y* symmetry, with identical categorizations for *z*. These symmetries are visualized in fig. 2.3. The modes with the simplest polarization are those with opposite *y* and *z* symmetries, i.e., *y*-even-*z*-odd and *y*-odd-*z*-even symmetries, which we will categorize as TM-like (vertical polarization) and TE-like (horizontal polarization), respectively. Even-even and odd-odd modes have more complicated field patterns. A more general discussion about the consequences of these symmetries can be found in chs. 3 and 7 of ref. [49].

Figure 2.4 shows the dispersion diagram for the fundamental mode of each of these symmetries. The black region is called the light cone, and comprises the continuum of unguided modes [49]. Due to its finite transverse dimensions, the PNW exhibits waveguide dispersion: at low frequencies the evanescent field for each mode is large, and the effective index $n_{\text{eff}} \equiv c\beta/\omega$ is small; in the limit of high frequency, the modes are well-confined to the waveguide and n_{eff} approaches the bulk refractive index of the waveguide. In a small spectral region sufficiently far from these extremes, the dispersion relation for a given mode can be reasonably approximated as

$$n(\omega)\omega = c\beta_x,$$

where $n(\omega) = n^{(0)} + \frac{n^{(1)}}{\omega},$ (2.1)

where $n(\omega)$ is the effective index, $n^{(0)}$ and $n^{(1)}$ are constants, ω is the angular frequency, *c* is the speed of light and β_x is the longitudinal component of the propagation vector. The dispersion model is linear, but accommodates a non-zero frequency at $\beta_x = 0$. We will see later that waveguide dispersion partially accounts



Figure 2.3: **PNW mode symmetries.** Arrows indicate direction and magnitude of electric field vector. The top (bottom) row has even (odd) *y* symmetry, and the left (right) column has even (odd) *z* symmetry. Quasi-TE and -TM modes are indicated. Calculated at $\beta_x = 8.55 \text{ rad/}\mu\text{m}$, w = 290 nm, A = 0 nm, g = 220 nm, t = 200 nm.



Figure 2.4: **PNW band diagram.** Simulated with dimensions w = 290 nm, t = 200 nm, g = 220 nm.



Figure 2.5: **APCW band diagram.** Blue bands have even z symmetry, red bands have odd z symmetry. Solid lines represent the modes which can be directly excited in the experiment. Shaded gray region represents the partial gap for the y-odd-z-even (TE) mode, f_0 refers to the center frequency of this gap and Δf the frequency width. (*t*=200 nm, *w*=288 nm, *A*=140 nm, *g*=220 nm.)

for the discrepancy between dispersion relations given by a mode solver and by the naïve dielectric stack model; see sec. 2.4.

2.2 Bloch Modes and Band Structure of APCW

The introduction of discrete translational symmetry in the dielectric function of the APCW maintains the transverse symmetry of the PNW, allowing similar categorization of modes, but results in Bloch modes with different dispersion properties, shown in fig. 2.5. The propagation constant is now expressed conveniently as a multiple of the reciprocal lattice constant $b = 2\pi/a$, and only the region near the edge of the first Brillouin zone is shown. We now additionally categorize modes by their longitudinal properties: the lower branch of each transverse symmetry is called the dielectric band, because for $\beta_x \approx \pi/a$ the fields are concentrated in the thick part

of the dielectric; the upper branch is called the air band, and is concentrated in the the thin part of the dielectric. Intensity profiles at $\beta_x = \pi/a$ for each symmetry are shown in figs. 2.6 and 2.7. Each symmetry has a spectral region between dielectric and air bands without propagating modes, resulting in a partial band gap. As β_x approaches π/a , the bands flatten to accommodate this partial band gap, resulting in spectral regions of "slow light."

The most important symmetry for our experiment is the TE-like y-odd-z-even symmetry. It can be excited directly by modes launched from the fiber and has a large field in between the nanobeams when excited near the band edge (fig. 2.6). Appropriate choice of geometry aligns the band gap to the D_1 and D_2 transitions of cesium (335.1 and 351.7 THz, respectively), allowing us to leverage the modal properties of the APCW for atomic physics [31–34]. At the band edge the group index $n_{\rm g}$ diverges; far from the band edge, the group index for the y-odd-z-even mode is $n_{\rm g} \approx 1.5$ for the dielectric band and $n_{\rm g} \approx 2.6$ for the air band. In the following discussion, references to band properties refer to this symmetry unless otherwise specified. The TM-like y-even-z-odd mode is the other symmetry we can deliberately excite in the experiment. For our geometry it sits a little higher in frequency. While our system is not designed to allow the excitation of the yeven-z-even or y-odd-z-odd modes—y and z symmetries are nominally preserved across the entire length of the waveguide, and fiber-launched modes couple to TEand TM-like modes-they can be observed if symmetries are broken by scatterers or geometric defects (see ch. 4). Most apparent is the dielectric band edge of the yeven-z-even mode, which falls in the y-odd-z-even band gap. These symmetries are very weakly guided in the single-nanobeam section of the waveguide, and coupling dominantly manifests as loss.

Controlling Band Placement

Precise band placement is crucial in our experiment. Figure 2.8 shows the change in group index and mode visibility as the excitation is tuned away from the band edge. Here mode visibility is defined as

$$V(\omega) \equiv \left| \frac{I(\mathbf{r}_1; \omega) - I(\mathbf{r}_2; \omega)}{I(\mathbf{r}_1; \omega) + I(\mathbf{r}_2; \omega)} \right|$$

where $I(\mathbf{r}; \omega)$ is the intensity, $\mathbf{r}_1 = (0, 0, 0)$ is the center of "thick" part of the unit cell, and $\mathbf{r}_2 = (a/2, 0, 0)$ is the center of the "thin" part of the unit cell. Visibility for the dielectric band, and group index for both bands, decays rapidly as the excitation is detuned from the band edge, complicating the construction of stable GM optical



Figure 2.6: *z*-even mode profiles at $\beta_x = \pi/a$. (a) shows the *y*-odd-*z*-even mode and (b) shows the *y*-even-*z*-even mode. In both figures, (i)–(iii) show cuts of $|E|^2$ for the air band, and (iv)–(vi) for the dielectric band.



Figure 2.7: *z*-odd mode profiles at $\beta_x = \pi/a$. (a) shows the *y*-even-*z*-odd mode and (b) shows the *y*-odd-*z*-odd mode. In both figures, (i)–(iii) show cuts of $|E|^2$ for the air band, and (iv)–(vi) for the dielectric band.

dipole traps and reducing Γ_{1D} . Achieving a group index $n_g \ge 5$ requires placing the band edge within 500 GHz of the atomic transition frequency.

Band placement is achieved by modifying the geometry. All dimensions specified in fig. 2.2 are readily changed, except thickness *t*, as mentioned above. By convention we do not change the lattice constant *a*; this leaves nanobeam width *w*, amplitude *A* and gap *g*. As it turns out, *g* is a relatively weak knob, so bands are primarily placed by changing *w* and *A*. Figure 2.9 shows the simulated change in TE-like band gap center frequency f_0 and stopband width Δf as *w* and *A* are varied. Changes in *A* dominantly affect Δf , while changes in *w* dominantly affect f_0 , though there is appreciable effect on Δf as well. Dashed and dot-dashed lines represent the $f_0 = (f_{D_2} + f_{D_1})/2$ and $\Delta f = (f_{D_2} - f_{D_1})$ isocurves, respectively; the idealized crystal dimensions (those which align the band gap to the cesium transitions) occur at their intersection.



Figure 2.8: Mode visibility and group index. (a) Group index n_g (above) and mode visibility $V(\omega)$ (below) for the *z*-even-*y*-odd mode near the band edge. Red (blue) curves represent dielectric (air) band. (b) $|E|^2$ in the z = 0 plane for dielectric (left) and air (right) bands, at different detunings from the band edge. Good visibility persists for the air band, but visibility rapidly disappears for dielectric band.

These isocurves inform the geometry of the photonic crystal taper, which is specified as a linear ramp from zero amplitude to A and from a start width w_0 to w: the instantaneous width Ω in the taper can be written

$$\Omega(x) = (1 - \xi(x)) w_0 + \xi(x)w + \xi(x)A\cos\left(\frac{2\pi x}{a}\right),$$

$$\xi(x) = \frac{x}{N_t a}, \quad 0 \le x \le N_t a,$$

where N_t is the number of taper cells. The linear taper geometry is illustrated in fig. 2.9c. Following the advice in ref. [50], w_0 is chosen so that f_0 remains approximately constant throughout the tapering region, ensuring the cesium transition frequencies are never inside the instantaneous band gap. Linear sweeps of w and Ado not result in linear sweeps of local band properties: as is apparent from fig. 2.9b, Δf remains relatively small until $A/w \approx 0.3$, reducing the "effective" length of the taper. Linear tapering taper is not optimal, but simplifies drawings and produces "robust" understandable behavior; more details about tapering are discussed



Figure 2.9: **Band gap placement: width and amplitude.** Change in (a) f_0 and (b) Δf with amplitude and width. Dashed line marks f_0 isocurve corresponding to D_1-D_2 mid-frequency; dot-dashed line marks Δf isocurve corresponding to D_1-D_2 frequency difference. For these simulations t = 200 nm, g = 220 nm, a = 370 nm. (c) Illustration of tapering region: initial waveguide width w_0 is smaller than nominal crystal width w, approximating the f_0 isocurve in (a).

in sec. 2.4. Linear coefficients for crystal dimensions near the target band placement are reported in tab. 2.1.

Because precise band placement is difficult in the presence of significant process variance, we have developed a post-processing technique based on atomic layer deposition (ALD) to fine tune the APCW bands by conformally depositing an alumina adlayer (n = 1.75) on all surfaces of the waveguide. This process is described in more detail in secs. 3.6 and 4.2. The effect on the geometry is perturbative for the thin (ca. 10 nm) films we deposit—were we to deposit very thick films, the we might have to consider more complicated effects, like roughness "healing" [51] on our intentionally modulated exterior. Modeled shift rates per adlayer thickness θ are also reported in tab. 2.1.

| | $\delta(f_0)$ | $\delta(\Delta f)$ | unit |
|-----------------|---------------|--------------------|---------|
| δw | -243 ± 1 | -108 ± 1 | GHz/nm |
| δA | 28 ± 1 | 249 ± 2 | GHz/nm |
| δg | 50 ± 10 | 30 ± 10 | GHz/nm |
| ба | -310 ± 10 | -70 ± 10 | GHz/nm |
| δn | -86 ± 1 | -8 ± 1 | THz/RIU |
| δt | -240 ± 10 | -2 ± 1 | GHz/nm |
| $\delta \theta$ | -740 ± 15 | -180 ± 5 | GHz/nm |
| δB | 263 ± 7 | 160 ± 10 | GHz/nm |
| $\delta \phi$ | 2.0 ± 0.1 | 2.0 ± 0.1 | THz/rad |
| δs | -108 | -50 | GHz/nm |

Table 2.1: **Tuning rates of band gap with geometry (simulated).** Shift of band gap center f_0 and width Δf with nanobeam width w, modulation amplitude A, intrabeam gap g, lattice constant a, refractive index n, nanobeam thickness t, alumina adlayer thickness θ , shape distortion B, and phase shift ϕ .



Figure 2.10: Geometric distortions. (a) non-sinusoidal modulation; (b) phase shift; (c) width asymmetry.

Geometric Distortions

In addition to modeling how the bands behave as geometry is intentionally modified, it is also instructive to model how unintended geometric distortions affect the band structure. Here I consider a few common perturbations to the unit cell, depicted in fig. 2.10. The magnitude and signatures of distortions inform at what point we ought to be concerned, particularly for distortions which violate y or z symmetries, which can result in band mixing, and may help account for discrepancies between modeled and measured band structure (see sec. 4.1).

First, I consider a common error resulting from over- or under-dosed patterns: non-sinusoidal corrugation (see sec. 3.3). In this case the exterior corrugation remains periodic, but the contour deviates from the intended sinusoid, as depicted



Figure 2.11: **Phase shift.** Modes pictured have even z symmetry, with solid (dashed) lines indicating odd (even) y symmetry. (a) Band structure for nominal crystal ($\phi = 0$ rad), (b) $\phi = 1.6$ rad, and (c) $\phi = 3.2$ rad.

in fig. 2.10a. Ideally, the width of a single modulated nanobeam Ω can be written

$$\Omega(x) = w + A\cos\left(\frac{2\pi x}{a}\right).$$

A simple way to capture deviations from this sinusoid is to include the next term in the Fourier series:

$$\Omega(x) = (w - B) + A\cos\left(\frac{2\pi x}{a}\right) + B\cos\left(\frac{4\pi x}{a}\right),$$

where $4|B| < |A|$.

Parameterizing the constant width term this way ensures the peak-to-peak amplitude remains constant. This distortion maintains y and z symmetry in the dielectric function, so PNW symmetries continue to be sensible categories. Simulated band movement for small B is summarized in tab. 2.1. Another observed distortion is a relative phase shift of the two nanobeams, depicted in fig. 2.10b. For modest phase shifts ϕ , band movement can be captured as a linear shift rate, as in tab. 2.1, but for sufficiently large phase shifts, the violation of y symmetry results in nonperturbative changes to band structure, shown in fig. 2.11. While the phase shift is unintentional in our project, a similar technique has been used intentionally to create systems of waveguides and emitters with directional emission [52]. Hybridization of y symmetries first occurs around $\phi \approx \pi$.

A third possible distortion is nanobeams with asymmetric width, depicted in fig. 2.10c. This can be captured by assigning different widths to the top and bottom nanobeams: $w_{top} = w - s$ and $w_{bottom} = w + s$. Like phase shifts, this distortion violates *y* symmetry. Band shift rates with *s* are in tab. 2.1, and band diagrams are shown in fig. 2.12. Hybridization of *y* symmetries first occurs around $s \approx 30$ nm.



Figure 2.12: Width asymmetry. Modes pictured have even z symmetry, with solid (dashed) lines indicating odd (even) y symmetry. (a) Band structure for nominal crystal (s = 0 nm), (b) s = 29 nm, and (c) s = 58 nm. For sufficiently large s, second and third z-even modes hybridize.

2.3 Monogators

The APCW geometry has been carefully designed to support large Purcell factors and GM atomic traps; realizing these structures is the primary focus of our fabrication efforts. However, over the course of this project other photonic crystal structures have been designed and fabricated: Fabry–Pérot cavities with corrugated "fishbone" mirrors, used primarily as diagnostic structures for process characterization, and 1D slot waveguides, which provide more control over spontaneous emission and larger n_g , are discussed in ref. [48].

Double nanobeam and slot waveguide structures are attractive because certain modes (in the APCW, the TE mode) have large fields between the nanobeams, necessary for deep GM trapping potentials and large atom–light coupling. However, the dual beam geometry leads to increased fabrication complexity. Pursuing simplicity, one might consider single-nanobeam corrugated structures: the y junction, a source of scattering, is no longer necessary; stiction between APCW nanobeams, which sometimes occurs in wet chemical processing or as the result of SEM imaging, is no longer possible. From a fabrication perspective, it's attractive to have a structure for which all geometric parameters are lithographically defined—in the APCW, mechanical effects change g after release from substrate (see fig. 4.16 and surrounding discussion).

In certain applications the advantages afforded by single-beam structures may outweigh the reduction in atom–photon coupling. In this section I briefly discuss the properties of a corrugated single-nanobeam photonic crystal waveguide, referred to as the *monogator* (MPCW).



Figure 2.13: **APCW band edges vs. nanobeam separation.** Plots show band edge frequency ($\beta_x = \pi/a$) for different APCW symmetries as nanobeam separation *g* is varied. Red (blue) lines indicate odd (even) *z* symmetry, with solid lines indicating TM (TE) mode. Vertical black solid line indicates nominal *g* for the APCW.

Symmetries and Band Structure

Instead of directly considering the band diagram of an MPCW, it's instructive to consider limiting cases for the bands of an APCW. In the limit where the nanobeams are infinitely far apart $(g \rightarrow \infty)$, single nanobeam modes cannot hybridize, and energies of the *y* symmetries become degenerate; in the opposite limit where the nanobeams touch $(g \rightarrow 0)$, the energies of these modes are "pushed" apart. This trend is shown in fig. 2.13, which plots the band edge frequencies ($\beta_x = \pi/a$) for each APCW symmetry as the gap *g* is varied. Keeping uncoupled symmetries far away can be advantageous, as it reduces the potential for crosstalk.

The APCW is designed to align the TE bands to cesium D_1 and D_2 transitions. The TE symmetry was chosen over others as it produces the largest field in between the nanobeams (see figs. 2.6 and 2.7). In an MPCW this region does not exist, meaning TE is not preferred a priori. Figure 2.14a shows the parameterized geometry of an MPCW; fig. 2.14b (c) shows an MPCW band diagram with TE (TM) modes aligned to the cesium D line transitions.

Parameters for MPCWs shown here were iteratively and coarsely optimized "by hand"—structures with better band and field properties certainly exist. A few qualitative comments about the parametric landscape bear mentioning. First, opening a wide band gap for the TM-like modes is difficult; the system pulls towards unrealizable geometries: $A \rightarrow w$, resulting in the minimum width 2(w - A) approaching



Figure 2.14: **MPCW.** (a) Diagram of MPCW geometry, parameterized identically to APCW with g = 0 nm (see fig. 2.2). (b) TE- and (c) TM-aligned band structures. Solid lines indicate TE (blue) and TM modes (red), and cesium transition frequencies (black); dashed lines indicate other symmetries. Apparent movement of light cone (black region) results from different lattice constants, a: TE-aligned MPCW dimensions are (w, A, a) = (199, 124, 350) nm, and TM-aligned are (w, A, a) = (251, 216, 400) nm. Thickness is fixed at t = 200 nm.

zero. In my design I have constrained 2(w - A) = 70 nm. Secondly, to align the bands to the cesium transitions requires changing the lattice constant *a* from the nominal 370 nm for the APCW to 350 nm for the TE-aligned MPCW and 400 nm for the TM-aligned MPCW. This change in lattice constant affects the relationship of the band structure to the light cone. The maximum frequency ω_{max} for which guided modes exist is

$$\omega_{\max} = \frac{c\pi}{a},$$

hence increasing *a* reduces ω_{max} , moving it closer to the atomic transition frequencies. For a = (350, 370, 400) nm, $\omega_{\text{max}} \approx 2\pi \times (428, 405, 375)$ THz. As modes approach the light line, they are increasingly weakly guided, resulting in significant evanescent fields. Evanescence is necessary for atoms to interact with the mode, but also makes modes more susceptible to scattering loss.

Field profiles

Figures 2.15, 2.16 and 2.17 show estimated Purcell factor (eq. 1.2) absent group index enhancement (i.e., P_{wg}/n_g), using the mode area [31]:

$$A_{\rm m}(\mathbf{r}) \equiv \frac{\int \mathrm{d}^3 r' \,\epsilon(\mathbf{r}') |E(\mathbf{r}')|^2}{a\epsilon(\mathbf{r}) |E(\mathbf{r})|^2},$$



Figure 2.15: **Purcell factor, APCW (TE).** Top (bottom) row shows dielectric (air) bands, and left (right) column shows thick (thin) cross-section. Color indicates $P_{\rm wg}/n_{\rm g}$ (logarithmic scale). Top left cross-section is to be compared to ref. [32], fig. 1(c) when multiplied by $n_{\rm g} = 11$.

where $\epsilon(\mathbf{r})$ is the relative permittivity and $E(\mathbf{r})$ the electric field magnitude at location \mathbf{r} . Calculating a total Purcell factor requires group index information, illdefined at the X-point due to the van Hove singularity. The top left panel of fig. 2.15 corresponds to the spatial Purcell factor shown in ref. [32], fig. 1(c). In that work, the group index is reported to be $n_g \simeq 11$, corresponding to a Purcell factor of $P_{wg} = (1.0 \pm 0.1)$ at the trapping position $(x_1, y_1, z_1) = (0, 0, 220)$ nm.

As evidenced by fig. 2.15, maximal P_{wg} for the APCW TE modes occurs in between the nanobeams, at the thick part of the structure for the dielectric band and thin part for the air band, and falls off roughly exponentially with distance from the nanobeams. For MPCWs, P_{wg} is maximized at the surfaces—left and right sides for TE polarization (fig. 2.16), and top and bottom for TM (fig. 2.17).

As demonstrated in ref. [32], atoms can be trapped above an APCW in the potential created by an external beam. Figure 2.18a shows a similar optical potential for



Figure 2.16: **Purcell factor, MPCW (TE).** Top (bottom) row shows dielectric (air) bands, and left (right) column shows thick (thin) cross-section. Color indicates P_{wg}/n_g (logarithmic scale).

a single-beam waveguide, with analogous trapping locations indicated. Unfortunately, the z_1 potential minimum used in ref. [32] coincides with the surface of the single-beam structure, reducing its feasibility as an atom trap. The z_{-1} minimum could conceivably be used, but is quite far (z = -520 nm) from the waveguide. Figure 2.18b shows P_{wg}/n_g for the APCW and MPCWs considered above. The intersection of solid black and blue lines indicates the conditions in ref. [32] (220 nm trap position, dielectric band APCW). The increased evanescence of TM MPCWs (green lines) extends this high-field region to nearly 400 nm, but not to the z_{-1} position (black dashed line). Two-color GM traps like those described in ref. [22], or a combination of external and GM potentials, may produce trap potentials with minima closer to the structure, resulting in better coupling.

2.4 Dielectric Stack Model

Up to this point we have examined systems with continuous or discrete translational symmetry along a given axis, implicitly with infinite extent. The waveguides used in


Figure 2.17: **Purcell factor, MPCW (TM).** Top (bottom) row shows dielectric (air) bands, and left (right) column shows thick (thin) cross-section. Color indicates $P_{\rm wg}/n_{\rm g}$ (logarithmic scale).

our experiment are finite—only $N_n = 150$ nominal unit cells—and many signatures we use to assess our photonic crystals are direct results of their finite length. While in general finite-difference time-domain (FDTD) methods are necessary to obtain detailed information about the properties of finite or higher-dimensionalⁱ systems, many important features of our system can be extracted from a less computationallyintensive, one-dimensional transfer matrix model (TMM), allowing intuition to be gained rapidly.

A transfer matrix relates the incident and outgoing electric fields on one side of an optical element to the other side. The formalism is convenient because the matrix for a complete system is simply the matrix product of constituent elements. Detailed information on the formalism can be found in ref. [53]; refs. [12, 54] also have relevant discussions. Conventions used here have the transfer matrix operate on a column vector of the electric fields on the right side of an optical element

ⁱi.e., two- or three-dimensional band gap structures.



Figure 2.18: **MPCW external illumination.** (a) FDTD simulation of intensity pattern (color indicates $|E|^2$, arb. units) formed by a single nanobeam structure illuminated by a plane wave. As in ref. [31], fig. 1(b), the incident wave is near detuned from the cesium D₂ line, propagates in the negative z direction and is x polarized. This pattern creates optical potentials for atoms ($U \propto |E|^2$). Simulation is two-dimensional; waveguide is 718 × 200 nm. (b) P_{wg}/n_g for different modes along the z axis. Blue lines correspond to APCW (fig. 2.15), red to TE MPCW (fig. 2.16) and green to TM MPCW (fig. 2.17) modes, with solid lines indicating dielectric and dashed lines air bands (top left and bottom right, respectively, in referenced figures). Solid black line indicates position of atoms trapped in ref. [31], and dashed black line position of z_{-1} in (a).

to give the fields on the left; the first element of this vector represents the fields which propagate toward the left; and primed reflection and transmission amplitude coefficients indicate what happens to fields incident from the left. These conventions are summarized in fig. 2.19a. The general transfer matrix **M** can be written [53]

$$\underbrace{\begin{pmatrix} E_L^{\text{out}} \\ E_L^{\text{in}} \\ \mathbf{l} \end{pmatrix}}_{\mathbf{l}} = \underbrace{\frac{1}{t} \begin{pmatrix} tt' - rr' & r' \\ -r & 1 \end{pmatrix}}_{\mathbf{M}} \underbrace{\begin{pmatrix} E_R^{\text{in}} \\ E_R^{\text{out}} \\ \mathbf{r} \end{pmatrix}}_{\mathbf{r}}.$$
(2.2)

Often the sought quantities are the (complex) reflection and transmission amplitude coefficients for a system of many elements. One can verify from eq. 2.2 that

$$r = -\frac{M_{21}}{M_{22}},$$
 $t = \frac{1}{M_{22}},$ $r' = \frac{M_{12}}{M_{22}},$ $t' = \frac{\det \mathbf{M}}{M_{22}}.$

Modeling a wave traveling in a stratified medium only requires two basic transfer matrix elements. The first is a matrix representing propagation in an isotropic



Figure 2.19: **TMM conventions.** (a) Diagram of the conventions used for the transfer matrix formalism. (b) Unit cell represented by eq. 2.3.

medium:

$$\mathbf{P}(n,L) \equiv \begin{pmatrix} e^{\mathrm{i}nkL} & 0\\ 0 & e^{-\mathrm{i}nkL} \end{pmatrix},$$

where *n* is the refractive index of the medium, $k \equiv \omega/c$, and *L* is the distance traveled through the medium. The other required element represents the interface between two materials with different refractive indices. The reflection and transmission amplitudes at such an interface can be calculated from the Fresnel equations [55]. At normal incidence, these amplitudes can be written

$$r = \frac{n_a - n_b}{n_a + n_b},$$
$$t = \frac{2n_a}{n_a + n_b},$$

where a subscripts indicate the originating medium. Using these coefficients and eq. 2.2 we can create a transfer matrix for an interface:

$$\mathbf{I}(n_a, n_b) \equiv \frac{1}{2n_a} \begin{pmatrix} n_a + n_b & n_a - n_b \\ n_a - n_b & n_a + n_b \end{pmatrix}.$$

Using elements **P** and **I** we can construct a unit cell for a 1D photonic crystal:

$$\mathbf{U}(n_a, n_b, a) = \mathbf{P}\left(n_a, \frac{a}{4}\right) \mathbf{I}\left(n_b, n_a\right) \mathbf{P}\left(n_b, \frac{a}{2}\right) \mathbf{I}\left(n_a, n_b\right) \mathbf{P}\left(n_a, \frac{a}{4}\right).$$
(2.3)

The unit cell is illustrated in fig. 2.19b. A few additional constraints will allow us to calculate the band structure for this unit cell. If $\{n_a, n_b\} \in \mathbb{R}$, the matrix U satisfies time-reversal symmetry, and its eigenvalues can be written $\lambda_{\pm} = e^{\pm i\beta a}$ [53]. By

taking the trace we can write the characteristic equation

$$2\cos(\beta a) = \operatorname{Tr} \mathbf{U}$$
$$= 2\cos\left(\frac{n_a k a}{2}\right)\cos\left(\frac{n_b k a}{2}\right) - \frac{n_a^2 + n_b^2}{n_a n_b}\sin\left(\frac{n_a k a}{2}\right)\sin\left(\frac{n_b k a}{2}\right). \quad (2.4)$$

For fixed n_a and n_b , this equation relates the propagation constant β to frequency $\omega = ck$, allowing us to construct a dispersion diagram. Band edge frequencies are obtained by solving for k with $\beta a = \pi$. Note that the LHS of eq. 2.4 is constrained to [-2,2] for $\beta \in \mathbb{R}$ —in the band gap solutions for β are imaginary, as expected. For a = 370 nm, the band gap center f_0 and width Δf are approximately given by

$$f_0(n_a, n_b) \approx (673.9 - 140.0(n_a + n_b))$$
 THz, (2.5)

$$\Delta f(n_a, n_b) \approx (166.1 |n_a - n_b|)$$
 THz. (2.6)

By concatenating these elements, we can create a model of the finite crystal, including the taper between parallel nanobeam and nominal crystal.

Group Velocity and Index

The group velocity

$$v_{\rm g} \equiv \frac{\mathrm{d}\omega}{\mathrm{d}\beta} \tag{2.7}$$

is of central importance in our experiment through its relationship to the local density of states. It is relatively straightforward to extract group velocity information from the transfer matrix transmission coefficient [56, sec. II]: as the propagation constant β accounts for how quickly phase accumulates as a wave propagates, in a finite structure we can to define a mean propagation constant $\overline{\beta}$ via

$$\phi = \int_0^L \mathrm{d}x \,\beta(x)$$
$$= \bar{\beta}L$$

and thence a mean group velocity

$$\bar{v}_{g} = \frac{d\omega}{d\bar{\beta}} = L\frac{d\omega}{d\phi},$$
(2.8)

or equivalently group index $\bar{n}_g \equiv c/\bar{v}_g$. This quantity can be readily calculated from a transfer matrix, as $\phi = \arg(t)$.



Figure 2.20: Group index calculations using TMM. Dashed line represents the group index for an infinite structure calculated using the dispersion relation eq. 2.4; solid line is calculated using a finite TMM ($N_t = 15$ and $N_n = 150$) and eq. 2.8; × points are inferred via eq. 2.9 using the full length of the waveguide ($L = (2N_t + N_n)a$), and \circ points are inferred using only the length of the nominal crystal ($L = N_n a$).

In the experiment, phase information is difficult to extract, but some inferences can be made from transmission or reflection spectra. In many ways, the finite photonic crystal behaves like an etalon made of a highly dispersive material. Resonances in an etalon occur whenever the round-trip phase accumulation $\Phi = 2\beta L$ advances 2π . Thus the group velocity can be estimated at frequencies between resonances to be

$$\bar{v}_{g}\left(\omega_{i} + \frac{\Delta_{i,i+1}}{2}\right) = \frac{d\omega}{d\bar{\beta}}$$

$$\approx \frac{\Delta_{i,i+1}}{\frac{d\Phi}{2L}}$$

$$= \frac{L\Delta_{i,i+1}}{\pi},$$
(2.9)

where ω_i is the angular frequency of the *i*th resonance and $\Delta_{i,i+1}$ is the angular frequency difference between the *i*th and (i + 1)th resonance. In fig. 2.20 I plot the group index $n_g \equiv c/v_g$ for an infinite structure, for a finite structure using 2.8, and estimates using 2.9. Because the length *L* of the etalon is tricky to define with the "soft" tapered boundaries, I plot estimates both with and without the added length of the taper. These estimates are roughly consistent with the group index for the infinite system given by solving eq. 2.4.

The structure of band edge resonances can also be fit to a dispersion model. Here

I present a simple model, similar to the one discussed in the supplementary information of ref. [32]. Near the band edge, the dispersion can be approximated as a hyperbola:

$$\left(\frac{\Delta\omega}{\delta\omega}\right)^2 - \left(\frac{\Delta\beta}{\delta\beta}\right)^2 = 1.$$
 (2.10)

Here $\Delta \omega = \omega - \omega_0$, $\Delta \beta = \beta - \beta_0$, where (β_0, ω_0) specifies the center of the hyperbola and $\delta \omega$ and $\delta \beta$ are parameters which determine curvature and asymptotic behavior. To calculate the group velocity $v_g = \frac{\partial \omega}{\partial \beta}$, we can take a derivative with respect to β :

$$v_{\rm g} = \pm \left(\frac{\Delta\omega}{\delta\omega^2}\right)^{-1} \frac{1}{\delta\beta} \sqrt{\left(\frac{\Delta\omega}{\delta\omega}\right)^2 - 1},\tag{2.11}$$

with the minus sign for dielectric bands and plus sign for air bands. From v_g we can calculate a group index $n_g \equiv c/v_g$

$$n_{\rm g} = \pm \left(c \times \frac{\delta \beta}{\delta \omega} \right) \frac{\left(\frac{\Delta \omega}{\delta \omega}\right)}{\sqrt{\left(\frac{\Delta \omega}{\delta \omega}\right)^2 - 1}}.$$
 (2.12)

By taking the limit of this expression as $\Delta \omega \to \infty$, we can identify $c \times \delta \beta / \delta \omega$ as the asymptotic group index $n_g^{(\infty)}$ (≤ 2 in our structures), leaving $\delta \omega$ and ω_0 as parameters to be fit. Note the van Hove singularity occurs when $\Delta \omega = \pm \delta \omega$, giving the band edge frequency $\omega_{be} = \omega_0 \pm \delta \omega$. These fit parameters can also be used to crudely reconstruct the band diagram; using $\delta \beta = \delta \omega n_g^{(\infty)} / c$ and identifying $\beta_0 = \pi/a$, we have

$$\beta = \frac{\pi}{a} - \frac{\delta\omega \, n_{\rm g}^{(\infty)}}{c} \sqrt{\left(\frac{\Delta\omega}{\delta\omega}\right)^2 - 1}.$$
(2.13)

Taper Length

As mentioned above, we terminate our crystal with tapering regions which transition between the double nanobeam and nominal crystal. The gradual transition reduces the impedance mismatch between standard and slow light waveguides. Currently the number of tapering cells N_t is 15, though only about half of them have $\Delta f \gtrsim 0$ (see fig. 2.9 and surrounding discussion). Figure 2.21 shows the TMM reflectivity decreases as N_t increases. A decrease in free spectral range is also apparent, resulting from increased effective length. Even with large N_t , reflectivity at the band edge is difficult to suppress. At fixed frequency, reflectivity decreases as N_t^{-2} for linear tapers [50]. Linear tapering simplifies fabrication and produces a system with simple behavior, but optimization techniques may produce better impedance matching (lower reflectivity) for fixed N_t [57].



Figure 2.21: **Taper length: reflectivity and group index.** (a) Reflectivity *R* vs. taper length $N_t a$. Impedance matching improves as the number of taper cells N_t increases, as evidenced by the decreased reflectivity. (b) Average group index vs. taper length. Blue line shows \bar{n}_g for $N_t = 5$, red shows \bar{n}_g for $N_t = 30$, and dashed black line shows n_g for nominal unit cell. As impedance matching improves, \bar{n}_g more closely resembles n_g .

Intensity Profile

The TMM also allows us to model the envelope of the electric field inside the APCW. Consider a system composed of two elements, $\mathbf{M} = \mathbf{K} \mathbf{L}$, excited from the right ($E_L^{\text{in}} = 0$):

$$\begin{pmatrix} t E_R^{\rm in} \\ 0 \end{pmatrix} = \mathbf{K} \mathbf{L} \begin{pmatrix} E_R^{\rm in} \\ r E_R^{\rm in} \end{pmatrix}.$$

If the transmission for the total system M is known, the fields at the interface between K and L are given by

$$\begin{pmatrix} E_{\leftarrow} \\ E_{\rightarrow} \end{pmatrix} = \mathbf{K}^{-1} \begin{pmatrix} t E_R^{\text{in}} \\ 0 \end{pmatrix}$$
$$= \frac{E_R^{\text{in}}}{M_{22}} \begin{pmatrix} K_{11}^{-1} \\ K_{12}^{-1} \end{pmatrix}.$$
(2.14)

We can use these fields to calculate the intensity envelope I at this interface as

$$I = \left| \frac{E_R^{\text{in}} K_{11}^{-1}}{M_{22}} \right|^2 + \left| \frac{E_R^{\text{in}} K_{12}^{-1}}{M_{22}} \right|^2.$$
(2.15)

By "breaking open" the dielectric stack at the interface of each unit cell, we obtain mode profiles.



Figure 2.22: **Supermodes (TMM).** Color indicates \sqrt{I} (instead of *I*) for improved contrast. Waveguide is excited from the left. *Top*: air band edge. *Bottom*: dielectric band edge.

Figure 2.22 shows eq. 2.15 calculated near the band edges. The intensity envelope (referred to as a *supermode*) of the band edge resonances are apparent: on each side, the resonance nearest the band gap has a single antinode, the next-nearest two antinodes, and so on. The existence of supermodes has obvious relevance to the atomic physics experiment—an atom at a node interacts with the waveguide in a very different way to an atom at an antinode. Experimentally-obtained intensity data can also be used to determine propagation constant inside and outside the band gap by fitting with a simple model [33].

Long-Range Distortion and Loss

The qualitative character of the band edge resonances is strongly affected by slow changes to the local band gap. These sorts of long-range deviations may result from slow changes in the background dose during *e*-beam lithography, from mechanical effects after undercut, or from inhomgeneous cesium accumulation. Figures 2.23ab shows the effect of changes in the local band structure, here modeled as a Gaussian deviation of f_0 : $f_0(j) = \bar{f}_0 + \delta f_0(j)$, where *j* indicates the index of the unit cell in the TMM, \bar{f}_0 is the unperturbed center frequency, and

$$\delta f_0(j) = \Delta \exp\left(-\frac{(j-j_0)^2}{2\sigma^2}\right),$$

where Δ indicates the maximum frequency perturbation, j_0 the location of this maximum, and σ the width of the distribution. This is achieved by modifying n_a and n_b via eqs. 2.5 and 2.6 to produce the desired perturbation. Here I fix $\sigma = 25a$, $j_0 = (N_{\rm n} + 2N_{\rm t})/2$, and vary Δ .

At one band edge (b) these sorts of perturbations can result in higher-Q resonances as reflection is increased at the ends of the structure; at the other (a) the impedance matching "improves," resulting in increased transmission between first and second resonances.

At our operating frequency, silicon nitride has very low optical loss: studies of Fabry-Pérot-like resonators in our group [48] have demonstrated losses of ≤ 0.6 dB/cm around 850 nm; losses as small as 0.04 dB/cm have been measured at 1550 nm [58]. Still, our system has losses, resulting from fabrication imperfections and, significantly, from cesium exposure (sec. 4.6). Systems with losses (or gains) can be modeled in the transfer matrix paradigm using complex refractive indices $\{n_a, n_b\} \in \mathbb{C}$. This invalidates the assumptions of eq. 2.4, but can be added post hoc to the indices calculated for a lossless system.

The spatial distribution of losses impacts the character of their effect on reflection and transmission spectra. Figure 2.23c shows the how the distribution of the absorber affects band edge features. Loss is modeled as an imaginary index with magnitude normally distributed about the center of the crystal: all refractive indices are replaced with complex values $(n \rightarrow n + ik)$, with spatial distribution

$$k(j) = k_0 \frac{\exp\left(-\frac{(j-j_0)^2}{2\sigma^2}\right)}{\sqrt{2\pi}\sigma}.$$

The normalization ensures that as the distribution is changed attenuation remains nearly constant. For the loss model $k_0 = 0.02$ is fixed, and distribution σ is varied: the top plot has the narrowest distribution ($\sigma = 10a$), and exhibits disproportionate loss for the first resonance (and more generally, odd-numbered resonances) when compared with the bottom plot, which has a more distributed absorber ($\sigma = 100a$). This can be understood from the supermode distribution in fig. 2.22: modes with peaks that overlap with an absorber are more heavily attenuated.



Figure 2.23: Long-range band gap deformation and losses. (a) and (b) Effects of long-range band gap deformation. Black shaded area represents local band gap across length of crystal. Top has undistorted band structure ($\Delta = 0$ THz), middle and bottom have Gaussian deviation of f_0 with magnitude $\Delta = -0.5$ THz and $\Delta = -1$ THz, respectively. In all plots Gaussian has width $\sigma = 25a$ and is centered with respect to the crystal. (c) Localized losses. In all plots integrated loss is identical, but distribution changes. Loss has Gaussian distribution about center of crystal, with width (from top to bottom) $\sigma = 10a$, 50a and 100a. In all cases loss is modeled as imaginary refractive index, with peak value $k_0 = 0.02$.

Accommodating Waveguide Dispersion

While the dielectric stack system is useful for qualitative understanding of finite PCWs, we have implicitly assumed that the system consists of elements with infinite transverse extent. This gives the wrong dispersion, as is evident in fig. 2.24a.

The refractive indices in the TMM can be modified to incorporate waveguide dispersion by giving the refractive indices frequency dependence. We can fit a dispersion model to a band diagram given by an eigenmode solver by reformulating eq. 2.4 as an optimization problem: minimize

$$\epsilon[\mathbf{n}(\omega)]^2 = \sum_{\omega_i} |\operatorname{Tr} \mathbf{U}(\mathbf{n}(\omega_i); \omega_i) - 2\cos(\beta(\omega_i)a)|^2,$$

where ω_i are frequencies at which the dispersion is known, $\beta(\omega)$ is the propagation constant given by an eigenmode solver and $\mathbf{n}(\omega_i)$ contains the refractive indices with an appropriate dispersion model. Figure 2.24b shows experimental reflection data with a fit to a dispersive TMM model, with refractive indices modified according to eq. 2.1 (i.e., linear dispersion with non-zero frequency intercept).



Figure 2.24: **Comparison of TMM to MPB.** (a) Dashed line shows dispersion relation for transfer matrix model calculated using eq. 2.4, dot-dashed line shows fit with compensated refractive indices (eq. 2.1), and 'o' points show APCW band diagram calculated using MPB. Here compensated indices are $n_A(v) = 1.73 - (174 \text{ THz})/v$ and $n_B(v) = 1.64 - (174 \text{ THz})/v$. While imperfect, the transfer matrix model with compensated indices more closely reproduces the band structure of the APCW. (b) Fit of dispersion-compensated TMM (black, dot-dashed) to spectral data from DCOR28D, device 2 (red 'o' points), showing good agreement.

2.5 Mode Overlap Calculations

The sensitivity of band placement to the "critical dimensions" of the unit cell has been discussed. Dimensions of the coupler are also critical—differences of tens of nanometers change the coupling efficiency between fiber and waveguide significantly. Power transfer efficiency at the interface of two translationally-invariant waveguides is given by

$$\eta_{1\to 2} = \left| \operatorname{Re} \left\{ \frac{\mathscr{O}(\mathbf{E}_1, \mathbf{H}_2) \mathscr{O}(\mathbf{E}_2, \mathbf{H}_1)}{\mathscr{O}(\mathbf{E}_2, \mathbf{H}_2)} \right\} \frac{1}{\operatorname{Re} \left\{ \mathscr{O}(\mathbf{E}_1, \mathbf{H}_1) \right\}} \right|, \qquad (2.16)$$

where
$$\mathscr{O}(\mathbf{A}, \mathbf{B}) = \int d\mathbf{S} \cdot (\mathbf{A} \times \mathbf{B}^*),$$
 (2.17)

dS is the surface normal at the interface, and subscript 1 (2) indicates the incident (outgoing) fields. The expression results from projecting an eigenmode of the fiber into the eigenbasis of the waveguide; a good discussion can be found in ref. [45, Sec. 2.2]. For fields from the same waveguide, $\mathcal{O}(\mathbf{E}_i, \mathbf{H}_i)$ gives the integral of the Poynting vector in the transverse plane, which we can identify as twice the power. Refs. [59, Sec. 8.5] and [60] provide insight into why eq. 2.17 is the appropriate choice for the inner product. Power transfer efficiency between arbitrary waveguides must be calculated numerically, but some insight can be obtained by considering

the analytical form of $\eta_{1\to 2}$ for Gaussian beams, as is done in ref. [61]. The modes of single-mode optical fiber are well-approximated by Gaussian beams of the form

$$\mathbf{E}_{y}(r,x) = \hat{\mathbf{y}}E_{0}\exp\left(-\frac{r^{2}}{w^{2}} - \mathrm{i}\beta x\right),$$

where $r = \sqrt{y^2 + z^2}$ and β is the propagation constant in the fiber. The one-way power transfer efficiency between Gaussian beams with different waists is given by

$$\eta_{1\to 2} = \frac{4w_1^2 w_2^2}{(w_1^2 + w_2^2)^2},$$
(2.18)

where w_i represents the beam waist (i = 1, 2). Similarly, for modes with matched waists ($w_1 = w_2 = w$), a transverse displacement results in an efficiency of

$$\eta_{1\to 2} = \exp\left(-\frac{d^2}{w^2}\right),\tag{2.19}$$

where d is the displacement.

The modes of our waveguide coupler are very poorly approximated by Gaussian beams—fig. 2.25a–d show the intensity profile of our rectangular waveguide coupler with different coupler widths. As the coupler increases in width confinement improves, resulting in a decreasingly evanescent field. At sufficient distance from the coupler, intensity I(y, z) decays exponentially:

$$I(y,0) \approx I_0 \exp\left(-\frac{y}{w_p}\right),$$
 (2.20)

where I_0 is a constant that depends on the waveguide dimensions and w_p is the intensity decay length. Figure 2.25e shows the roughly exponential decay for modes depicted in fig. 2.25a–d.

Power transfer efficiency between fiber and coupler are calculated numerically; an MPB script for this calculation is given in appendix B.2. Figure 2.25f compares $\eta_{1\rightarrow 2}$ for the fiber with a Gaussian mode and the fiber with the waveguide coupler as a function of the characteristic intensity length scales, w_p .ⁱ While peak efficiency occurs at different length parameters for fiber–Gaussian and fiber–coupler coupling,

ⁱNB: For consistency with ref. [61], Gaussian width parameters w in eqs. 2.18 and 2.19 are expressed in terms of the electric field E, while eq. 2.20 and fig. 2.25e–f are given in terms of the intensity profile, $I \sim |E|^2$. Quantities relating to intensity profiles are here distinguished by the subscript 'p.' For Gaussian modes, field and intensity length scales are related by $w = \sqrt{2}w_p$, and for exponential decay by $w = 2w_p$.



Figure 2.25: **Coupler mode profiles.** (**a**–**d**) \sqrt{I} profiles for silicon nitride couplers of different widths: (a) 100 µm; (b) 130 µm; (c) 160 µm; (d) 190 µm. In all simulations coupler thickness is fixed at 200 nm. Confinement improves as coupler width increases. (**e**) Cuts of *I* along the *y* axis for the mode profiles shown in (a–d): red, green, blue and black curves correspond to (a), (b), (c) and (d), respectively. At sufficient distance from coupler, field decays exponentially. (**f**) $\eta_{1\rightarrow 2}$ between 780 HP fiber and (solid) Gaussian beam, calculated using eq. 2.18; and (dashed) coupler. For Gaussian beam, w_p is width parameter for intensity profile; for coupler, w_p is 1/e intensity decay length for modes (a–d) obtained from fitting curves in (e).

both exhibit similar qualitative behavior, consisting of a sharp increase towards the optimum and slower falloff as w_p is increased.

More quantitatively useful, however, is fig. 2.26, which shows $\eta_{1\rightarrow 2}$ for TE- and TMlike modes launched from a Nufern 780HP fiber into a silicon nitride waveguide as a function of excitation frequency and coupler width. One immediately observes that TE and TM modes do not have simultaneous width optima, and that, depending on the coupler width, $\eta_{1\rightarrow 2}$ can change appreciably across the operating bandwidth, particularly for the TM mode.

Fibers are aligned to the coupler using silicon vee grooves. If grooves are improperly sized, or if fiber cores are not centered in the cladding or cladding diameters differ from 125 μ m, there may be a transverse displacement of fiber and waveguide modes.ⁱ

ⁱPrivate correspondence with Nufern suggests core/clad offsets are very small—they report 0.5 μ m represents 6 σ in their core/clad offset.



Figure 2.26: Mode overlap: geometry and frequency. Mode overlap $\eta_{1\to 2}$ between (1) 780HP fiber and (2) silicon nitride coupler, as a function of excitation frequency and coupler width in nm for (a) TE and (b) TM polarizations. In all simulations coupler thickness is fixed at 200 nm.



Figure 2.27: Mode overlap: displacement. Mode overlap $\eta_{1\to 2}$ between (1) 780HP fiber and (2) silicon nitride coupler (f = 343 THz, coupler width 140 nm, thickness 200 nm) as a function of (a) y and z displacement. (b) shows a cut along the $\Delta z = 0$ axis.

Figure 2.27a shows $\eta_{1\to 2}$ for transverse displacements of fiber and coupler at fixed dimensions and frequency, and fig. 2.27b shows a one-dimensional cut. As in the case of fibers, fall off in coupling is approximately Gaussian, with width $\sigma = 2.8 \,\mu\text{m}$.

Chapter 3

FABRICATION

In this chapter I outline the procedure to fabricate a chip of APCWs, describing in detail the processing steps, summarized in fig. 3.1, as well as techniques for process characterization. Many steps in this procedure descend from those described in ref. [62, sec. 3.3.1]. UV lithography, *e*-beam lithography and oxygen plasma steps were performed in the Kavli Nanoscience institute at Caltech (KNI); all other processing steps were performed in the Painter group cleanroom.

3.1 Substrate Preparation

The silicon substrates used in this experiment were obtained from Silicon Valley Microelectronics, Inc. (SVM). Two different runs of 200 μ m-thick single-crystal silicon wafers were purchased for this project: the first batch of wafers were low-resistivity (50 Ω ·cm) three-inch wafers produced using the Czochralski (CZ) process; the second batch were high-resistivity (5900–8000 Ω ·cm) four-inch wafers produced using



Figure 3.1: **Fabrication procedure overview,** divided into four subprocedures: *back side patterning*, consisting of substrate preparation (sec. 3.1), UV photolithog-raphy (sec. 3.2), and reactive ion etching (RIE, sec. 3.4); *device layer patterning*, consisting of electron beam lithography (sec. 3.3) and a second RIE; an *undercut*, consisting of a potassium hydroxide (KOH) etch, subsequent wet chemistry, critical point drying (CPD), and an O₂ plasma clean (sec. 3.5); and post-characterization iterative *trim tuning* and *packaging* (described in ch. 4) for the vacuum chamber.

the float-zone (FZ) technique. These wafers are produced with the face normal to the $\langle 100 \rangle$ direction. In both cases SVM deposited a (200 ± 14) nm thick film of high-stress (ca. 800 MPa), stoichiometric silicon nitride on both sides of the wafers using low-pressure chemical vapor deposition. After depositing the nitride, SVM coats the wafer with photoresist and dices the wafer into 10×10 mm squares ("chips") with the edges (nearly) parallel to the $\langle 110 \rangle$ direction.

Solvent Rinse and Piranha Clean

To prepare the substrates for patterning, we perform a solvent rinse and piranha clean. Chips are loaded into a PTFE caddy with device layer facing radially outward. The caddy is submerged in solvent baths consisting of trichloroethylene (TCE), acetone and isopropyl alcohol (IPA), and, after rinsing with deionized (DI) water, a piranha solution consisting of a 3:1 ratio of sulfuric acid and 30% hydrogen peroxide. The sulfuric acid is pre-heated to ca. 75 °C before adding the peroxide. After adding the peroxide, the solution will bubble, and its temperature will rise to ca. 120 °C.ⁱ When bubbling partially subsides (after around 10 s) chips are submerged. We use the caddy to "stir," incorporating the peroxide into the sulfuric acid. Chips remain in solution for 10 min. before being multiply rinsed with DI water: caddy is transferred from water bath to water bath until acid residue no longer visibly diffuses (around two fresh 60 mL baths). Caddy is then removed from the DI water, rinsed under the DI faucet, and each chip is sequentially removed from the caddy, rinsed under the DI faucet, and blown dry with nitrogen. This cleaning sub-procedure is performed many times throughout the fabrication process.

3.2 Photolithography

After cleaning,ⁱⁱ the next step in our procedure is to create an opening in the nitride on the back side of the chip, allowing us to open a rectangular through-window beneath the waveguides in a later step. In contrast to the device layer, the back side pattern has comparatively large features and is the same for all chips, and hence can be efficiently patterned using ultraviolet (UV) photolithography. We begin by

ⁱThis process is sensitive; on one occasion I had a violent reaction, resulting in an acid mess, but thankfully no injuries. Adequate PPE must be worn for this process (acid-resistant apron, face shield and TRIonic ® gloves) and it must be performed in a fume hood. The DI water rinse before piranha *must* be adequate to remove all solvents, which react violently [63] with sulfuric acid.

ⁱⁱOn some occasions I have performed the *e*-beam patterning (sec. 3.3) before UV lithography without obvious detrimental effects. Additionally, if UV lithography is performed first, the piranha may optionally be omitted—the pattern is a macroscopic window, and does not require precision lithography.



Figure 3.2: **Illustration of UV mask** (not to scale). White represents transparent regions, and blue, masked regions. The mask used for this project is a 6 in. square chrome-on-glass mask, which allows many chips to be exposed in parallel.

spinning a film of AZ 5214 E photoresistⁱ onto the front side of the chip. Four drops of the resist are deposited on the chip, one in each corner. The spinner ramps up to 4000 RPM at 2500 RPM/s, and spins for 60 s. Our spin results in a ca. 1400 nm film [64]. After the spin, chips are transferred to a clean glass microscope slide. When all chips have been spun, we cure the photoresist by baking at 100 °C for 60 s. This first layer of resist serves to protect the silicon nitride on the front side of the chip. We then invert the chip and spin more photoresist on the back side, this time baking for 90 s. The resist on the back side will be used to define the window.

We create our pattern by masking the photoresist and exposing it to a UV (405 nm) source. Our chrome-on-glass mask (depicted in fig. 3.2) has a repeating 10×10 mm cell with a 2×6 mm window displaced 2 mm from the top and left sides of the cell. The repeated pattern allows us to expose several chips in parallel. We expose using a Karl Süss Microtech MA6 mask aligner in the KNI, using "soft" contact mode and an alignment gap of 30 µm. A 15 s exposure is sufficient to clear the resist.ⁱⁱ Chips are developed using MF-CD-26 for ca. 100 s,ⁱⁱⁱ rinsed in DI water and blown dry with nitrogen. This pattern is transferred to the nitride via reactive ion etch (RIE), described in sec. 3.4. After the window has been etched into the chip, we remove the resist by solvent rinse and piranha clean, described above. At this point substrates

ⁱObtained from Integrated Micro Materials, 8141 Gateway Drive, Suite 240, Argyle, Texas 76226. http://www.imicromaterials.com

ⁱⁱNB: The two mask aligners in the KNI have different doses; for a process developed on Suss1, the exposure times must be doubled on Suss2.

ⁱⁱⁱThis timing isn't critical; we just monitor the chips in solution and remove them after visually confirming that the resist has cleared.

are ready for device layer patterning, achieved via electron beam lithography.

3.3 Electron Beam Lithography

All features in the device layer are defined via electron beam lithography (EBL). Here I outline the lithography process, including the CAD drawing, fracturing, loading the chips in the EBL tool, pattern exposure, and resist image formation.

Pattern Generation

There are 69 different parameters specifying the geometry for our APCW chips. The most commonly changed parameters are those affecting band placement and coupling efficiency, in an effort to maintain optimal performance in the presence of process drift. These parameters are used to create a CAD drawing. In this work, parameters were tabulated in a text file (.param), fed into a MATLAB function which writes an AutoLISP app (.lsp), which was then loaded in AutoCAD to create the AutoCAD Drawing Exchange Format (DXF) drawing (.dxf). The DXF file is then converted into a format which can be read by the pattern generator on the EBL tool, a process called "fracturing." This sequence of software operations and intermediate file formats are depicted in fig. 3.3. Folders with example patterns (including all intermediate files) are attached as .zip files at CaltechTHESIS.

The parameters file is just a specially-formatted plain text file, and can be opened by any text editor—I specified the extension .param for the purpose of organization, but .txt (or no extension) would function identically. There are many ways to organize this information in a EBL project, but text files have been convenient as they are portable, and interface well with command line utilities, particularly find, grep and diff. The .param file consists of eight sections, each beginning by the character sequence // and a descriptive name. Within these sections, parameters are specified by keywords (strings the MATLAB parser looks for) followed by values. Different parameters have different data types: the majority are physical dimensions, specified in microns, but there are also Booleans (specified as 'yes' and 'no' strings), integers, and polymorphic parameters, which change the behavior of the MATLAB interpreter depending on the supplied type (dimension or string).

Two sections are special: //global, which specifies settings and dimensions pertaining to the whole print; and //scan, which specifies a scan parameter (by convention the string xx). The scan parameter string assumes a different value for each waveguide as parameters are extracted, simplifying parametric sweeps (e.g., coupler width, or APCW width or amplitude). The parameter string can be used in



Figure 3.3: Software flow for *e*-beam patterns. (I) Generating CAD layout. .param file is interpreted by MATLAB writeAutoCAD script, producing an AutoLISP app (.1sp). This app can be run by typing appload at the AutoCAD command line and selecting the desired .1sp file. (II) Fracturing CAD files. This step must be performed on a computer with the software licence for Layout BEAMER. The Python script (.py) must be edited to specify the .dxf file and pattern rotation (top or bottom). This produces two .gpf files, one for each exposure. (III) Layout and exposure. Job is assembled in CJob on the BEAMS computer connected to the EBPG, and queued in Cebpg to perform the exposure.

any section excepting //global and //scan. All other sections in the .param file pertain to elements in the waveguide. Tables 3.1 and 3.2 provide annotations for the keywords in a .param file; letters in the left column correspond to dimensions indicated in fig. 3.4.

Some keywords merit additional explanation. To expedite writes and reduce the total exposed area of our pattern, we utilize the gridding technique described in ref. [62, sec. 1.3.1.3]. This results in nitride "tiles" in the window separated by exposed lines. As tiles are disconnected from the rest of the nitride, they are undercut and float away in the potassium hydroxide etch (sec. 3.5). The width of these grid lines are specified by the keyword tiling gap. The fda length keyword refers to the region of the vee groove at the edge of the chip: the resist in this region is a bit thicker than the resist in the middle of the chip (a phenomenon referred to as "edge bead"). Consequently, the dose in this region is increased by to ensure the resist clears.

Our pattern is partitioned into two exposures: a fine exposure, which uses a low

current beam (300–500 pA) and finer resolution (2.5 nm) to define important geometry (waveguide, photonic crystal, etc.); and a coarse exposure, which uses a higher current beam (5 nA) and coarser resolution (10 nm) to expose less-critical regions (window, vee grooves, etc.). The **fine exposure length** keyword indicates the length of vee groove above and below the window to include in the fine partition. The **hole recess** keyword describes the length of siderail near the top and bottom of the window that has no holes; this feature is discussed in some more detail in sec. 3.5 and fig. 3.19. The tcc keyword refers to whether the roundabout-style coupler should be used (i.e., a coupler with a hole at the intersection of waveguide and tether). This geometry helps reduce the reflection of the tether, and is described in more detail in ref. [48]. Finally, the **inverse crystal** and **inverse** y keywords refer to whether the geometry in critical regions around the y-junction and photonic crystal should indicate the waveguide or the space *around* the waveguide; this subtlety is described in more detail in fig. 3.6 and the surrounding discussion.

Fracturing and Exposure

Before a pattern can be exposed, it must be translated into a format the EBL tool can interpret in a process called fracturing. In the case of the Raith EBPG5000+ tool used in this project, the format is Generic Pattern Format (.gpf). GPF partitions a pattern into exposures, mainfields and subfields.

An exposure consists of a binary pattern of exposed and unexposed regions.ⁱ Exposures are segmented into mainfields, which are squares representing the largest area over which the electron beam can be deflected. The maximum achievable deflection for the EBPG5000+ at 100 keV acceleration voltage is 579.2 μ m [65], but frequently mainfield size is instead limited by the desired pattern resolution. Beam deflection is achieved by two electrostatic actuation systems in the electron beam column, called the mainfield and subfield deflection. The actuators are controlled by digital-to-analog converters (DACs) with 16 and 14 bits, respectively.ⁱⁱ At maximum mainfield resolution (2.5 nm steps for the EBPG5000+), the mainfield deflection is limited to

$$(2.5 \text{ nm}) \times 2^{16} = 163.8 \ \mu\text{m}.$$

If a pattern exceeds the size of a mainfield, the stage on which the substrate is

ⁱThis is not strictly true—the dose can be modulated to achieve proximity correction or grayscale lithography, but for many projects, including ours, this pattern can be thought of as binary.

ⁱⁱAt the time of writing, there are plans to upgrade the system to 20 bit DACs, increasing the mainfield size at high resolution.

| | Keyword | Default | Description |
|---|---|------------------|--|
| | //global | | Describes window, global settings |
| | filename: | | Name for .1sp output |
| | window: | yes | Whether print should have a window |
| | number of devices: | 16 | How many waveguides to print |
| Α | device spacing: | 320.0 | Pitch between waveguides on chip |
| В | window length: | 1800.0 | Length of window |
| С | window width: | 5000.0 | Width of window |
| | window fillet radius: | 20.0 | Radius of curvature for window corners |
| | tiling gap: | 1.0 | Window grid width (see text) |
| | //grooves | | Describes vee groove, siderails |
| D | vee groove width: | 149.17 | Width of the vee grooves |
| Е | vee groove length: | [2235.0, 3500.0] | Length of vee grooves [front, back] |
| F | fda length: | 300.0 | "Flexible dose array" length (see text) |
| G | length of device past window: | -415.0 | Distance between coupler and window |
| Н | fine exposure length: | [600.0, 600.0] | Extent of fine exposure past window [front, back] |
| | | | (see text) |
| Ι | safety rail width: | 5.0 | Width of siderail |
| J | hole lattice constant: | 3.0 | Pitch of holes in siderail |
| | number of hole rows: | 1 | Number of rows of holes in siderail |
| Κ | hole radius: | 0.5 | Radius of safety rail holes |
| L | hole recess: | 330.0 | Distance from window edge to first hole (see text) |
| | <pre>//coupler parameters</pre> | | Describes fiber coupler |
| Μ | nominal width: | | Width of coupler |
| Ν | tether center width: | 0.200 | Width of coupler tether |
| 0 | prelength: | 2.0 | Distance from coupler end to tether |
| | postlength: | 2.0 | Distance from tether to single nanobeam |
| | fillet radius: | 1.0 | Radius of fillet at tether anchor point |
| | tcc: | yes | Use "roundabout" style coupler? (see text) |
| Р | hole radius: | 0.113 | Radius of roundabout hole |
| Q | tcc radius: | 0.300 | Fillet radius between tether and waveguide |
| | <pre>//single nanobeam parameters</pre> | | Adiabatic taper to nominal waveguide width |
| R | length: | 199.0 | Length between coupler and angled node array |
| S | end width: | 0.370 | End width of single nanobeam |
| | function: | | MATLAB anonymous function describing transi- |
| | | | tion from start to end width. Default is cubic: |
| | | | @(w0,wf,y) w0 + (wf - w0).*y.^3, where w0 |
| | | | is start width, wf is final width and y fraction be- |
| | | | tween start $(y=0)$ and end $(y=1)$. |

Table 3.1: CAD parameters, part 1. Letters correspond to dimensions indicated in fig. 3.4.

mounted must move, potentially resulting in stitching errors. Stage trajectory for a coarse exposure is shown in fig. 3.7b.

Features within a mainfield are segmented into subfields. Though the GPF format specifies several subfield types, the default subfield type is a trapezoid, sometimes referred to as a trap, which is parameterized by a base, height, an orientation (X or Y, depending on which sides are parallel to the tool's coordinate system), and a left and right shift of vertices relative to the implied rectangle to make the shape a trapezoid. The mainfield deflection is used to drive the beam to the lower left corner of a subfield, and the subfield deflection is used to drive the beam through the subfield, exposing the pattern. See fig. 3.5 for illustration. Maximum subfield size is limited by the bit depth of the subfield DAC in an analogous way to the limits on mainfield size. Subfields are pixelated into shots (sometimes called exels) separated

| | Keyword | Default | Description |
|----|--|---------|--|
| | <pre>//angled node array parameters</pre> | | |
| U | end width: | 0.525 | Waveguide width after angled node array |
| V | taper length: | 10.0 | Length over which waveguide widens from start |
| | | | width to blank length |
| | number of tethers: | 16 | Number of tethers in array |
| W | blank length: | 1.0 | Distance between end of taper and first tether |
| Х | width at tethering: | 1.125 | Width of waveguide at tethering |
| | tether center width: | 0.143 | Width of individual tether |
| | lattice constant: | 0.220 | Pitch of tethers at anchor to waveguide |
| | ext. lattice constant: | 0.320 | Pitch of tethers at anchor to siderail |
| Y | tether displacement: | -11.1 | Distance between anchors on waveguide and siderail |
| - | <pre>//cooled single nanobeam parameters</pre> | | Long waveguide between angled node array and PhC |
| | end width: | inherit | inherits end width from angled node array, or number |
| | | | in microns |
| | start length: | 5.0 | Distance to first node array |
| | stop length: | 5.0 | Distance from last node array to next element |
| Ζ | tether displacement: | 0.0 | Distance between anchors on waveguide and siderail |
| AA | spacing: | 150.0 | Spacing between node arrays |
| | //photonic crystal | | |
| | inverse crystal: | yes | Negative-space fracture of photonic crystal (see text) |
| BB | end width: | 0.280 | w_0 |
| CC | gap: | 0.270 | g |
| | inverse y: | yes | Negative-space fracture of y-junction (see text) |
| DD | taper length: | 30.0 | Y-junction length |
| EE | split length: | 15.0 | Y-split length |
| | lattice constant: | 0.370 | a |
| | amplitude: | | A |
| | width: | | W |
| | number taper: | 15 | Number of taper cells $N_{\rm t}$ |
| | number: | 150 | Number of nominal cells N_n |
| | //scan | | |
| | scanned parameter: | XX | Typically used for parametric sweeps of amplitude, |
| | | | width or nominal width |
| | start value: | 1 | |
| | end value: | 16 | |

Table 3.2: CAD parameters, part 2. Letters in right column correspond to dimensions indicated in fig. 3.4; w_0 , g, a, A and w are illustrated in figs. 2.2 and 2.9.

by the beam step size (BSS). Though the subfield deflection can make much smaller steps than the mainfield deflection (as small as 0.08 nm), it is common to set the BSS to the mainfield resolution (called BSS fracturing), as the subfield origin is limited to the resolution of the mainfield DAC. Exposing a shot consists of directing the electron beam to a location on the substrate, and dwelling there until a specified dose (charge per unit area) is achieved. The dwell time τ is given by

$$\tau = \frac{(\text{dose})(\text{beam step size})^2}{(\text{beam current})}.$$

The minimal possible dwell time on the EBPG5000+ is 20 ns (50 MHz exel frequency) [65].

The mainfield size for the fine exposure is limited by the DAC to 160 μ m, whereas the coarse exposure is deflection limited (579.2 μ m). Exposure partitioning is



Figure 3.4: **.param dimensions.** (a) Global overview. (b) Angled node array. Vertical lines at exterior represent edge of vee groove, to which tethers are anchored. (c) Coupler and tapered single nanobeam. (d) Node array and siderail. (e) Negative y-junction. This has been rotated 90° clockwise with respect to (a–d). (f) Negative taper and crystal. This has been rotated 90° clockwise with respect to (a–d). In (e–f) color fill indicates shape defined by geometry ("negative space" geometry).

necessary due to the size of the pattern we are trying to expose: increasing the beam current reduces the dwell time by the same factor, but also increases the beam spot size [66], resulting in process blur; similarly, decreasing the mainfield resolution in most cases increases the mainfield area, reducing the number of stage movements and thereby significantly reducing the total write time, but also reduces the precision with which subfields can be placed. Fine and coarse exposures are combined in a job file (.job) file on the EBPG5000+ and exposed sequentially. The current pattern consists of 16 waveguides, and takes about 90 min. at a dose of 172 μ C/cm².

We create .gpf files using a proprietary fracturing program called Layout BEAMER. BEAMER accepts an input CAD file format (.dxf, .gds), allows a user to perform operations on layers (adding/subtracting layers, e.g.), and then fractures the geom-



Figure 3.5: **Deflection in EBPG5000+.** Illustration of how patterns are exposed on the EBPG. The lower left corner is addressed by the main deflection (red) and the subfield is rastered through using the subfield deflection (blue). Subfields are often trapezoidal (traps). Subfields are pixelated into shots (gray).

etry to .gpf at a specified resolution. The common interface to the BEAMER library is to create a flow using their VisualFLOW[™] GUI, but the complexity of our fracture makes this interface cumbersome. Conveniently, the BEAMER library also supports a Python interface, which is used for this project.

BEAMER supports a few different fracturing algorithms optimized for different use cases [67]. Our pattern consists of many different shapes, and we use each algorithm for different parts of the pattern. Ideally the fracturing algorithm would not influence the resulting pattern, but unfortunately over-fracturing can result in significant errors. Figure 3.6a shows a fractured pattern alongside a scanning electron micrograph (SEM). In the transition region between double nanobeam and photonic crystal, the fracturing pattern comprises several small rectangles. In the SEM of the released device, this region exhibits undesirable roughness. In contrast, the region with a simpler fracturing pattern has smooth sidewalls. A second concern about this pattern is the obvious difference in the fracturing of adjacent, nominally identical unit cells, which potentially could lead to undesirable, disorder-induced photon localization [68, 69]. To enforce that each unit cell is fractured identically we must first constrain the lattice constant to a multiple of the



Figure 3.6: Errors introduced by fracturing. (a) *Right detail*: The interior sidewalls in the photonic crystal taper exhibit roughness correlated with over-fracturing in the fractured pattern. *Left detail*: Illustration of adjacent unit cell fractures. (b) *Right detail*: Simplified fracture reduces interior sidewall roughness and (left detail) regularizes unit cell fracture.

mainfield resolution. Next, we must ensure the geometry is specified in a way which preserves the periodicity. This can be achieved by modifying the CAD drawing. In our process (described below) we remove resist in regions exposed to the electron beam, and hence GPF files specify the "negative space" around the structures we are trying to produce (see fig. 3.4e–f). Originally, the shapes in our CAD drawings defined the waveguide, but exacting finer control over the fracture requires that we instead define the regions around the waveguide, segment unit cells into individual shapes, and perform no operations in the fracture that erase the information about these segments. In Layout BEAMER, this means the layer cannot go through any Boolean, heal, or bias operations before export. Functionally, this is achieved by creating an intermediate GPF export of these shapes, and later merging (without healing) with the rest of the pattern.

This level of control over the fracture is adequate for this project. If finer control



Figure 3.7: Loading chips in EBPG. (a) Chips (purple) are clamped to wafer tables (green) using either copper ground clips (orange) or height reference clamps (blue). Ground clips rotate with table when rotation screws (red) are adjusted; height reference clamps remain fixed. At the top of the holder is the calibration block (dark gray), on which the Faraday cup and reference markers are located. (b) Coarse exposure for "top" layout; red indicates subfields, black squares mainfields, and the blue arrow the order in which mainfields are exposed.

were necessary, .gpf files could be created or edited using the gpfgtx and gtxgpf tools included with Raith's BEAMS software.

Mounting Chips and Pattern Alignment

On the EBPG5000+, Holder 2 (H2) is designed to mount two, 2-in. wafers, one on each wafer table. A cartoon depicting some features of this holder is shown in fig. 3.7a. Commonly this holder is used to mount samples smaller than wafers ("chips"). There are two different types of clamps under which chips can be mounted: copper ground clips at the top (90°) and bottom (270°) of each wafer table, and height reference clamps, located at 0°, 120°, and 240° on the left wafer table and 60°, 180°, and 300° on the right wafer table. When copper clips are used, height reference clamps must be filled with shim pieces to ensure that samples are at the correct height for exposure. Each table has an individual rotation adjustment screw. Pieces clamped under the copper clips will co-rotate with the table when the screw is adjusted—this can be useful to compensate for the dice angle (discussed below).

The height reference clamps are not attached to the wafer table, and will not rotate, so the rotation of chips clamped this way must be adjusted manually and individually with tweezers.

Each type of clamp necessarily obscures part of the substrate, limiting the area of the chip which can be patterned. To accommodate different clamping locations, two different pattern rotations are used: for the chips at the "top" of the table $(0-180^{\circ})$ the pattern is rotated so the window is nearer the bottom of the chip, while for chips at the "bottom" of the table $(180-360^{\circ})$ the pattern is rotated to be nearer the top of the chip. These exposures have been named "top" and "bottom," respectively, though in CJob the window location will appear to be on the right for top exposures and left for bottom exposures, as the EBL tool's coordinate system is rotated 90° from the microscope view. The part of the chip masked by the clamp cannot be patterned, which commonly means that back vee grooves terminate before the end of the chip.

The copper clamps have a very small contact area with the chip surface. This contact point must be sufficiently distant from the edge of the chip to ensure the chip lies flat on the table—chips clamped too close to the edge will cant upwards, due to the edge bead of the resist.ⁱ The contact area for the height reference clamps is larger, and chips can generally be clamped closer to the edge, e.g., on a corner, while maintaining good contact with the wafer table.

When clamped on a corner, it is possible to pattern back vee grooves which extend to the end of the substrate. This prospect is attractive, as it may improve coupling efficiency for back couplers.ⁱⁱ Another possibility could be to rotate the pattern 90 degrees into "left" and "right" patterns, but this might require additional changes in the fracturing—the default mainfield trajectory meanders along the X direction, which for top and bottom fractures results in individual waveguides being written in sequence (see fig. 3.7b). A naïve left or right fracture would write mainfields *across* waveguides, potentially leading to stitching errors.

Our patterns are designed to create vee grooves in the silicon substrate to align an optical fiber to the device layer. Many details about this technique are discussed in ref. [44]. Because this technique relies on etch anisotropy in crystalline silicon, it is important that the device layer pattern be aligned to the crystal planes in the silicon

ⁱThis will register as a height range error in the EBPG: either %ENG_W_SHOUHM or %ENG_W_HEOURA.

ⁱⁱSince November 2016, a few prints have been written with full-length back vee grooves, but improperly-compensated rotation prevented conclusive testing of improved coupling efficiency.

substrate. Ordinarily, simply aligning the chip edges to the EBPG5000+ coordinate system would suffice. Unfortunately, the second batch of wafers received from SVM were diced at a small but varying angle with respect to the $\langle 110 \rangle$ direction, requiring calibration of the offset angle for each wafer and compensation in the EBL tool.

A pattern dry etched into the silicon nitride behaves as a mask during a potassium hydroxide etch; for example, a rectangular mask results in vee-shaped grooves with faces normal to the $\langle 111 \rangle$ direction. Measuring the angle between these crystal planes and the wafer dice allows us to determine the dice angle offset. Figure 3.8 shows a micrograph depicting the crystal planes exposed by the KOH etch, and a sample pattern used to determine the rotation of substrate.ⁱ The potassium hydroxide etch is discussed at length in sec. 3.5.

Once the crystal/dice offset is determined, the chips must be rotation compensated in the EBL tool. The rotation of a chip can be determined by locating the corners using the alignment microscope or EBPG5000+ in SEM mode. Insufficient manual alignment can be compensated for by applying a small rotational correction (up to 0.2 degrees) in the EBPG5000+.ⁱⁱ

Creating an Image in Electron Beam Resist

The electron beam resist used for this project is ZEP520A, produced by Zeon Chemicals. We apply four pipetted drops to the center of the chip and spin at the same speed and ramp as the UV resist spin. Resist is cured on a 180 °C hotplate for 120 s, resulting in a film thickness of ca. 350 nm. ZEP520A is a positive-tone resist, meaning the regions where it is dosed are removed when developed. We develop using Zeon ZED-N50 (n-amyl acetate) for 150 s and rinse in methyl isobutyl ketone (MIBK) for 30 s. In August 2016 I purchased personal glassware to develop our chips, with the goal of improving repeatability; prior to that point I used the common Painter group developer and rinse. Shortly after making this purchase I did produce a crop of "good" science chips, but there is inadequate data to conclude personal glassware had a significant impact.

The behavior of the resist largely determines the transfer function between the input (computer) and output (developed) geometries, and it is informative to characterize

ⁱOne might also consider the "wagon wheel" pattern in ref. [70].

ⁱⁱThis is achieved by spoofing an aligned write in cjob using joystick (JOY) markers: use the SEM to find the coordinates of the chip corners and mathematically counter-rotate around the geometric center to determine the spoofed marker location. Ordinarily joystick markers require the operator confirm a marker has been correctly located, but this behavior can be overridden by an .ini file.



Figure 3.8: **Determining substrate rotation.** *Left:* (a) False-color micrograph of a vee-groove. Yellow planes are normal to the $\langle 111 \rangle$ direction, forming the sides of the vee grooves. (b) Cross test pattern to determine substrate/crystal orientation offset. Each pattern has a slightly different rotation: the top 16 patterns are rotated counter-clockwise, bottom 16 clockwise, in steps of 0.1°. *Right:* The rotational offset is inferred by measuring the angle between the pattern in the nitride and the etched shape below.

some of its properties. One important property is the dose-to-clear, which is the dose (charge per area) required to completely remove the resist when developed. The contrast curve for our process, shown in figure 3.9, indicates a value between 110 and 120 μ C/cm². In most applications the idealized contrast curve resembles a step function, maximizing the slope of the transition, "binarizing" the dose pattern into a mask, with "exposed" regions receiving a dose far above the dose-to-clear and "unexposed" regions a dose far below. This ensures that the pattern transferred to substrate is also binary. The contrast curve is temperature-dependent; there is some evidence that process latitude could be improved by reducing the temperature of the developer [71, 72]; an additional brief discussion about cold develop is in appendix D.

The transfer function between pattern and dose image has some very long-range components: fig. 3.10 shows the simulated electron beam point spread function (PSF) for our ZEP520A-silicon nitride-silicon stack. This simulation is performed using PENELOPE, a Monte Carlo package which calculates electron and photon



Figure 3.9: **Dose-to-clear.** *Left:* test pattern, with dose increasing from (i) 60 μ C/cm² to (ix) 140 μ C/cm² in increments of 10 μ C/cm². Film thicknesses determined using spectral reflectance (Filmetrics F40). *Right:* Contrast curve *t*(*d*). Curve from 14-Dec-2015 (01-Nov-2016) determined by writing an array of 500 μ m (200 μ m) squares and varying the dose. It is important to use test patterns much larger than the longest-range components of the PSF to obtain an accurate contrast curve (see fig. 3.10).

transport in materials [73]. The short-range blur is attributed to forward scattering as the beam interacts with the resist, and the long-range tail ($\sigma \approx 30 \ \mu m$) is attributed to backscattering from the substrate [66]. For small, isolated patterns, the long range tail of the PSF is unimportant, but for large or dense patterns the contributions from backscattering become significant. The dose-to-clear is the threshold value for large, uniform exposures, relying on the large background dose from neighboring regions to exceed the threshold for the resist.

In addition to the simulated point spread function many other factors, including the spot size for the EBPG5000+ and developer temperature, contribute to the total PSF for the process [72]. It is difficult to determine the individual contributions of each of these factors, collectively referred to as process blur.

The total point spread function PSF relates the dose image I and the expressed dose pattern O through convolution:

$$I(x, y) = \iint d\xi \, dv \, O(\xi, v) \text{PSF}(x - \xi, y - v).$$

The resist thickness T after develop is just the function composition of I with the



Figure 3.10: **Electron beam PSF** for charge distribution in resist layer for our stack. Simulated with PENELOPE.

contrast curve t(d) (see fig. 3.9):

$$T(x, y) = t(I(x, y)).$$

This pattern is then transferred to the substrate via reactive ion etch.

If the PSF and contrast curve were known perfectly, we could calculate an O which produced the ideal T. In practice O can be modified to produce something closer to the desired output using technique called *proximity effect correction* (PEC) [67], which compensates the dose of each subfield (increasing it in sparse areas and decreasing it in dense ones) to approximate the idealized, binary dose image.

At what dose should we expose our pattern to minimize the discrepancy between edges in O and T? In general the dose-to-clear is too small to properly expose most patterns, relying on a large accumulated background dose. While the answer to this question is pattern-dependent, the canonical test pattern is the 200 nm pitch, 1:1 line–space pattern. GenISys recommends using this pattern to determine the base dose, D_b , which has the circular definition as the dose which properly exposes this pattern [72]. D_b also has the nice property that the thresholds fall near the intended place for a wide range of process blurs, illustrated in fig. 3.11.

Determination of the base dose is the first step to implement PEC in Layout BEAMER. To this end, line–space patterns were exposed. The pattern consists of a $120 \times 120 \,\mu\text{m}$ swatch of alternating lines and spaces, each programmed to have 100 nm width (i.e., 1:1 space:line ratio with 200 nm pitch). Figure 3.12a shows SEM images of this pattern exposed at different doses, and the corresponding thresholded images. As



Figure 3.11: **Process blur for 1:1 line–space pattern (simulated).** In all figures red represents an exposure with small process blur, blue, large process blur and the black dashed line the threshold for the resist. If a pattern is under-dosed (left), features will come out smaller than intended, with larger discrepancies for larger process blur. If a pattern is over-dosed (right), features will come out too large, with larger discrepancy for larger blur. If a pattern is dosed at the base dose (center), regardless of the blur, features on the interior will be sized correctly.



Figure 3.12: Line-space dose array. (a) *Top*: SEM images of resist pattern. *Bottom*: Images are converted to binary using ImageJ [74] median filter and Make Binary utilities. As dose is increased, the ratio of the resist ridge (space) to the cleared line decreases. (b) Line:space ratio as a function of dose. Blue points are for patterns dosed with a 300 pA beam, and red with a 500 pA beam. The gray region indicates points below the nominal dose-to-clear. A near-unit ratio indicates the base dose for this pattern, obtained for ca. 135 μ C/cm².

the dose is increased, the line width increases, resulting in a non-unit space:line ratio. Figure 3.12b shows this ratio as the dose is increased, revealing the base dose for this pattern occurs at around $D_{\rm b} \approx 135 \,\mu {\rm C/cm}^2$.

The line-space pattern is a clean test, but bears little resemblance to our actual geometry. In fig. 3.13 I plot the simulated process latitude for the APCW dimensions.



Figure 3.13: **Simulated process blur.** In all plots the black dashed line is the nominal (written) dimension and colored lines are expressed dimension as a function of relative dose, given in units of the base dose. The process blur σ_b is varied between 10 and 100 nm in steps of 10 nm (see eq. 3.1); lines closest to the written dimension are 10 nm blur, and discrepancy for fixed dose increases monotonically with σ_b .

For simplicity, I assume the PSF consists of a single Gaussian with blur $\sigma_{\rm b}$

$$PSF(x, y) = \left(2\pi\sigma_{b}^{2}\right)^{-1} \exp\left(-\frac{x^{2} + y^{2}}{2\sigma_{b}^{2}}\right),$$
(3.1)

and that the $D = 0.5D_b$ isodose contours represent the edge of the structure. As the dose is increased the width (a negative feature) decreases and the gap (a positive feature) increases in size; these effects become more severe as the process blur radius increases. Notably, the amplitude is comparatively insensitive to changes in dose, but still changes as the process blur increases. We can use this fact to estimate process blur by comparing the SEM-measured amplitude and pattern dimensions.

Figure 3.14 shows measured dimensions for a dose array. By fitting the measured dimensions to simulation, we can extract process blur of $\sigma_b \approx 40$ nm and $D_b \approx 168 \,\mu\text{C/cm}^2$. Better estimates may be obtained by including long range effects. As a result of this analysis, I reduced the dose for our structures from 230 $\mu\text{C/cm}^2$ to 172 $\mu\text{C/cm}^2$ in January 2016. Some of the test patterns for this study are shown in fig. 3.15.

Though proper dosing reduces object–image discrepancy in the crystal, residual offsets unfortunately remain: discrepancies in measured *A* are dominated by process blur and are only weakly dependent on dose; additionally the dose which minimizes discrepancies for *w* and *g* depends on the local geometry of the photonic crystal, and is not the dose which minimizes discrepancies in the region near the coupler. These discrepancies are referred to as "magic factors," defined as the *written* dimension (i.e., what is specified in the CAD drawing) minus the *measured* dimension (what



Figure 3.14: **Measured process blur.** Dashed line represents the written dimension, circles, measured data, and solid line, fit using process blur simulation (fig. 3.13).



Figure 3.15: **APCW dose array.** Visual inspection confirms general trends in fig. 3.14—gap (width) increases (decreases) with dose, while amplitude is relatively unaffected.

is measured in the SEM images). Positive (negative) magic factor indicates that measured dimensions are smaller (larger) than what is specified in the CAD drawing.

Figure 3.16 shows a comparison of written and measured dimensions for *w* and *A* crystal parameters, with the diagonal dashed line indicating equal written and measured dimensions (zero magic factor). Black points denote measurements of crystals written before January 2016, which were dosed at 230 μ C/cm²; both *w* and *A* are undersized with respect to written dimensions (positive magic factor). Blue points indicate crystals written after lowering the dose to 172 μ C/cm², but before obtaining personal developer glassware in August 2016; *w* magic factors exhibit some noise from write to write, but have a smaller mean, while *A* magic factors, and may also exhibit smaller process variance, though the time period over which these chips were written was comparatively short, meaning they are less susceptible to process drift. Magic factors for these data, defined as the difference of written and measured parameters, are given in tab. 3.3.



Figure 3.16: **Magic factors.** *Left:* written vs. measured width *w. Center:* written vs. measured amplitude *A. Right:* written vs. measured coupler width. In all plots black (I) points indicate crystals written at 230 μ C/cm², blue (II) points indicate crystals written at at 172 μ C/cm² and developed using common Painter group glassware, and red (III) points indicate crystals written at at 172 μ C/cm² and developed using bersonal glassware. Each point represents a single waveguide, with error bars representing one standard deviation of the measured values across different unit cells. Magic factors for these data are given in tab. 3.3.

| | Ι | II | III |
|---------------|---------------|-----------------|---------------|
| width w | (42 ± 11) | (-2.5 ± 19) | (-2 ± 13) |
| amplitude A | (21 ± 3) | (10 ± 3) | (9 ± 2) |
| coupler width | | (24 ± 12) | (24 ± 8) |

Table 3.3: **Magic factors** (written minus measured dimension) for w, A and coupler width. (I) corresponds to the black, (II) to the blue and (III) to the red points in fig. 3.16. Fewer micrographs of the coupler were taken for (I), so no estimate is provided.

Magic factors could be further reduced by dosing crystal and coupler differently; in the limit of infinite segmentation, this technique approaches PEC. Preliminary attempts were made to implement PEC for this project using the base dose determined in fig. 3.12 and simulated PSF (fig. 3.10), but did not result in global magic factor zeroing as I hoped. This is possibly the result of inaccuracies in the blur estimate (not captured by PENELOPE), PSF simulation, or base dose. If implemented properly (i.e., with the correct base dose and accurate PSF including blur), PEC could help reduce process variance. Additionally, as mentioned above, cold developer could help improve process stability and reduce the effects of process blur.

| Parameter | Setpoint | Value |
|---|----------|-------|
| Forward power (W) | 30 | 30 |
| ICP power (W) | 1000 | 969 |
| DC bias (V) | | 81-82 |
| C ₄ F ₈ flow (sccm) | 25 | 25 |
| SF ₆ flow (sccm) | 12 | 12 |
| Chamber pressure (mTorr) | 15 | 14.3 |
| Helium pressure (mTorr) | 10 | 10 |
| Helium flow (sccm) | | 5 |
| Temperature (°C) | 15 | 19–21 |
| Duration (s) | 180 | 180 |

Table 3.4: **RIE parameters.** Helium flow and DC bias are set indirectly by helium pressure and forward power, respectively. Temperature indicates the table temperature as reported by the Oxford software; its range represents the slow rise over the course of the etch.

3.4 Reactive Ion Etch

We transfer the patterns in the resist (both UV and EBL) to the chip by reactive ion etching (RIE).ⁱ Our etch is a mixed-mode RIE process, using SF₆ as the etch gas and C₄F₈ as the inhibitor. Chips are adhered to a four-inch silicon carrier wafer using thermal paste and loaded into an Oxford Plasmalab 100 ICP-RIE. Parameters of our etch are summarized in table 3.4. In contrast to many Painter group users, in our project we have chosen to track DC bias instead of forward power. The RIE is run in manual mode, allowing the user to adjust parameters mid-etch; we manually servo forward power to keep the DC bias in the 81–82 V range.

Afterwards chips are unloaded and removed from the carrier wafer. To remove the thermal paste, we saturate a cleanroom wipe with IPA, and push the chip along the wipe using tweezers. We then perform a solvent rinse and piranha clean to remove residual thermal paste and *e*-beam resist. Unfortunately, despite intensive cleaning, the thermal paste permanently discolors the nitride. This is fine as long as the device-layer nitride remains clean, i.e., is protected by *e*-beam or photoresist when patterns on the back side of the substrate are etched.

Empirically-determined etch rates for ZEP520A, silicon nitride and silicon are listed in tab. 3.5. These rates are determined by varying etch time and measuring etch depth using SEM images, shown in fig. 3.17. The selectivity (ratio of mask and silicon nitride etch rates) for our process is ca. 1.7. The parallel nanobeam etch

ⁱRIE is a rich and complicated subject and will not be covered in detail here; refs. [75] and [76] provide a good introduction.
| set Material | Etch rate (nm/s) |
|-----------------|------------------|
| ZEP520A | 1.3 |
| silicon nitride | 2.2 |
| silicon | 4.7 |

| Table 3.5: | Empiri | cally | determined | DRIE | etch | rates. | Rates | are | estimated | by |
|------------|-----------|--------|----------------|----------|--------|----------|-------|-----|-----------|----|
| measuring | etch dept | th for | different etc. | h times; | see fi | igure 3. | 17. | | | |



Figure 3.17: Etch test SEM. (a) No etch; (b) 90 s etch; (c) 180 s etch; (d) 270 s etch. Scale bar is 400 nm. Micrographs have false color; green corresponds to ZEP520A, fuchsia to silicon nitride, and gray to silicon. Etch depth between nanobeams and outside differ, illustrating the ARDE effect. By (d) the ZEP520A mask has completely eroded, causing the top surface of the nitride to be etched. Cross-sections were created by mechanically cleaving the chip.

test pattern exhibits aspect-ratio dependent etching (ARDE): the trench in between the nanobeams etches more slowly than on the exterior. The corrugated exterior of the photonic crystal may also cause ARDE, as depicted in fig. 3.18: troughs of the photonic crystal may have sloping sidewalls (up to ca. 20°) when compared with the peaks. This breaks *z* symmetry in the dielectric, which may lead to band mixing (see sec. 2.2).

3.5 Potassium Hydroxide Etch and Wet Chemistry

Potassium hydroxide solution (KOH) is an anisotropic etchant of silicon, meaning that different crystal planes etch at different rates [70]. It is useful in our process



Figure 3.18: **ARDE in crystal.** Focused ion beam (FIB) cross-sections of photonic crystal. Micrographs have false color: yellow corresponds to *e*-beam deposited tungsten (W), green to ZEP520A, fuchsia to silicon nitride (Si₃N₄), and gray to silicon (Si). The angle of the exterior sidewall apparently changes at different cross sections. (a) is the cross-section at thin part of crystal, (b) between and (c) at the thick part of crystal. Images are taken at the FIB angle of 52°. Scale bar represents 300 nm.



Table 3.6: **KOH etch rates.** Empirically determined etch rates for 30% KOH solution at 73 °C. Determined by monitoring groove depth and undercut. Rates are inferred by measuring patterned groove width *A*, trench width *B* and undercut *C* as etch progresses. Depth *D* is calculated using known angle for $\langle 111 \rangle$ direction.

for many reasons: the Si etch rate is higher than for RIE, allowing us to etch through hundreds of microns of silicon relatively quickly; it has an unmeasurably small etch rate for silicon nitride [70]; and the etch anisotropy allows us to create precise vee grooves in the silicon substrate to register an optical fiber to the device layer. Table 3.6 gives experimentally-determined etch rates for the $\langle 100 \rangle$ and $\langle 111 \rangle$ directions.

Chips are loaded into a PTFE caddy designed by Su-Peng Yu, based on an earlier design [62, ch. 2]. The caddy is designed to hold chips vertically, and has a removable

PTFE shield on its exterior, allowing it to be transferred from one liquid to another while keeping the chips submerged. In recent KOH etches, I have performed a solvent rinse and piranha clean immediately prior to ensure substrates were clean. A 30% KOH solution is heated to 73 °C and stirred at 200 RPM until temperature stable. The caddy is then submerged until the through window is completely open, typically ca. 150 min. Unfortunately, the etch rate is not perfectly uniform across all samples. This inhomogeneity may be the result of variations in KOH concentration or temperature across the solution, or the result of pattern/crystal misalignment—misaligned samples will undercut more quickly.

The effects of etch timing and sample alignment can be simulated: fig. 3.19a shows a simulationⁱ of the KOH etch for the mask shown in fig. 3.19b. This mask represents the region of the chip where the vee groove meets the window; truncating to this region expedites the simulation while still capturing the essential behavior. At a distance of ca. $300 \,\mu m$ from the edge of the window the side rails are perforated to accelerate the undercut. The first column of fig. 3.19a shows a properlyaligned substrate: as the etch progresses (top to bottom), the hole-induced undercut completes, producing an etch front at the first hole. This etch front progresses toward the end of the window, and, upon arrival, undercuts the corner silicon at the end of the vee-groove, resulting in a "flap" of suspended nitride.ⁱⁱ Proper timing allows the side rails to undercut, keeping the window clear, but minimizes the size of the flaps. If the patterns are not properly aligned, the etch front does not proceed symmetrically, resulting in one etch front reaching the vee groove before the other, as shown in the second and third columns, and in the microscope image in fig. 3.19c. Often one siderail per device will break—a study in fall 2016 revealed that this occured in about 40% of devices. The break is located within ca. 250 μ m from the edge of the window, and occurs during the KOH etch. The underlying mechanism is unknown. Commonly all waveguides on a given chip will have this break occuring in the same location, as illustrated in fig. 3.20, suggesting it may be related to substrate rotation. Most of the time waveguides with one broken siderail still exhibit normal spectral properties, and can be used in the experiment.

When undercut is complete, the PTFE shield is installed and the solution is transferred to heated DI water to rinse. In the water bath the shield is removed to dilute the KOH. Water dilution is performed twice. The sample is then transferred to

ⁱACES: Anisotropic Crystalline Etch Simulator, for Windows 95.

ⁱⁱActually, the astute observer may notice that the corner undercuts *before* the etch front arrives; this timing is sensitive to the etch rate for the $\langle 111 \rangle$ direction and hole size, but may very well happen.



Figure 3.19: **KOH etch simulation.** (a) Total simulated etch time increases from top to bottom, corresponding to 45 m, 90 m, 135 m and 180 m, respectively; and (clockwise) rotation angle increases from left to right. Blue rectangle indicates planned location for optical fiber, white indicates undercut nitride and fuchsia indicates nitride in contact with silicon (i.e., not undercut). (b) Diagram of mask used in (a), highlighting the presence of holes in the side rail. Blue scale bar represents $300 \ \mu\text{m}$. (c) Optical microscope image of the effects of rotation. Pattern/crystal rotation is ca. 0.3° , resulting in an asymmetric undercut. Region pictured is between vee grooves.



Figure 3.20: Siderail failure. Job #22863 (Hedda C), devs. 1–3.

Cyantek Nano-Strip® heated to 160 °C for 20 min. This is to clean organic residues without the added turbulence of piranha solution. The samples are then again twice rinsed in water.

In November of 2014 we added a buffered hydrofluoric acid (BHF) solution to our process, in response to an unidentified contaminant, shown in fig. 3.21. We believe this material is silicon dioxide. It is certainly not organic (it did not respond to Nano-Strip, piranha or O_2 plasma), and, while much smaller than that of silicon, silicon dioxide has non-zero etch rate in KOH—at our process temperature it is around 1/200th of the rate for $\langle 100 \rangle$ silicon [70]—and any glass components we use repeatedly will etch appreciably. In particular, the glass dish used for the KOH etch will eventually show signs of aging and need to be replaced. The BHF solution we use is very dilute (1:40 BHF:H₂O) and exposure is limited to 2 min., to minimize the etch rate for silicon nitride while maintaining an appreciable rate for silicon dioxide—the selectivity for BHF solution of silicon dioxide over silicon nitride is ca. 100:1, and, at a ratio of 1:5 BHF:H₂O, the reported etch rate for LPCVD silicon nitride is less than 1.3 nm/min. [77].

After removing from the BHF solution, the sample is rinsed with water three times.



Figure 3.21: **Importance of BHF.** (a) Before BHF. Micrograph shows the end of the vee groove, where waveguide is suspended over window; scale bar is 25 μ m. Contaminant is apparent. (b) Same region as in (a), but after BHF. Contaminant has been removed. (c) KOH glassware. Bottom of right-most dish shows obvious signs of etching, and should be replaced.

The water is then diluted three timesⁱ with chromatography-grade IPA, and transferred to a critical point drying tool (CPD), which removes the liquid without surface tension or evaporation, which can damage delicate structures. After the CPD, undercut samples are removed from the caddy and cleaned in an oxygen plasma.

Some chips have been successfully released without the CPD, instead using a slowlydraining beaker. This technique was explored in fall of 2015. At the time it was believed that avoiding the CPD resulted in cleaner, lower-loss structures, evidenced by quality factor measurements of fishbone cavities [48]. This technique was not adopted as standard procedure, due to reduced yield: only about 30% of the waveguides survived, as broken siderails would often wrap around the waveguide, rendering it unusable.

ⁱI check the dilution visually—when removing the shield produces no more visible BHF, dilution is adequate.

3.6 Atomic Layer Deposition

The final fabrication step is atomic layer deposition (ALD) of a thin, conformal alumina layer. ALD was initially developed for use as a dielectric layer in thin film electroluminescent displays [78], but in recent years has been employed as a fabrication technique in nanophotonics [79]. The technique has been used to tune spectral features in photonic crystals and photonic crystal cavities [80–84]; to heal surface roughness, reducing propagation loss in waveguides [51, 85]; and to improve quality factors of cavities [86]. Additionally, ALD-deposited alumina has been shown to protect against alkali corrosion [87, 88]. In our experiment ALD allows us to fine-tune the spectral features of the structure after optical characterization. We choose to use alumina because of its similar refractive index ($n \approx 1.75$), low optical loss, and the potential to protect our structures against cesium corrosion, discussed in sec. 4.6.

All ALD processes consist of two precursor chemicals. To make alumina, these precursor chemicals are trimethylaluminum (TMA) and water vapor, which react to form alumina (and methane gas) according to [89]:

$$\operatorname{Al}(\operatorname{CH}_3)_3 + \frac{3}{2}\operatorname{H}_2\operatorname{O} \to \frac{1}{2}\operatorname{Al}_2\operatorname{O}_3 + 3\operatorname{CH}_4.$$

In an ALD process (illustrated in fig. 3.22), a sample is exposed to one of these precursors in a heated reactor at low pressure (around 0.3 torr). This precursor is able to bond to atoms on the surface, but is limited to forming a single monolayer. After saturating the surface, the reactor is purged with an inert gas (argon in our case), and the second precursor is allowed to react with the chemisorbed monolayer, but also forms precisely one additional monolayer. By cycling precursors and purge, atomically-thin layers can be formed on the substrate, allowing very precise control.

For our process (250 °C), a single cycle produces an alumina layer with a thickness of $\theta \approx 0.11$ nm. Generally we target around 100 cycles (11 nm) of ALD-alumina, resulting in an adlayer of comparable thickness to those in ref. [88], with the hope that it will slow degradation due to cesium exposure. The empirical tuning rates for this process are given in sec. 4.2.



Figure 3.22: **Illustration of ALD. (a)** Reactor is pulsed with water vapor, exposing sample (gray) and resulting in a layer of hydroxyl groups. Unreacted water vapor is then purged. **(b)** Reactor is then pulsed with TMA, which reacts with the hydroxyl groups, expelling methane. Unreacted TMA is purged. **(c)** Reactor is again pulsed with water vapor, which reacts with the surface, expelling methane. Cycle is repeated, controllably growing a conformal adlayer of alumina.

Chapter 4

CHARACTERIZATION AND INSTALLATION

In this chapter I describe the characterization of our APCW waveguides, summarizing their properties and providing comparisons to models presented in ch. 2. I further describe the procedure used to package chips and install them into the atomic physics setup, and discuss the effects of cesium exposure on the waveguides.

4.1 Characterization Setup and Spectra

Optical characterization of devices is performed primarily using a fiber-coupled broadband superluminescent diode (SLED),ⁱ covering the range of 815-910 nm (330–368 THz). Finer spectral detail can be obtained by using diode laser sources near the cesium D1 and D2 transitions,ⁱⁱ or a Ti:sapphire laser.ⁱⁱⁱ Sources are polarized on a polarizing beamsplitter cube (PBS), and an achromatic quarter and half waveplate are used to compensate for downstream birefringence produced by bends in the fiber, allowing the excitation of specific modes in the waveguide. Exciting the TE mode requires the light exiting the fiber be polarized horizontally (in the device layer plane); the TM mode is excited when the light is polarized vertically. These polarizations are distinguished by monitoring scattering into the z direction—the coupler tether scattering into this direction is minimized when the polarization is aligned to the TM polarization, and a 45° rotation of the half waveplate aligns the polarization to TE. Scattering is monitored using a microscope and NIR-sensitive camera. After polarization control, light is sent through a 50/50 fiber beamsplitter—one port of the beamsplitter is sent on to the waveguide, and another port allows us to monitor reflection. A schematic is shown in fig. 4.1.

In practice, the wavelength dependence of the fiber birefringence [90] results in slightly different polarization across the spectrum when illuminating with a broadband source. Fig. 4.2a shows the extinction ratio (ER) of the dark port of a PBS for light exiting the polarization control setup. Red data show the ER with the quarter and half waveplates aligned to minimize *total* transmission, as measured by a photodiode. In the part of the spectrum relevant to our experiment this method results

ⁱInphenix Model IPSDM0822-0318

ⁱⁱNewport TLB-6716 and TLB-6718, respectively.

iiiM2 SolsTiS®



Figure 4.1: **Diagram of polarization control setup.** Dot-dashed mirror is long-pass dichroic at 875 nm, and dotted mirror is on a removable magnetic base.

in ≤ 20 dB extinction. The ER was then optimized in a narrow bandwidth around 338 THz using the OSA, shown in blue, which recovers the ca. 30 dB ER expected for a PBS. In these data the ER degrades by 0.32 dB/THz about the optimum, though in general this should depend on how much "twist" is present in the fiber. Polarization filtering of light transmitted through an APCW shows ca. 5–10 dB improved extinction in the band gap (fig. 4.2b), though this will also suppress birefringence introduced in the APCW.

We use Nufern 780HP single-mode fiber to couple light into the waveguides. To reduce reflections at the fiber facet, we use fibers with an anti-reflective coating.ⁱ Fiber positioning is achieved by mounting fibers in vee groove fiber holders mounted on five-axis micrometer positioners. During characterization the chip is held to an aluminum mounting block using an SEM spring clip. Fibers are mounted at a ca. 5° incline. At a safe distance from the coupler, the inclined fiber is lowered into the groove until the fiber begins to cant upwards, and then raised again until the fiber lies flat in the groove. The Rayleigh length for the mode launched by the fiber is ca. 20 μ m, and the fiber must be positioned within this distance from the coupler to obtain good coupling. As the fiber is brought toward the coupler, interference fringes appear in the microscope; the spacing of this pattern can be used as feedback for the fiber–coupler distance. An image of the fiber positioning apparatus is shown in fig. 4.3.

When using the SLED, transmission or reflection spectra are obtained using an

ⁱOz Optics; R<0.3% in 790–890 nm.



Figure 4.2: **Spectral polarization variation of SLED.** (a) Extinction ratio for PBS. Quarter and half waveplates were first adjusted to minimize transmitted light as measured by a photodiode (red). Extinction ratio (ER) was then optimized for a narrow bandwidth (blue). Blue lines are fits to optimized ER about optimum, with slope ca. 0.32 dB/THz. (b) Transmission spectra for APCW with (blue) and without (red) PBS filtering before OSA.



Figure 4.3: **Characterization setup.** False color image of fiber positioning apparatus. Green highlights the aluminum mounting block, blue the vee groove fiber holders, red the microscope objective, and purple the "glue bot" applicator. Each positioning element is mounted on a stack of micrometer positioners.



Figure 4.4: **OSA spectra.** From top to bottom: transmission spectrum; reflection spectrum; inferred group index using eq. 2.9 and fit; and reconstructed band diagram. Around 364 THz, transmission goes to zero, corresponding with the intersection of the air band with the light line.

optical spectrum analyzer (OSA).ⁱ For the tunable diode laser sources, the laser frequency is swept and transmission and reflection signals are recorded on a photoreceiver.ⁱⁱ To account for losses in the optical system, reflection spectra are normalized to a flat cleaved fiber, assumed to have a reflection of 3.5% due to the index mismatch between optical fiber and air; transmission spectra are normalized by splicing input to output, bypassing the chip. Examples of transmission and reflection spectra are shown in fig. 4.4.

Many elements in our waveguide produce unintentional reflections, which manifest as fringes in transmission and reflection spectra. Each pair of reflectors can be considered an etalon with free spectral range (FSR) Δv [91];

$$\Delta v = \frac{c}{2nL},$$

where c is the speed of light in vacuum, and the product of refractive index n and etalon length L gives the optical path length. Using this relation and the Fourier transform of OSA spectra allows us to relate fringes to optical path lengths, shown

ⁱAnritsu MS9740A

iiNewport 2051-FC



Figure 4.5: **PSD of reflection and transmission spectra.** Power spectral density (PSD) vs. optical path length. Red curve shows PSD for reflection, blue for transmission. Gray highlighted regions represent peaks associated with crystal and coupler–crystal lengths.

in fig. 4.5. There is a broad peak around $nL = 140-320 \ n \cdot \mu m$ which we can associate with the chirping band edge resonances; assigning a crystal length 150a =55.5 µm, this is consistent with a group index $n_g = 2.5-5.8$. In reflection, there is another prominent pair of peaks associated with an etalon with $nL \approx 2.7 \ n \cdot mm$, corresponding to a length $L \approx 1.4 \ mm$, the distance between coupler and crystal, with n = 2. We can associate this etalon with the interference of backscattering off the coupler tether and reflection from the photonic crystal. Free-space coupling schemes enable different geometries for the coupler tether, can be used to suppress this etalon [48].

Spectra can also reveal the presence and features of other symmetries. Figure 4.6 shows both transmission spectra for both the TE (red) and TM (blue) modes. The frequency-dependent coupling efficiency for TM mode is evident, and results from coupler geometry optimized for TE (see sec. 2.5). The TM mode can be deliberately excited by rotating the polarization of light at the coupler; if transverse symmetries are broken in the waveguide, light can also couple to unintended modes. A dramatic example is shown in fig. 4.7: a scatterer in the photonic crystal taper projects the input mode into both *y*-even and *y*-odd symmetries, allowing the *y*-even-*z*-even dielectric band to be observed in reflection spectra. Features inside the band gap are particularly noticable, but signatures of crosstalk are also present in the frequency range above the band gap, where there is a confluence of modes. Symmetries other



Figure 4.6: **TM and TE transmission spectra.** Red trace shows TE transmission, while blue shows TM. TM transmission remains high throughout TE band gap. Increased coupling efficiency at lower frequency is consistent with mode overlap calculations in sec. 2.5.

| | $\delta(f_0)$ | $\delta(\Delta f)$ |
|------------|---------------|--------------------|
| δw | -320 ± 10 | -130 ± 20 |
| δA | 70 ± 20 | 270 ± 30 |

Table 4.1: **Tuning rates of band gap with geometry (empirical).** Shift of band properties with w, A in GHz/nm. Coefficients determined using OSA data for 24 different substrates with scans of A or A and w.

than the fundamental TE and TM are only weakly guided in the single nanobeam leading up to the photonic crystal, so crosstalk with these modes manifests dominantly as loss.

Empirical Tuning Rates and Coupling Efficiencies

Aggregate spectral data is used to determine empirical tuning rates for f_0 and Δf relative to width and amplitude and given in tab. 4.1. Values are determined using OSA spectra from one- and two-dimensional parameter scan chips; tuning rates are assumed to be linear in each scanned parameter, and average fit coefficients are reported. Fits use *written* (i.e., CAD-specified) dimensions. Coefficients here differ only slightly from those determined by simulation (see tab. 2.1).

Aggregate data also reveal the change in coupling efficiency with coupler width. Figure 4.8 shows the mean band gap reflectivity as a function of written (CAD) coupler width. The simulated round-trip coupling efficiency in our operating band-





Figure 4.7: Scatterer coupling to other symmetries. *Above*: micrograph showing taper of APCW and a prominent scatterer attached to the waveguide. Scale bar represents 1 μ m. *Bottom left*: reflection spectrum when excited from the left. The *y*-odd, *z*-even band gap spans 338–354 THz, but there is a prominent intra-band gap feature around 348 THz, consistent with the location of the *y*-even, *z*-even dielectric band edge (see fig. 2.5). Features above 360 THz may show signs of interaction with quasi-TM modes as well. *Bottom right*: reflection spectrum when excited from the right. Intra-band gap feature is not apparent, as light from the right is attenuated before reaching scatterer.

width is shown below for comparison. Recall from tab. 3.3 that the written coupler width exceeds the true, measured, coupler width, indicating that the written optimum should occur at a larger value than simulations predict. Indeed, the offset between optimal widths for simulated and measured efficiencies is well captured by the measured "magic factor" offset between written and measured coupler dimensions (ca. 25 nm) reported in tab. 3.3. Downstream inefficiencies in our millimeter-length waveguide account for differences in coupling efficiency [48]. The preponderance of data points for couplers with widths smaller than optimal is not accidental—devices are designed to have ALD-deposited alumina claddings with significant thickness, requiring initially undersized couplers.ⁱ

ⁱIt would also be worthwhile to validate coupling efficiency as a function of *measured* coupler width. Measured dimensions are not used here because SEM images of the coupler do not exist for every waveguide, reducing the size and statistical significance of that dataset.



Figure 4.8: **TE coupling efficiency vs. coupler width.** *Above*: mean band gap reflectivity \overline{R} vs. *written* coupler width. Couplers are deliberately under-sized to accommodate ALD (described below). Dashed line is scaled and shifted simulation curve, provided as a guide for the eye. *Below*: mean, round-trip TE coupling efficiency $|\eta|^2$, averaged over 330–360 THz (see fig. 2.26 and surrounding discussion). Peak offset between plots is consistent with ca. 25 nm magic factor for coupler width (see tab. 3.3).

Correspondence of Simulated and Measured Band Placement

We can determine the accuracy of simulations by comparing empirical and simulated band edge placement. Figure 4.9a shows an SEM of an APCW with parameterizing dimensions indicated. Empirical dimensions are obtained by assuming the lattice constant *a* is exactly 370 nm, and then measuring several unit cells to determine a mean value and uncertainty, which is typically on the $\approx 2\%$ level. It is also assumed that pixels are square; this was confirmed by imaging the same sample horizontally and vertically. These dimensions can be fed into an eigenmode solver to obtain band structure (see sec. 2.2), and compared with experimentally-obtained spectral data. Here I use the positions of the first dielectric and air band resonances as a proxy for the location of the band edge, which slightly overestimates Δf .

Figure 4.9b compares simulated and measured band structure for several different APCWs on different chips. Here I have used *measured* dimensions (those obtained via SEM) as inputs to the simulation, instead of *written* dimensions, meaning that,



Figure 4.9: **Simulated and measured band placement.** (a) SEM image with dimensions indicated. For APCW geometry, we assume a = 370 nm to determine scale, as the EBPG is believed to be better calibrated than SEM. (b) Simulated vs. measured band placement. Each color represents a different APCW on a different chip (see table below for fabrication details). APCW pictured in (a) is Holde B, dev. 9. \circ represents simulated band positioning using measurements from SEM images, and \times the positions of the first resonances on air and dielectric bands. Solid lines show changes to simulated band structure within measurement uncertainty: magenta and orange lines indicate changes in *w* and *A*, respectively; cyan represents ± 10 nm changes to *t*.

unlike in fig. 4.8, we need not account for the magic factor. There is an uncertainty in the measured dimensions, whose effect on band structure is here indicated by solid colored lines: magenta and orange lines represent the effect of measurement uncertainty in w and A, respectively, while the cyan line indicates the effect of a ± 10 nm uncertainty on the silicon nitride thickness about the assumed thickness t = 200 nm. There is a variable discrepancy across different chips, though in all cases the simulation places band structure at lower frequencies (redder) than is measured. In most cases small changes to geometry cannot entirely account for the discrepancy.

Inhomogeneity of Nominal Duplicates

Comparing nominally identical waveguides—those with the same CAD-specified dimensions—can help identify potential sources of variance. Figure 4.10 shows two sets of nominally identical waveguides: reflection spectra on the left come from two consecutively-written, consecutively-etched substrates, while transmission spectra on the right come from identical devices on the same substrate. The band gaps for devices on separate substrates are shifted with respect to each other by ca. 2.3 THz, while the band gaps on the same substrate exhibit much smaller disparity, only about 0.5 THz.

More recent patterns do not duplicate crystal dimensions on a single chip—instead, each pattern has a parameter scan, maximizing the possibility that some of the devices on the chip are in the correct spectral location. Figure 4.11 shows the relative spectral positioning for a parameter scan chip. The scan is square in w-A space, with a 5 nm step in written geometry between adjacent devices, corresponding to a parallelogram in $f_0-\Delta f$ space, with some process noise. ALD thickness was chosen conservatively, ensuring that no devices "overshoot" the cesium D₁ frequency (black dashed line).

On many occasions multiple prints of the same pattern file were made. Figure 4.12 compares relative spectral locations of duplicate prints, revealing a sometimes large but directional disparity. The direction is very similar to that of the simulation disparity (fig. 4.9). Each pair of substrates was written in the same beam write, so EBL parameters (beam current, spot size, etc.) should be very similar, as should RIE conditions.

There are several possible reasons for the spectral inhomogeneity of nominally identical prints. One potential cause is a variance in the silicon nitride film thickness.



Figure 4.10: **Comparing nominally identical waveguides.** In each column the waveguide above and below are nominally identical, and lines are aligned with the first resonance on either band edge, with red indicating the resonance position of the top waveguide and blue indicating the resonance position of the bottom waveguide. (**a–b**) Reflection spectra from two separate substrates. Waveguides are from consecutive prints (jobs 21063 and 20164), and show a ca. 2.3 THz frequency offset. Coupling efficiency is also notably different. (**c–d**) Transmission spectra from the same substrate (job 20597). They are offset in frequency by ca. 0.5 THz.



Figure 4.11: **APCW parameter scan.** Midgap frequency f_0 and stopband width Δf for job # 21063 (Lil F) after ALD. Written width *w* and amplitude *A* are programmed as a square array, with step size 5 nm. Blue arrow indicates mean displacement of 5 nm change in *A*, and red arrow 5 nm change in *w*. Black dashed line indicates cesium D₁ frequency.

SVM specifies a $\pm 7\%$ accuracy for silicon nitride film thickness, and private correspondence suggests the variance in film thickness across a single wafer and in a wafer run may also be this large. According to tab. 2.1, a 2 THz shift in f_0 corresponds to ca. 8 nm difference in film thickness, which should be observable by eye (see appendix C); significant color variation across nominally identical prints was not observed. Additionally, the simulated shift of band structure with film thickness does not predict a significant change in Δf , though we have already seen that empirical and simulated shift rates are not in perfect agreement.

Comparing the dominant shift rate to the simulated shifts listed in tab. 2.1 suggests an isotropic growth or shrinkage of dimensions would also account for the disparity. For example, a 3 nm ALD-alumina cladding would account for a 2 THz discrepancy in the f_0 well. However, it's difficult to attribute this type of deformation to any processing step—the only known isotropic silicon nitride etchant to which the chips are exposed is BHF, and the concentration and timing used corresponds to an etch that is too small to account for such a discrepancy.



Figure 4.12: **Comparing nominally identical prints.** Each arrow represents the displacement in band structure between a device and its nominal duplicate partner on a different substrate. Each color represents a different GPF file. Each pair of substrates was written in the same EBL session (see table). The slope of the shift is approximately $\delta[\Delta f]/\delta[f_0] \approx 0.23$. Dashed line is a guide for the eye, indicating the dominant shift direction.

4.2 Atomic Layer Deposition

Regardless of the underlying cause, the spectral variance of nominally identical prints is a serious impediment to producing devices with a band structure aligned to the cesium D lines. A common strategy of dealing with process variance is simply to produce a parameter scan with enough devices that some subset, statistically, will have the desired properties. Unfortunately, our fiber-coupling strategy limits the efficacy of this strategy, as the linear device density is constrained by the diameter of an optical fiber, limiting the number of devices per chip to ca. 20.

A second strategy is to produce devices, characterize them, and the perform a postcharacterization tuning process to correct their spectral properties. As devices on the same chip exhibit reduced spectral spread, global post-characterization tuning can result in many usable devices. Before the ALD tool, Jonathan Hood and I attempted



Figure 4.13: **ALD shift.** *Left*: change in f_0 and Δf per ALD-alumina cycle. Empirically determined shift rates for f_0 and Δf are $-(86 \pm 7)$ GHz/cycle and $-(18 \pm 2)$ GHz/cycle, respectively. For 0.11 nm-thick monolayers, these rates are in reasonably good agreement with tab. 2.1. *Right*: comparison of reflection spectra before ALD (red) and after 115 cycles (blue), illustrating spectral redshift and increased coupling resulting from improved mode overlap. Reflection spectra from job # 23062 (Heide B) dev. 11.

to develop a controllable, isotropic RIE process to achieve spectral tuning. We developed a CF₄-based etch, which shifted the band structure around 1 THz/min. One generation of devices installed in the atomic physics experiments was tuned this way, but the process is very time sensitive, and does not result in uniform tuning across the chip. More information about etch rates and chemistry for silicon nitride in CF₄ can be found in ref. [92].

Our current post-characterization tuning technique relies on on atomic layer deposition (ALD), described in sec. 3.6. ALD is the finest spectral control we have for our structures, and is probably also the most reliable step in our fabrication process. Empirical tuning rates are summarized in tab. 4.13, and are in good agreement with the simulated tuning rate (tab. 2.1).

Unfortunately, ALD results in a global geometric modification, meaning we must compensate the coupler dimensions to maintain high coupling efficiency. Also, because adlayer thickness constitutes a single degree of freedom, we are not able to independently tune dielectric and air bands. Commonly we choose to align the dielectric band edge to the cesium D_1 transition, as both the atomic level and APCW band structures are simpler [93]. While the ALD tool is able to narrow the spread of band placement across substrates, producing a substrate with several acceptable

APCWs is still primarily a matter of producing *many* parameter scan chips. Knowing the spectral inhomogeneities are directional (fig. 4.12) along roughly the same axis as ALD allows one to strategize by writing parameter scans orthogonal to this axis, i.e., dominantly scans of Δf .

4.3 Image Data

Spectral data, described above, is the primary characterization tool for the APCW. Complementary information can be obtained through images obtained by optical and electron microscopy, described below.

Optical Microscopy

Scattered light from the waveguide provides data complementary to reflection and transmission spectra. Image data has been used to empirically recreate dispersion relations [32]. Figure 4.14a shows a composite image of a waveguide under broad-band illumination (330–368 THz). Scattering is apparent at the input (red arrow) and output couplers at the extremes of the waveguide, and at the beginning of the photonic crystal (blue arrow). A large portion of the broadband source lies in the stopband of the APCW resulting in significant reflection, causing scattering to occur dominantly on the input side of the APCW.

A tunable light source can be used to observe frequency-dependent scattering. Figure 4.14b shows scattering from the APCW when the excitation is tuned to bandedge resonances. While supermode structure can partially be made out, it's clear that some elements in the waveguide scatter more than others—in this data set there is a very large scatterer at the left of the crystal, and many other point scatterers can be seen on the crystal interior. Appropriate normalization allows us to reconstruct supermode structure (see sec. 2.4). In the transfer matrix model, integrating the frequency-dependent intensity profile I(x, f), starting at the band edge frequency f_{BE} gives

$$p(x;\delta) = \int_{f_{\rm BE}}^{f_{\rm BE}+\delta} \mathrm{d}f \, I(x,f),$$

which gives an increasingly flat profile as more resonances are included ($\delta \rightarrow \infty$), as evidenced in fig. 4.14c. Performing an analogous integral over the acquired image data provides a reference against which scattering can be normalized.



Figure 4.14: Intensity profiles from image data. In all scattering images tone has been inverted for improved print quality, i.e., black represents highest scattering. (a) Scattering from entire crystal. $L \approx 2.6$ mm. Red arrow indicates scattering at coupler, blue arrow scattering at beginning of APCW. (b) (i)–(iv) show scattering from crystal when excited at 1st, 2nd, 3rd and 4th dielectric band resonances, respectively. Some parts of the waveguide scatter more light than others, leading to uneven sampling of the intensity profile. Normalizing against integrated intensity profiles enables supermode reconstruction. (c) Integrated intensity profiles using the TMM (see fig. 2.22). Lowest curve is integral over 1st resonance, next over 1st and 2nd resonances, etc. As more resonances are included profile is increasingly "flat," allowing us to normalize scattering.

Electron Microscopy

Many critical features in our structure require cannot be resolved by optical microscopy. For this reason we use a scanning electron microscope (SEM) to image our structures. SEMs are our dominant source of information about APCW and coupler geometry, etch profiles, and device cleanliness. A practical introduction to SEM can be found in ref. [94], while ref. [66] provides more comprehensive information. All SEM images presented in this work were acquired with SEMs in the KNI: the Nova NanoLab 200 and Nova NanoLab 600. While images are acquired both on substrate and after KOH etch, I've found it difficult to distinguish the silicon nitride and silicon boundaries for devices on substrate when imaging from above. For this reason, whenever possible quantitative information is acquired from images after KOH.



Figure 4.15: **SEM settings.** Image (a) was acquired using "standard" settings in this work: 2 kV acceleration voltage, 13 pA beam current and imaged using the "through-the-lens" detector (TLD). These settings provide good surface detail and minimize charging artifacts. Images (b–d) show the same region as in (a), but imaged using the Everhart Thornley detector (ETD) at 2 kV, 3 kV, and 5 kV, respectively. Blue and red arrows indicate detritus on the waveguide; the visibility of these features changes with acceleration voltage. Image (e) shows a wider view of the same sample imaged at 10 kV; at this acceleration voltage, charging is sufficient to cause the beams to repel.

While SEM techniques are quite standard and modern microscopes are very userfriendly, some features of our waveguides make them difficult to image: silicon nitride is an electrical insulator, which can result in charging, producing artifacts in SEM images. At sufficient acceleration voltage, sample charging can be severe enough to produce mechanical effects: at 10 kV and high magnification, charging can cause the nanobeams to repel. Conductive coatings can be used to mitigate charging [94], but would ruin optical properties of our structures. I have found that charging effects can be reduced by imaging at low (ca. 2 kV) acceleration voltage; fig. 4.15 shows a comparison of the same sample imaged under different conditions.

The extreme aspect ratio of our structure makes imaging challenging: in an image wide enough to capture the entire structure, one pixel corresponds to about 50 nm, which is too coarse to obtain useful information. However, by "stitching" many images at higher magnification together into a composite, one can collect information about the entire structure. Figure 4.16 shows edge data extracted from such a sequence of images (ca. 5 nm/pixel resolution), and extracted crystal dimensions. Images were median filtered and binarized using ImageJ [74], and stitched together in MATLAB using a custom script. Reflection spectra for this device are shown in

fig. 4.17.

These data show that width w and amplitude A remain relatively uniform across the nominal length of the crystal (excluding tapers), except at one location around $x = 45 \ \mu\text{m}$ where there is a defect on the interior sidewall; top and bottom beams exhibit relatively small discrepancies between upper and lower nanobeams ("good y symmetry"). Linear fits reveal small slopes along the nominal crystal: $\Delta w / \Delta x =$ $(0.1 \pm 0.06) \ \text{nm/}\mu\text{m}$, $\Delta A / \Delta x = -(0.03 \pm 0.03) \ \text{nm/}\mu\text{m}$, and $\Delta g / \Delta x = -(0.1 \pm$ $0.02) \ \text{nm/}\mu\text{m}$. g exhibits large discrepancy between written (dashed) and measured dimensions, which is more pronounced in the tapering regions. This is not the result of lithographic error, but rather a mechanical effect. Some efforts have been explored to address this issue (see ref. [48]), though any geometrical modification to address mechanical effects necessarily affects optical impedance matching as well (see fig. 2.9).

4.4 Science Chips

According to the EBPG job manager (jman), since April 2015 I have logged more than 300 hours of beam time, resulting in 182 APCW prints. Installing a sample in the vacuum chamber requires quite a lot of effort, meaning the bar is quite high for any chip to be installed: several waveguides on the same chip must have the desired spectral properties. Unfortunately, this is a rare occurrence, due to mechanical failures, poor spectral placement, poor coupling, or other non-idealities. Still, a subset of these chips have been deemed sufficiently attractive to be installed. In this section I present the spectral data for these *science chips*.

Table 4.2 shows the five chips selected for installation in Lab 2 between June 2016 and January 2017. The columns of the table show the EBPG job number, chip "nickname," the EBPG write date, KOH release date, number of cycles of ALD deposited on the chip, date of "butt-coupling" for the experiment (see sec. 4.5), device number, and spectral properties: the frequency of the first dielectric (f_{diel}) and air band (f_{air}) resonances, one-way power transfer efficiency for front (η_f) and back (η_b) couplers, and transmission $T = \eta_f \eta_b$. One-way efficiencies are inferred via bandgap reflection, i.e. $\eta_i = \sqrt{R_i}$, and incorporate both the fiber–coupler power transfer efficiency $\eta_{1\rightarrow 2}$ (sec. 2.5) and all other losses between coupler and APCW. Reported uncertainties represent one standard deviation of transmission or reflection fluctuations, due to coupler etalon (sec. 4.1) and frequency dependence of mode overlap across the spectrum.



Figure 4.16: Geometric uniformity across crystal. In all plots red indicates bottom beam, blue top beam, and dashed black line "written" dimension. (a) Edges of photonic crystal, acquired from composite SEM image. (b) Fitted w across length of crystal; $w = (291 \pm 4)$ nm and (290 ± 3) nm for top and bottom beams, respectively. (c) Fitted A across length of crystal. $A = (121 \pm 2)$ nm for both beams in nominal crystal. w and A are extracted by fitting each period of the crystal to a sinusoid. (d) Fitted g across length of crystal. $g = (216 \pm 5)$ nm. Analysis performed on job # 22773 (Hilde B), device 1; images were acquired before ALD process.



Figure 4.17: **Reflection spectra for Hilde B, dev. 1.** (a) Before ALD, and (b) after 65 cyc. ALD. Red (blue) lines indicate cesium D_1 (D_2) transition frequency.

ALD is typically done in multiple steps, which here is represented by addition: 70+27 indicates 70 layers were initially deposited, and after further characterization an additional 27 were deposited to trim the band edge to the correct location. All structures are roughly aligned to the cesium D₁ transition ($f_{D_1} = 335.1$ THz). On parameter scan chips, band edges for separate APCWs on a given print are not aligned, meaning compromises must be made: tuning f_{diel} to f_{D_1} for one APCW may place f_{D_1} in the band gap of another. Additionally, it is desirable to position the f_{diel} at a small blue detuning from f_{D_1} , anticipating redshift resulting from cesium exposure (see sec. 4.6). Param files for all chips discussed in this section are included electronically at CaltechTHESIS.

June–July 2016

Three chips were butt-coupled in summer 2016. The standard procedure for butt coupling (described in sec. 4.5) involves a heat-curing epoxy, requiring the temperature of chips to be elevated to ca. 110 $^{\circ}$ C. Unfortunately, at the time these chips were processed we encountered problems related to heating, including substrate cracking and waveguide stiction, now attributed to chemical volatiles evaporating off the heated mounting block and condensing on the chips. Consequently, the

| Т | 0.18 ± 0.02 | 0.13 ± 0.02 | 0.14 ± 0.02 | 0.18 ± 0.02 | 0.21 ± 0.01 | 0.19 ± 0.01 | 0.21 ± 0.02 | 0.22 ± 0.02 | 0.18 ± 0.02 | 0.18 ± 0.02 | 0.15 ± 0.02 | 0.19 ± 0.02 | 0.21 ± 0.02 | 0.20 ± 0.02 | 0.20 ± 0.02 | 0.17 ± 0.01 | 0.13 ± 0.02 | 0.22 ± 0.03 | 0.23 ± 0.03 | 0.23 ± 0.03 | 0.20 ± 0.03 | 0.19 ± 0.03 | 0.20 ± 0.03 | 0.19 ± 0.02 |
|------------------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|
| $\eta_{ m b}$ | 0.36 ± 0.03 | 0.26 ± 0.04 | 0.30 ± 0.03 | 0.35 ± 0.03 | | | | | 0.42 ± 0.03 | 0.39 ± 0.03 | 0.31 ± 0.03 | 0.41 ± 0.03 | 0.32 ± 0.03 | 0.33 ± 0.02 | 0.38 ± 0.03 | 0.47 ± 0.04 | 0.32 ± 0.04 | 0.44 ± 0.04 | 0.41 ± 0.04 | 0.42 ± 0.04 | 0.40 ± 0.04 | 0.36 ± 0.04 | 0.37 ± 0.04 | 0.37 ± 0.03 |
| η_{f} | 0.49 ± 0.03 | 0.49 ± 0.04 | 0.47 ± 0.04 | 0.51 ± 0.04 | | | | | 0.43 ± 0.03 | 0.46 ± 0.03 | 0.48 ± 0.03 | 0.47 ± 0.03 | 0.53 ± 0.03 | 0.53 ± 0.04 | 0.49 ± 0.03 | 0.42 ± 0.03 | 0.42 ± 0.04 | 0.51 ± 0.04 | 0.55 ± 0.04 | 0.54 ± 0.04 | 0.51 ± 0.04 | 0.52 ± 0.04 | 0.54 ± 0.04 | 0.51 ± 0.04 |
| f _{air} (THz) | 349.8 | 352.0 | 349.0 | 350.1 | 349.8 | 351.0 | 352.0 | 350.5 | 350.1 | 348.0 | 350.5 | 351.3 | 350.0 | 350.6 | 350.5 | 350.4 | 352.1 | 352.1 | 352.4 | 353.0 | 350.1 | 350.8 | 350.6 | 349.9 |
| $f_{\rm diel}$ (THz) | 337.4 | 336.8 | 335.9 | 335.8 | 337.2 | 337.1 | 336.8 | 335.7 | 337.5 | 336.2 | 335.9 | 335.5 | 335.9 | 336.2 | 335.9 | 335.9 | 336.9 | 336.8 | 336.8 | 336.9 | 335.7 | 336.0 | 335.6 | 335.1 |
| Device | 10 | 12 | 15 | 16 | 10 | 11 | 12 | 16 | S | 6 | 11 | 12 | S | 7 | 8 | 11 | S | 9 | 7 | 8 | 6 | 10 | 11 | 14 |
| Glue | 2016-06-27 | | | | 2016-07-06 | | | | 2016-07-13 | | | | 2016-11-18 | | | | 2017-01-23 | | | | | | | |
| ALD (# cyc.) | 70+27 | | | | 100+18 | | | | 30+38+20 | | | | 110+12 | | | | 100 | | | | | | | |
| КОН | 2016-06-01 | | | | 2016-06-30 | | | | 2016-06-01 | | | | 2016-09-20 | | | | 2016-12-21 | | | | | | | |
| EBPG | 2016-05-27 | | | | 2016-05-27 | | | | 2016-05-27 | | | | 2016-09-13 | | | | 2016-09-13 | | | | | | | |
| Name | Lil F | | | | Lil H | | | | Dil D | | | | Harald C | | | | Hilde E | | | | | | | |
| Job # | 21063 | | | | 21064 | | | | 21065 | | | | 23060 | | | | 23059 | | | | | | | |

Table 4.2: Science chips. See text.

first two install attempts (Lil F and Lil H) failed. The procedure was modified slightly, reducing the temperature to around 100 $^{\circ}$ C, and the third attempt (Dil D) was installed successfully.

Figures 4.18, 4.19 and 4.20 show reflection or transmission spectra for the buttcoupled waveguides on these chips, and estimates of group index. The first two chips, Lil F and Lil H, are identical patterns, compared above in figs. 4.10 and figs. 4.12. Chips of this era were rectangular scans of width *w* and amplitude *A* with 5 nm step size for each dimension (see fig. 4.11), with an interleaved coupler scan: written coupler width for these prints is 0.123 µm for all devices with odd index, and incremented in 0.004 µm steps for even devices, starting at 0.111 µm for dev. 2 and increasing to 0.139 µm for dev. 16. "Lil" and "Dil" prints are differentiated by the center of their *w*–*A* scan (Dil has 5 nm larger *A* and 13 nm larger *w*) and written gap *g* (270 nm for Lil and 250 nm for Dil). A larger initial *w* parameter produces band structure which is initially closer to the target location, resulting in fewer cycles of ALD—88 cycles for Dil D vs. 97 and 118 for Lil F and Lil H, respectively.

November 2016–January 2017

A second install was performed in the fall and winter. Wary of the problems encountered with the heat curing epoxy, we decided to revert to UV curing glues, which were used in the earliest stages of this project, but were abandoned in 2014 due to concerns about vacuum compatibility. Unfortunately, the install attempted with the UV curing glue in November 2016 (Harald C) ultimately failed—after installing the chip in the chamber, some bonds between chip and fibers broke, resulting in many devices without transmission. We troubleshot the heat curing epoxy procedure (described below), and successfully installed a chip (Hilde E) in January 2017. This install resulted in 8 devices coupled in transmission and reflection, resulting in a record number of fibers (16) fed into the vacuum chamber—prior to this point the maximum number of addressed devices was 4 (8 fibers).

Like substrates considered in the previous install, Harald C (fig. 4.21) and Hilde E (fig. 4.22) are rectangular (w, A) scans, this time with reduced step size: Harald C, dev. 1 has written parameters (w, A) = (0.134, 0.292) µm, with step size (δw , δA) = (0.5, 2) nm, whereas Hilde E is a slightly wider scan, with dev. 1 (w, A) = (0.1355, 0.282) µm and (δw , δA) = (1, 4) nm.



Figure 4.18: Lil F reflection spectra. From top to bottom: device 10, 12, 15 and 16. Left (right) column shows dielectric (air) band edge; red (blue) line indicates Cs D_1 (D_2) transition frequency; × and \circ are estimates of group index n_g using eq. 2.9 with $L = (2N_t + N_n)a$ and $L = N_n a$, respectively.

4.5 Packaging and Installation

After characterization and ALD, chips used in the experiment must be packaged for installation in the vacuum chamber. Fewer than ten chips have been installed in the Lab 2 chamber (see appendix A). This is too few installations to generate meaningful statistics, but I will describe what we've found has worked, what has not, and what we've determined are "best practices" for chip packaging.

Fibers and chip are fixed to our mount assembly using glue. Over the course of this project we have used two different types of glue. Initially we used UV-curable adhesives—Dymax OP-29 to fix the chip to the mounting piece and to strain relieve, and Dymax OP-4-20632 to glue fibers to the chip. Concerns about vacuum compatibility of these adhesives led us to transition to a heat-curing epoxy, EPO-TEK 353ND. Despite complications described above, the thermal cure glue is preferred to the UV-curable adhesives—in our experience the bonds are more robust.



Figure 4.19: Lil H transmission spectra. From top to bottom: device 10, 11, 12 and 16. Left (right) column shows dielectric (air) band edge; red (blue) line indicates Cs D₁ (D₂) transition frequency; × and \circ are estimates of group index n_g using eq. 2.9 with $L = (2N_t + N_n)a$ and $L = N_n a$, respectively.

One drawback of the EPO-TEK epoxy is that curing requires that the temperature be elevated to around 110 °C, resulting in differential thermal expansion between silica fiber, silicon chip, and mounting piece; and thermo-optic effects, which complicate coupling optimization, in addition to the cracking and stiction problems described above. We have found these problems can be mitigated by sonicating relevant parts (mounting piece, mounting block, tweezers, etc.) in progressive solvent baths (acetone, methanol, isopropyl alcohol); by inserting a piece of clean vacuum foil between mounting block and mounting piece to protect the structure from condensates; by limiting the amount of adhesive used to glue the chip to the mounting piece; and by using mounting pieces made of Macor®, as its coefficient of thermal expansion ($9.3 \times 10^{-6}/K$) [95] is closer than aluminum's ($23 \times 10^{-6}/K$) [96] to silicon's ($2.3 \times 10^{-6}/K$) [97].ⁱ Chip and mount are heated using a 35 Ω resistive

ⁱThe investigation leading to these recommendations, particularly regarding the avoidance of condensates, was led by Su-Peng Yu and Szilard Szoke.



Figure 4.20: **Dil D transmission spectra.** From top to bottom: device 5, 9, 11 and 12. Left (right) column shows dielectric (air) band edge; red (blue) line indicates Cs D_1 (D_2) transition frequency; × and \circ are estimates of group index n_g using eq. 2.9 with $L = (2N_t + N_n)a$ and $L = N_n a$, respectively. Installed in lab 2 2016-07-13.

heater integrated into the mounting block. To avoid thermal shock we heat and cool the assembly slowly—a ramp of 2 V/min. is adequately conservative. At 20 V (ca. 11 W) the stage has reached a sufficient temperature to cure the glue. At the end of the session, the integrated heater is ramped down at the same rate. To accelerate the curing schedule, epoxy is mixed and "pre-cured" for a few minutes on a hotplate. We limit the amount of glue used to attach the chip to the mounting piece—uncured glue has very low viscosity and can "flood" the vee grooves. We have found that tacking the chips at three corners is sufficient.

After the chip has been tacked to the mounting piece, fibers are aligned in the vee grooves and glued in place. Proper alignment is confirmed by acquiring transmission or reflection spectra. Each fiber is first glued into the groove by applying glue at a displacement of ca. 500 μ m from the coupler, and then again over the mounting piece, to "strain relieve" the fiber, so that gentle pulling does not stress the fiber at the coupler. A very small amount of glue is used in the vee groove (ca. 5 nL)



Figure 4.21: Harald C transmission spectra. From top to bottom: device 5, 7, 8 and 11. Left (right) column shows dielectric (air) band edge; red (blue) line indicates Cs D₁ (D₂) transition frequency; × and \circ are estimates of group index n_g using eq. 2.9 with $L = (2N_t + N_n)a$ and $L = N_n a$, respectively.

to ensure it does not interfere with the coupler. The best bonds are formed if the uncured glue wicks into the groove underneath the fiber; this will occur reliably around $110 \,^{\circ}$ C and with the correct volume of glue. A device glued to the mounting piece with one device coupled is shown in fig. 4.23.

A more substantial amount of glue is applied to fix the fibers to the mounting piece (strain relief). A few millimeters of acrylate coating has been removed from the fibers, allowing them to be properly aligned in the vee grooves, which are sized to the uncoated fiber cladding diameter. The length of stripped fiber must be at least as long as the vee groove for the fiber to properly align, but it is recommended that the strain relief bond be located where the fiber's acrylate coating is retained to improve durability. A sufficient volume of glue should be used to completely envelop the coated fiber, but not so much glue that the mound interferes with the positioning or strain relief for the next fiber.

Glue is applied using a clean (stripped and wiped with IPA), cleaved optical fiber.



Figure 4.22: Hilde E reflection spectra. From top to bottom: device 5–11 and 14. Left (right) column shows dielectric (air) band edge; red (blue) line indicates Cs D₁ (D₂) transition frequency; × and \circ are estimates of group index n_g using eq. 2.9 with $L = (2N_t + N_n)a$ and $L = N_n a$, respectively. Installed in lab 2 2017-01-23.



Figure 4.23: **Gluing.** (a) Chip is first glued to Macor® mounting piece, fibers are glued into grooves, and strain relief glue is applied at locations indicated by red arrows. Image was taken before foil cover became standard operating procedure. (b) Image of a package chip mounted on the aluminum arm.

Gluing in the vee groove is delicate, and is performed using a computer-controlled actuator we've named "glue bot" (see fig. 4.3). The strain relief is applied with a manual applicator, traditionally a clean, cleaved fiber fixed to an ESD-safe screwdriver with Kapton® tape. After strain relief glue has cured, the bond is tested, using the fiber positioning micrometers to gently pull on the fiber. If bond is deemed adequate, the fiber is cut out of the characterization setup, the magnets holding the fiber in the vee groove are carefully removed, and the fiber spool is set aside. The next fiber is spliced into characterization, and the process repeats until all desired devices are coupled. It is recommended that all fibers on the "front" side of the chip be glued before proceeding to the "back" side, and that devices be coupled in order (from high to low or low to high device number) to minimize interference between strain relief bonds.

After all devices have been coupled, the chip assembly is attached to an aluminum arm mounted on a CF multiplexer using a groove grabber.ⁱ Fibers are then threaded through the multiplexer and fiber feedthroughs. Experimental apparatuses in the Kimble group requiring in vacuo optical fiber have, since 2009 [98, Sec. 2.3.1], used PTFE feedthroughs based on a design described in ref. [99]. The current design uses a CF-Swagelok® adapter with 1/8 in. tube diameter. PTFE feedthroughs have 1–4, 250 μ m-diameter holes drilled in them. If the Swagelok® nut compresses the PTFE feedthrough, fibers are held in place; increased tightness is sufficient to hold vacuum.ⁱⁱ Fibers which will be in the vacuum chamber are "groomed" to be loose

ⁱKimball Physics MCF275-GrvGrb-CB03.

ⁱⁱAnecdotally one full rotation beyond "finger-tight."
enough that they are not under stress but tight enough that they do not compress too much upon insertion into the spherical octagon. Some compression is inevitable, but will not affect coupling if strain relief is adequate.

Gluing is performed in a different room from the vacuum chamber. We have built a vessel to transfer the chip assembly from room to room, consisting of a hinged acrylic box with a CF flange to which the assembly can be bolted, and a gas fitting to purge the system with nitrogen. After installation, and a few days of pumping, the system is again ready to deliver atoms.

4.6 Alkali Corrosion

Unfortunately, exposing our waveguides to cesium results in significant degradation. Figure 4.24 shows SEM images of two waveguides removed from the vacuum chamber in early 2016. On the left is a waveguide which was not used in the experiment—it looks roughly as it did before the sample was installed. In the center is a device that was used heavily in the experiment, which consequently looks qualitatively different.

This degradation is associated with observable changes to the reflection and transmission spectra, measured in situ. Our first observations witnessed a rapid shift around 1 THz over the course of two weeks, shifting towards lower frequency (redshift). In both Labs 2 and 11 significant spectral shifts were observed only after switching to dipole trapping schemes—densities in the mini-MOT around the waveguide are insufficient to produce a measurable shift. As our experiment is very sensitive to the relative position of the Cesium D line transitions and the band gap, the wandering band structure is problematic. The useful lifetime of any APCW is shortened to the duration of time the D₁ transition remains outside the band gap.

We did not initially keep records of how heavily the experiment was used, making estimations of cesium load difficult, but soon after observing significant shifts we incorporated a system to count the number of experimental cycles in which atoms are delivered to the waveguide. While this accounting is flawed—it does not account for misalignment of the optical lattice and the device, nor does it incorporate any of the many parameters set in the experiment (GM powers and detunings, lattice speed, etc.)— it does give a rough sense of how quickly features tune.

Plots of the positions of band edge resonances and cumulative number of experimental cycles N_p for a single device (job # 21065 "Dil D," dev. 9) are shown in fig. 4.25.



Figure 4.24: Signs of cesium exposure. Scale bar represents 1 μ m. (a) DCOR28D, device 2 (in vacuum chamber but unused in experiment). (b) DCOR28D, device 3. waveguides are separated by 320 μ m. (c) Reflection spectra before (top), after (bottom) cesium exposure. This chip was tuned using a CF₄ etch (see sec. 4.2), and was not treated with ALD.

At times when the experimental is not run, e.g., between July 30 and August 7,ⁱ spectral features do not shift appreciably. Similarly, periods with large spectral shift are associated with more rapid accumulation of experimental cycles. This suggests shifts are not associated with cesium background present in the chamber, but rather directly correlated with operating the experiment. The black horizontal line indicates the frequency of the cesium D₁ transition, revealing that after around 5×10^4 experimental cycles the cesium transition has been shifted past the last visible resonance. A plot of the resonance shift, relative to their original ($N_p = 0$) frequency, reveals a linear trend. A linear fit yields the mean shift per experimental cycle, (-26 ± 1) MHz/cycle.

Transmission spectra at different N_p are also presented. These spectra share many features with models of long-range distortion and loss presented in sec. 2.4. The first resonance disappears at low cycle count ($N_p \approx 3 \times 10^5$) even as transmission far from the band edge remains high, consistent with a spatially-dependent loss. A second noteworthy feature is increased transmission at the anti-resonance between second and third band edge resonances, suggesting an "improvement" in impedance matching at that frequency, resembling the behavior of long length-scale band bending.

Naturally, it would be helpful to understand the physical process that causes the observed spectral shift. One reasonable hypothesis is that the nanobeams are being coated in cesium. The optical properties of metallic cesium can be obtained from

ⁱAt this time researchers were at the 2016 GRC on Quantum Science.



Figure 4.25: In situ cesium shift. *Left, top*: Resonance position and cumulative number of probe cycles, N_p , vs. time for job #21065 "Dil D," dev. 9. The frequency of the first dielectric band resonance is indicated by \circ , the second by \times , and the third by \Diamond . Solid black horizontal line indicates the D₁ transition frequency. *Left, bottom*: Frequency shift vs. N_p . Linear fit gives shift rate of (-26 ± 1) MHz/cycle. *Right*: Transmission spectra at various N_p . This chip was treated with 88 cyc. of ALD (see tab. 4.2).

the Drude model for metals [55, 100], which gives a frequency-dependent complex permittivity

$$\hat{\epsilon}(\omega) \equiv \hat{n}(\omega)^2 = 1 - \frac{\omega_{\rm p}^2}{\omega(\omega - i\tau^{-1})},\tag{4.1}$$

where ω_p is the plasma frequency, τ is the relaxation time and $\hat{n} = n + ik$ is the complex refractive index. In the model, below ω_p there is a frequency window for which n < 1; the width of this window depends on τ . For cesium $\omega_p = 2\pi \times 856$ THz and $\tau = 2.1 \times 10^{-14}$ s, giving $n \approx 0.3$ near the D line transitions.

Figure 4.26a–b show the effect of a 10 nm adlayer with variable *n* on APCW band structure. For any material with n > 1 (i.e., normal dielectrics) a redshift of f_0



Figure 4.26: **Simulating adlayers.** (**a**–**b**) Simulated change of APCW band gap position $(f_0, \Delta f)$ for a 10 nm adlayer with variable refractive index n. (a) shows shift of f_0 and (b) Δf . Dielectric adlayers have n > 1, and result in redshift of f_0 and shrinking Δf . In a frequency window below the plasma frequency, metals exhibit 0 < n < 1. (**c**–**d**) Simulated change of APCW band gap position for a cesium adlayer with variable thickness, confirming expected blueshift for metals: (c) shows shift of f_0 and (d) Δf . Data in (a–b) are obtained using MPB, while (c–d) were simulated in COMSOL by A. Asenjo-Garcia and S.-P. Yu.

is expected, while materials with n < 1 should exhibit blueshift. Figure 4.26 shows simulated band movement for a metallic cesium layer of variable thickness, confirming that the expected shift direction for metallic cesium is blue. As the experimentally-observed shift is consistently red, we do not believe the shift is due to a metallic cesium coating.

Protecting against Corrosion

There is some indication in the literature that ALD-alumina may be able to protect our waveguides from cesium's detrimental effects. A study of micro-fabricated cesium vapor cells found that a thin (ca. 6 nm) alumina-coating increases the vapor density lifetime by two orders of magnitude relative to uncoated cells, presumably by reducing the chemical interaction of cesium with the vapor cell walls [88]. Recent papers from the Pfau group describe a still more relevant apparatus in which they perform experiments with silicon nitride photonics (waveguides and ring resonators) and thermal rubidium vapor [101, 102]. They have pursued ALDalumina passivation with thickness similar to the vapor cell study (9 nm), reporting that this treatment allows resonances to "remain visible, although their linewidth is still increased after rubidium exposure" [102]. We are currently pursuing an ALDalumina passivation strategy to protect our structures, with coating thicknesses of 6–10 nm. Unfortunately, spectral shifts persist in our experiment.

Beyond surface passivation, we have also explored methods to actively reverse cesium corrosion. The Pfau group has observed that degradation can be partially reversed by elevating the temperature of the chip to 200 °C [101]. We have tried to elevate the temperature of the APCW using GM heating; the maximum achievable temperature is limited by the damage threshold for the devices, which is typically $\approx 300 \ \mu$ W. Temperature can also be increased by heating the silicon substrate with a free-space 940 nm laser, but not sufficiently to induce desorption.

We have also tried illuminating our structures with UV light in the hopes of removing cesium adsorbates. Light-induced atomic desorption (LIAD) is a technique used in some cold atom experiments to temporarily increase the partial pressure of an alkali to load a MOT. The increase in partial pressure is attributed to alkali adsorbates being driven off the walls of the vacuum chamber by UV light. The LIAD effect has been found to be surface- and frequency-dependent [103]. In our experiment we have seen partial spectral recovery, a ca. 1 THz blueshift, after imaging a 404 nm laser diode on an APCW overnight. Unfortunately, this seems to have reduced the power damage threshold for the waveguide, which failed soon afterward. Consequently, this technique is not currently recommended.

Finally, there is some evidence that surface coatings terminated in methyl groups can be used as chemical passivation against cesium [104, 105]. Similar coatings have been shown to reduce spin relaxation in in alkali vapor cells [106]; improved passivation and reduced spin relaxation are both attributed to the low polarizability of the coating. Certain self-assembled monolayer coatings with appropriate termination have successfully been integrated into fabrication procedures similar to ours [107].

Chapter 5

ATOMIC PHYSICS EXPERIMENT

In this chapter I briefly describe the Lab 2 atomic physics apparatus used to cool, transport and interrogate atoms around the APCW. This system shares many features with the Lab 11 setup—the principal difference is the atom delivery method. Many more details about this system may be found in ref. [93].

5.1 Vacuum System and Magneto-Optical Traps

Our vacuum system consists of two stainless steel Kimball Physics multi-CF[™] chambers separated by 70 cm. The first chamber is a 2.75 in. spherical hexagonⁱ attached through a valve to a cesium reservoir, referred to as the "source" chamber. The second chamber is a 6 in. spherical octagon,ⁱⁱ referred to as the "science" chamber. These two chambers are connected by a tube consisting of many components, including an all-metal inline valve,ⁱⁱⁱ which allows one side of the system to be brought to atmosphere while maintaining good vacuum on the other side. This valve is closed to isolate the source chamber when replacing a chip in the science chamber. On each side of the valve there is an assembly with an angle valve and Kwik-Flange[™] adapter to which a turbo pump station can be attached, as well as an ion getter pump, necessary to obtain sufficiently good vacuum ($\lesssim 10^{-8}$ torr) to load a magneto-optical trap (MOT). In the tube connecting the chambers there is also a thin differential pumping tube (8 in. long, 0.16 in. diameter), installed to reduce the partial pressure of room-temperature cesium in the science chamber. The isolation of the cesium source is intended to protect sensitive photonic devices from alkali corrosion (discussed in sec. 4.6). The axis connecting these two chambers runs east-west, corresponding to the y direction in APCW coordinates. A diagram of the vacuum system is shown in fig. 5.1.

A MOT is loaded in the source chamber from the room-temperature cesium background, using the usual six-beam configuration [108] and anti-Helmholtz coils with a gradient of ca. 10 G/cm and coil axis normal to the optical table. Atoms are cooled using light ca. 10 MHz red-detuned from the D₂ $F = 4 \rightarrow F' = 5$ transition, and

ⁱMCF275-SphHex-C2A6 ⁱⁱMCF600-SphOct-F2C8 ⁱⁱⁱMDC part #316001



Figure 5.1: **Diagram of vacuum system.** Red rectangles represent horizontal MOT beam paths and green circles locations of source and science MOTs. Source chamber (**a**) is connected to science chamber (**d**) by a channel in which a differential pumping tube (**b**) and inline valve (**c**) are present. A MOT is formed from cesium background vapor in (a) and pushed ballistically west to a secondary MOT in (d) using a near-resonant pushing beam. The separation of the chambers is L = 70 cm.

re-pumped using light resonant with the $D_2 F = 3 \rightarrow F' = 4$ transition. Atoms are transferred from the source to science chamber using a pushing beam, similar to the technique described in refs. [109, 110], and re-captured in a secondary "science" MOT. In-transit atoms are in free fall, and so must have sufficient velocity to cover the 70 cm distance without falling so far that they miss the differential pumping aperture, yet move sufficiently slowly to be recaptured by the science MOT. An in-depth study to optimize the pushing beam parameters was performed, and is discussed in ref. [93].

5.2 Transporting Atoms from Science MOT to APCW

In Lab 2 two different techniques have been used to transfer atoms from the science MOT to the chip. Here I briefly summarize each technique, and motivate our current, optical-lattice based transport scheme.

Mini-MOT

The first technique to localize atoms near the APCWs relied on forming a MOT in the 2×6 mm window in the chip. Atoms were transferred from science MOT to this "mini-MOT" by rapidly changing the position of the magnetic field zero, causing atoms to be "kicked" from the science MOT to the chip, and subsequently recaptured in the mini-MOT. The kick has peculiar behavior—atoms can only be pushed along specific directions due to velocity selective resonance [111], setting constraints on the relative position of science MOT and chip. In order to form a MOT in the



Figure 5.2: **Mini-MOT kick.** Images (a)–(f) show absorption images for atoms in transit, each separated by 1 ms. Position of window is indicated by black rectangle for clarity. By frame (f), the mini-MOT beams have been turned on and the atoms are decelerated. Cloud center-of-mass velocity is ca. 1 m/s.

window of the chip, six beams must be passed through the window. To achieve this geometry, the surface normal of the chip was oriented at an angle of 45° with respect to the optical table surface normal. Figure 5.2 shows sequential absorption images of a cloud of atoms transiting between science MOT and mini-MOT.

Operating a MOT with millimeter-scale beams leads to some complications. The capture velocity for atoms in a MOT is proportional to the square root of the beam diameter [112], meaning atoms with sufficiently high velocity or temperature will be lost. We have also found empirically that the MOT lifetime varies roughly linearly with the beam diameter in this regime. Threading beams through a millimeter-scale window also poses problems—scattering on the chip makes the mini-MOT very sensitive to beam pointing.

Still, the mini-MOT technique was sufficient for initial experiments. In ref. [31] we inferred a Purcell factor enhancement $P_{wg} \equiv \Gamma_{1D}/\Gamma' = (0.32 \pm 0.08)$ for atoms localized around an APCW, with Γ_{1D} representing the decay rate of atoms into the TE GM and Γ' the decay rate into all other channels. This inference was made using measured reflection spectra and transfer matrix model (TMM) incorporating the effects of atoms coupled to the photonic crystal. In this experiment a GM blue-detuned 10 GHz from the D₁ $F = 4' \rightarrow F = 4$ transition was used to modify the spatial density profile, guiding atoms from the surrounding mini-MOT into the



105

Figure 5.3: Enhanced decay rate in mini-MOT. (a) Reflection data (black) and TMM fit (red) for atoms interacting with the probe mode of an APCW. Reflection data are normalized to the nominal reflection of the structure absent atoms. (b) Trajectory simulation for atoms near an APCW with blue-detuned GM optical potential. Color indicates density relative to surrounding mini-MOT. Blue GM reduces density on device exterior, but leaves a channel in between the nanobeams through which atoms can transit. Figures adapted from ref. [31].

region of high probe-mode intensity while repelling atoms from unwanted regions, e.g., the single nanobeam section of the waveguide and the corrugated exterior. The density of atoms around the APCW relative to the surrounding mini-MOT number density (ca. 2×10^{10} cm⁻³) can be simulated using a particle trajectory simulation. A reflection spectrum with model fit and relative density simulation results are shown in fig. 5.3.

Beyond the Mini-MOT

While initial experiments with the mini-MOT were encouraging, the kinds of experiments that can be performed with transient, dilute atoms are limited. At this juncture labs 2 and 11 began to pursue different techniques to trap atoms at higher densities near the device. In Lab 11 dipole traps relying on a single incident external illumination beam and its reflection off the APCW were pursued. The polarization of the incident beam is aligned parallel to the nanobeams, maximizing reflection. This scheme creates a cigar-shaped optical potential above the waveguide, ca. 220 nm above the center of the structure, into which atoms from the surrounding mini-MOT can be cooled. This trap increases the number density of atoms by an order of magnitude over the surrounding mini-MOT [32], allowing several (ca. 3) atoms to interact simultaneously with the probe mode. The regime of several atoms is very interesting, as evidenced by the observation of superradiance [32] and the study of waveguide-mediated atom-atom interactions on both sides of the band edge [33].



Figure 5.4: **Optical lattice simulation** (courtesy L. Peng). Frames (a)–(d) show atoms (black, red) conveyed in a travelling-wave potential (green, white) towards a double nanobeam structure, depicted with coordinate system in (e). Each snapshot is separated by 0.3 μ s; atoms travel left to right ($-\hat{z}$ direction), arriving periodically with 1.25 μ s period (800 kHz lattice frequency). This technique guides an appreciable number of atoms into the region between the nanobeams, producing periodically-modulated transmission/reflection.

In Lab 2 we elected to pursue a different delivery strategy, avoiding the use of a mini-MOT entirely. At the time of writing our transport scheme consists of two counter-propagating beams which form a one-dimensional standing wave trap. Atoms are cooled from the science MOT into the optical lattice. By controllably detuning the frequency of one beam, we are able to create a slowly-moving optical potential that conveys atoms in the \hat{z} direction towards the APCW. The polarization of these beams is oriented nearly perpendicular to the nanobeams (along \hat{y}) in order to minimize reflection. The combination of this moving lattice and GM potentials can be used to guide a significant fraction of atoms into the target region between the nanobeams. The lattice cooling and transport techniques enable the delivery of millions of atoms at high number density, ca. 10^{11} cm⁻³. Because the parameter space for this family of traps is very large, our search is aided by numerical trajectory simulations performed by Lucas Peng; an example is shown in fig. 5.4.

Lattice delivery requires modifying the way the chip is mounted in the vacuum chamber. The current mounting configuration in Lab 2 is depicted in fig. 5.5: the chip is mounted on an arm which positions it at a small offset from the center of a spherical octagon, allowing the science MOT to be formed ca. 25 mm from the device plane. Before lattice delivery the chip was mounted on a rotary-linear feedthrough, similar to the system in Lab 11 [113], allowing the chip to be retracted. Sacrificing that degree of freedom allows the science MOT to be formed in the



Figure 5.5: Illustration of Lab 2 vacuum chamber. (a) Side view: Chip is positioned so that window is centered vertically and horizontally in viewports along the \hat{z} direction. Red arrows indicate paths for MOT beams, green arrows atoms arriving from source MOT, and orange rectangles location of quadrupole coils. (b) Top view: Chip is positioned at a displacement of 25 mm from center of chamber, allowing the science MOT to be formed. (c) Cutaway: Arm which holds chip is mounted to a groove grabber in a multiplexer, which has flanges for fiber feedthroughs while maintaining optical access for MOT beams.

center of the chamber, provides more optical access, and simplifies the installation process—fiber feedthroughs can be threaded in the characterization lab before breaking vacuum. CAD drawings of mounting piece and arm are included as an attachment at CaltechTHESIS.

5.3 Computer Control and Data Acquisition

Sequence Programming

An experimental cycle consists of many steps: loading the science MOT from the source, cooling into the optical lattice, additional lattice cooling, conveyance and interrogation of the reflection and transmission at the APCW. These procedures require that laser intensities and detunings, microwaves, and magnetic fields vary in time. Some instruments are controlled using serial communication (USB/GPIB); the majority of the experiment is controlled by National Instruments input/output modules with programmable digital and analog outputs.

Due to the complexity of experimental sequences, which require sub-millisecond timing resolution on ca. 50 different channels acting in concert, we have devised a system to disentangle sequence programming from execution, comprising a MAT-



Figure 5.6: Sequence composer GUI. See text.

LAB sequence *composer* and LabVIEW sequence *player*. This conceptual design is to my knowledge new for the Kimble group—when I joined, sequence timing was hard-coded into the LabVIEW circuit and programmed directly on the LabVIEW front panel, making record-keeping and procedural changes cumbersome. This sort of sequence programming is a general feature of AMO experiments, and our solution could be easily applied in other settings.

The sequence composer is shown in fig. 5.6. Panel (a) shows the voltage waveform for an analog channel over the course of the 550 ms experimental sequence. Panel (b) shows conceptual subroutines (MOT and polarization gradient cooling stages, e.g.), which, when selected display event timing in panel (c). Panel (d) shows a list of variables. When a sequence is created, variables in panels (b) and (c) are replaced with values before evaluation, simplifying parameter sweeps and events with relative timing.ⁱ The output of this program is a two-dimensional array of numbers, for which each column represents a different channel, and each row a 10 μ s time step. These arrays are saved as a text file, and read by a LabVIEW program. Variables can be swept linearly, resulting in an output file for each unique combination of variable values, creating meshes of arbitrary dimensionality. These sorts of *n*-dimensional sweeps are useful for optimization, e.g., magnetic field zeroing.

ⁱI was quite pleased by the flexibility of this interpreted variable design; its conceptual resemblance to the CAD drawing system discussed in sec. 3.3 is not accidental.

Optical Lattice Chirp Generation

The frequencies for the optical lattice beams are generated by direct digital synthesis (DDS), which is used to drive acousto-optic modulators (AOMs) to shift the frequency of light generated by a Ti:sapphire laser. Light sent to the AOMs is derived from the same source, and the AOMs are driven at around 80 MHz, which, in double-pass configuration results in a frequency shift of around 160 MHz. The first AOM is driven with a constant tone derived from a function generator.ⁱ To convey atoms, the second AOM is driven by a time-varying frequency, generated by mixing a 70 MHz tone with a time-varying tone around 10 MHz that is generated by an Analog Devices AD9854 DDS clocked at 200 MHz. All synthesizers are phase stabilized by a common 10 MHz clock. Atoms in an optical lattice can be transported by changing the frequency produced by the AD9854.

Arbitrary frequency ramps are achieved by operating AD9854 in "chirp" mode, which ramps the DDS-generated frequency by a specified frequency slew rate. The frequency ramps are programmed one at a time into on-board memory on the AD9854 and then updated, allowing any piecewise-linear frequency waveform to be created. Our AD9854 is on an evaluation board, and communication is achieved via Serial Peripheral Interface Bus (SPI) from an Arduino Due microcontroller. Frequency ramp commands are calculated and stored in the Arduino on-board memory, transferred to the AD9854 one at a time and updated using TTL pulses generated by the Arduino's on-board clock. The Arduino program is loaded before sequence execution, and triggered by a TTL from the sequence computer. The chirp is programmed in MATLAB, with a script to translate the arbitrary chirp into Arduino C. While the software allows arbitrary chirp waveforms, in practice piecewise linear chirps are used.

Image Acquisition and Photon Counting

The most common data produced by the experiment are absorption images and photon-counting data. Absorption images, like those shown in fig. 5.2, are created by illuminating with a weak, resonant probe beam—atoms absorb and rescatter light, casting "shadows" on a CCD. Image data are used to determine optical depth, allowing us to extract information about densities and number, atom temperature via time-of-flight techniques, and magnetic fields via microwave spectroscopy. Images are acquired using NIR-sensitive Sony XC-EI30 CCDs, and read-out by a dedicated imaging computer. Absorption imaging is ubiquitous in experiments with cold

ⁱTektronix AFG3251



Figure 5.7: **GM absorption spectrum.** (a) Depiction of raw data. An incident probe at fixed detuning transmits through the device. Photons are detected on an SPCM and recorded on a digitizer. Binary data contains time stamps of "clicks." (b) Transmission spectrum, created by binning data depicted in (a) into time and frequency bins. Spectrum shows absorption of atoms around APCW around 20 ms and 0 MHz probe detuning. For this dataset (2017-04-05 111600), the probe is TM polarized, lattice frequency is 800 kHz, and there is a blue-detuned guided mode potential. Data are taken on Hilde E, dev. 7 (see fig. 4.22).

neutral atoms; some details are discussed in ref. [114].

Spectra are acquired by monitoring the reflection and transmission of a GM probe on NIR-sensitive single-photon counting modules (SPCM).ⁱ The measured quantum efficiency of these modules is typically ca. 40% at 850 nm. A detection generates a TTL pulse with 50 ns width, which is sent to a five-channel digitizer with 200 ps resolution.ⁱⁱ Histograms of these data are used to produce spectra with time and probe-detuning information, like the example shown in fig. 5.7.

5.4 Search for Phase Sensitivity

Simulations like the one shown in fig. 5.4 can also be used to produce transmission spectra. These simulated spectra indicate that, under the correct conditions, transmission and reflection signals in the presence of localized atoms transiting between the nanobeams will be modulated at the frequency at which the lattice transits through the device. As recently as April 2017 we have begun to see signatures suggestive of this behavior. There are many techniques for extracting a periodic signal in the presence of noise; in our experiment the most successful techniques

ⁱPerkin-Elmer SPCM-AQR-14

ⁱⁱFAST ComTec MCS6A

rely on the combination of information in two channels: a "fringe" channel, which produces a click each time an intensity maximum passes through the device, and a "signal" channel, which monitors transmission of light through the device.

One possibility is to consider the cross correlation *X* between fringe a(t) and signal b(t):

$$X(\Delta t, t_0, t_1; \tau) = \int_{\tau - \Delta t/2}^{\tau + \Delta t/2} \mathrm{d}\tau' \, \int_{t_0}^{t_1} \mathrm{d}t \, a(t) b(t + \tau'), \tag{5.1}$$

where $[t_0, t_1]$ is the time interval under consideration, Δt is the coincidence window, and τ is the lag. Plots of X and its amplitude spectrum

$$Y(\Delta t, t_0, t_1, \Delta \tau; \nu) = \left| \frac{1}{\Delta \tau} \int_{-\Delta \tau}^{\Delta \tau} \mathrm{d}\tau' \left\langle X(\Delta t, t_0, t_1; \tau') \right\rangle e^{-2\pi \mathrm{i}\tau' \nu} \right|$$
(5.2)

over two different time intervals are shown in fig. 5.8. The maximum lag considered is $\Delta \tau = 200 \ \mu$ s, and the ensemble average (indicated by $\langle \rangle$) is taken over five repetitions with identical experimental conditions. In a period during which atoms are present at the APCW ([20,22] ms), a strong frequency-dependent signal is apparent, while at a later time ([38,40] ms) the contrast disappears. The amplitude spectrum for both intervals shows a peak at the lattice frequency (800 kHz). The peak at the later time interval is not the result of residual atoms—the cross-correlation of the fringe signal with a simulated Poisson source with identical count rate (black dashed line) shows a peak with similar magnitude. Still, the amplitude spectrum when atoms are present shows a peak with height exceeding what is expected for an unmodulated Poisson source, indicating that the presence of atoms does contribute to this signal.

The cross correlation and similar analyses show signatures of a periodic, atomdependent signal, but have a complicated interpretation. A more sophisticated effort to relate simulated and measured fringe-sensitive spectra, incorporating many effects (e.g., AC Stark shifts, number and spatial distribution of atoms, spatial mode profiles) has begun in Lab 2, and will be discussed in an upcoming manuscript.



Figure 5.8: **Cross-correlation signal.** *Top*: Red (blue) curve shows cross correlation (eq. 5.1) over the interval $[t_0, t_1] = [20, 22]$ ms ([38, 40] ms) with $\Delta t = 100$ ns. *Bottom*: Amplitude spectrum *Y* (eq. 5.2) of data shown in *top*, over the lag interval $\Delta \tau = 200 \ \mu$ s. The same dataset as in fig. 5.7 is used, with 0 MHz probe detuning. Red and blue curves correspond to data of the same color shown in *top*, while dashed black line shows averaged amplitude spectrum for cross correlation of fringe and simulated Poisson source with identical count rate. Spectrum is averaged over 1000 randomized intervals.

OUTLOOK

Integrating ultracold atoms with nanophotonics is a challenging endeavor. In our group we have designed photonic crystal waveguides with spatial modes and band structure that enhances atom–light interaction near the cesium D line transitions. We have performed process characterization and engineering to produce low loss structures with accurate band placement. We have also developed sophisticated atom delivery methods, recently observing phase-dependent signals consistent with the transport of atoms through the region between the nanobeams. With luck, we will soon have a dense array of atoms, trapped by and interacting through the guided modes of a photonic crystal waveguide.

The degradation and spectral shift of our APCWs resulting from delivering cesium is a serious impediment to our work: the finite usable lifetime and varying spectral properties makes a systematic search of delivery and trapping parameter space difficult. Finding a solution to this problem, either through surface passivation or active desorption, is imperative. Increasing the number of addressable devices on a chip may also alleviate some pressure, decreasing the frequency of sample exchange. The fiber butt-coupling technique employed in this work limits the density of devices on a single chip. Switching to a free-space coupling scheme [48] could increase waveguide density, and has other advantages as well, including the simplification of the installation process and reduction of unintended reflections in the waveguide.

The fabrication process described in this work is quite flexible, and could easily be extended to other waveguide geometries. In our group a great amount of thought has been devoted to planning future photonic crystal geometries, including waveguide defects in 2D photonic crystals [48] and true 2D photonic crystals [48, 93], which would potentially allow the exploration of increasingly exotic atom–photon interactions.

BIBLIOGRAPHY

- 1. Cronin, A. D., Schmiedmayer, J. & Pritchard, D. E. Optics and interferometry with atoms and molecules. *Rev. Mod. Phys.* **81**, 1051–1129 (2009).
- Ludlow, A. D., Boyd, M. M., Ye, J., Peik, E. & Schmidt, P. O. Optical Atomic Clocks. *Rev. Mod. Phys.* 87, 637 (2015).
- 3. Bloch, I., Dalibard, J. & Nascimbène, S. Quantum simulations with ultracold quantum gases. *Nat. Phys.* **8**, 267–276 (2012).
- 4. Kimble, H. J. The quantum internet. *Nature* **453**, 1023–1030 (2008).
- 5. Wallquist, M., Hammerer, K., Rabl, P., Lukin, M. & Zoller, P. Hybrid quantum devices and quantum engineering. *Phys. Scripta* **T137**, 014001 (Dec. 2009).
- 6. Chang, D., Vuletic, V. & Lukin, M. Quantum nonlinear optics—photon by photon. *Nat. Photon.* **8**, 685–694 (2014).
- Lodahl, P., Mahmoodian, S. & Stobbe, S. Interfacing single photons and single quantum dots with photonic nanostructures. *Rev.Mod. Phys.* 87, 347– 400 (2015).
- 8. Vahala, K. J. Optical microcavities. *Nature* **424**, 839–846 (Aug. 2003).
- 9. Purcell, E. M. Spontaneous Emission Probabilities at Radio Frequencies. *Physical Review* **69**, 681 (1946).
- 10. Yablonovitch, E. Inhibited Spontaneous Emission in Solid-State Physics and Electronics. *Phys. Rev. Lett.* **58**, 2059 (1987).
- 11. John, S. Strong localization of photons in certain disordered dielectric superlattices. *Phys. Rev. Lett.* **58**, 2486–2489 (1987).
- 12. Skorobogatiy, M. & Yang, J. Fundamentals of Photonic Crystal Guiding. (Cambridge University Press, 2008).
- 13. Manga Rao, V. S. C. & Hughes, S. Single quantum-dot Purcell factor and β factor in a photonic crystal waveguide. *Phys. Rev. B* **75**, 1–7 (2007).
- 14. Noda, S., Fujita, M. & Asano, T. Spontaneous-emission control by photonic crystals and nanocavities. *Nat. Photon.* **1**, 449–458 (2007).
- 15. Birnbaum, K. M., Boca, A., Miller, R., Boozer, A. D., Northup, T. E. & Kimble, H. J. Photon blockade in an optical cavity with one trapped atom. *Nature* **436**, 87–90 (July 2005).
- Casabone, B., Stute, A., Friebe, K., Brandstätter, B., Schüppert, K., Blatt, R. & Northup, T. E. Heralded Entanglement of Two Ions in an Optical Cavity. *Phys. Rev. Lett.* **111**, 100505 (Sept. 2013).

- 17. Reiserer, A., Kalb, N., Rempe, G. & Ritter, S. A quantum gate between a flying optical photon and a single trapped atom. *Nature* **508**, 237–240 (Apr. 2014).
- Aoki, T., Dayan, B., Wilcut, E., Bowen, W. P., Parkins, A. S., Kippenberg, T. J., Vahala, K. J. & Kimble, H. J. Observation of strong coupling between one atom and a monolithic microresonator. *Nature* 443, 671–674 (Oct. 2006).
- 19. Shomroni, I., Rosenblum, S., Lovsky, Y., Bechler, O., Guendelman, G. & Dayan, B. All-optical routing of single photons by a one-atom switch controlled by a single photon. *Science* **345**, 903.
- 20. Volz, J., Scheucher, M., Junge, C. & Rauschenbeutel, A. Nonlinear π phase shift for single fibre-guided photons interacting with a single resonator-enhanced atom. *Nat. Photon.* **8** (Nov. 2014).
- Vetsch, E., Reitz, D., Sagué, G., Schmidt, R., Dawkins, S. T. & Rauschenbeutel, A. Optical Interface Created by Laser-Cooled Atoms Trapped in the Evanescent Field Surrounding an Optical Nanofiber. *Phys. Rev. Lett.* 104, 203603 (May 2010).
- Goban, A., Choi, K. S., Alton, D. J., Ding, D., Lacroûte, C., Pototschnig, M., Thiele, T., Stern, N. P. & Kimble, H. J. Demonstration of a State-Insensitive, Compensated Nanofiber Trap. *Phys. Rev. Lett.* **109**, 033603 (July 2012).
- Béguin, J.-B., Bookjans, E. M., Christensen, S. L., Sørensen, H. L., Müller, J. H., Polzik, E. S. & Appel, J. Generation and Detection of a Sub-Poissonian Atom Number Distribution in a One-Dimensional Optical Lattice. *Phys. Rev. Lett.* 113, 263603 (Dec. 2014).
- 24. Jalali, B. & Fathpour, S. Silicon photonics. J. Lightwave Technol. 24, 4600–4615 (2006).
- Yoshie, T., Scherer, A., Hendrickson, J., Khitrova, G., Gibbs, H. M., Rupper, G., Ell, C., Shchekin, O. B. & Deppe, D. G. Vacuum Rabi splitting with a single quantum dot in a photonic crystal nanocavity. *Nature* 432, 9–12 (2004).
- Badolato, A., Hennessy, K., Atatüre, M., Dreiser, J., Hu, E., Petroff, P. M. & Imamoglu, A. Deterministic coupling of single quantum dots to single nanocavity modes. *Science* 308, 1158–61 (2005).
- Arcari, M., Söllner, I., Javadi, A., Hansen, S. L., Mahmoodian, S., Liu, J., Thyrrestrup, H., Lee, E. H., Song, J. D., Stobbe, S. & Lodahl, P. Near-unity coupling efficiency of a quantum emitter to a photonic-crystal waveguide. *Phys. Rev. Lett.* 093603, 1–13 (2014).
- Faraon, A., Santori, C., Huang, Z., Acosta, V. M. & Beausoleil, R. G. Coupling of Nitrogen-Vacancy Centers to Photonic Crystal Cavities in Monocrystalline Diamond. *Phys. Rev. Lett.* **109**, 033604 (July 2012).

- Sipahigil, A., Evans, R. E., Sukachev, D. D., Burek, M. J., Borregaard, J., Bhaskar, M. K., Nguyen, C. T., Pacheco, J. L., Atikian, H. A., Meuwly, C., Camacho, R. M., Jelezko, F., Bielejec, E., Park, H. & Lon, M. An integrated diamond nanophotonics platform for quantum optical networks. *Science* 6875 (2016).
- Thompson, J. D., Tiecke, T. G., de Leon, N. P., Feist, J., Akimov, A. V., Gullans, M., Zibrov, a. S., Vuletić, V. & Lukin, M. D. Coupling a single trapped atom to a nanoscale optical cavity. *Science* 340, 1202–5 (June 2013).
- Goban, A., Hung, C.-L., Yu, S.-P., Hood, J., Muniz, J., Lee, J., Martin, M., McClung, A., Choi, K., Chang, D. E., Painter, O. & Kimble, H. J. Atom–light interactions in photonic crystals. *Nat. Comm.* 5, 3808 (May 2014).
- Goban, A., Hung, C.-L., Hood, J. D., Yu, S. P., Muniz, J. A., Painter, O. & Kimble, H. J. Superradiance for Atoms Trapped along a Photonic Crystal Waveguide. *Phys. Rev. Lett.* **115**, 063601 (2015).
- Hood, J. D., Goban, A., Asenjo-Garcia, A., Lu, M., Yu, S.-P., Chang, D. E. & Kimble, H. J. Atom-atom interactions around the band edge of a photonic crystal waveguide. *P. Natl. Acad. Sci. USA* **113**, 10507–12 (2016).
- Hung, C.-L., Meenehan, S. M., Chang, D. E., Painter, O. & Kimble, H. J. Trapped atoms in one-dimensional photonic crystals. *New J. Phys.* 15, 083026 (Aug. 2013).
- Lalumière, K., Sanders, B. C., van Loo, A. F., Fedorov, A., Wallraff, A. & Blais, A. Input-output theory for waveguide QED with an ensemble of inhomogeneous atoms. *Phys. Rev. A* 88, 043806 (Oct. 2013).
- 36. John, S. & Wang, J. Quantum Electrodynamics near a Photonic Band Gap: Photon Bound States and Dressed Atoms. *Phys. Rev. Lett.* **64**, 2418–2421 (1990).
- Douglas, J. S., Habibian, H., Hung, C.-L., Gorshkov, A. V., Kimble, H. J. & Chang, D. E. Quantum many-body models with cold atoms coupled to photonic crystals. *Nat. Photon.* 9, 326–331 (2015).
- González-Tudela, A., Hung, C.-L., Chang, D. E., Cirac, J. I. & Kimble, H. J. Subwavelength vacuum lattices and atom-atom interactions in photonic crystals. *Nat. Photon.* 9, 320 (2015).
- Dayan, B., Parkins, A. S., Aoki, T., Ostby, E. P., Vahala, K. J. & Kimble, H. J. A Photon Turnstile Dynamically Regulated by One Atom. *Science* 319, 1062–1065 (2008).
- Aoki, T., Parkins, A., Alton, D. J., Regal, C., Dayan, B., Ostby, E., Vahala, K. & Kimble, H. J. Efficient Routing of Single Photons by One Atom and a Microtoroidal Cavity. *Phys. Rev. Lett.* **102**, 83601 (Feb. 2009).

- Alton, D. J., Stern, N. P., Aoki, T., Lee, H., Ostby, E., Vahala, K. J. & Kimble, H. J. Strong interactions of single atoms and photons near a dielectric boundary. *Nat. Phys.* 7, 159–165 (Nov. 2010).
- 42. Alton, D. J. Interacting single atoms with nanophotonics for chip-integrated quantum networks. PhD thesis (California Institute of Technology, 2013).
- 43. Cohen, J. D., Meenehan, S. M. & Painter, O. Optical coupling to nanoscale optomechanical cavities for near quantum-limited motion transduction. *Opt. Express* **21**, 11228 (May 2013).
- 44. Cohen, J. D. Fiber-optic integration and efficient detection schemes for optomechanical resonators. PhD thesis (California Institute of Technology, 2015).
- 45. Meenehan, S. M. *Cavity Optomechanics at Millikelvin Temperatures*. PhD thesis (California Institute of Technology, 2015).
- 46. Johnson, S. & Joannopoulos, J. D. Block-iterative frequency-domain methods for Maxwell's equations in a planewave basis. *Opt. Express* **8**, 173–190 (2001).
- 47. Yu, S.-P., Hood, J. D., Muniz, J. A., Martin, M. J., Norte, R., Hung, C.-L., Meenehan, S. M., Cohen, J. D., Painter, O. & Kimble, H. J. Nanowire photonic crystal waveguides for single-atom trapping and strong light-matter interactions. *Appl. Phys. Lett.* **104**, 111103 (Mar. 2014).
- 48. Yu, S.-P. PhD thesis (California Institute of Technology, 2017).
- 49. Joannopoulos, J. D., Johnson, S. G., Winn, J. N. & Meade, R. D. *Photonic Crystals: Molding the Flow of Light*. Second (Princeton University Press, 2008).
- Johnson, S. G., Bienstman, P., Skorobogatiy, M. A., Ibanescu, M., Lidorikis, E. & Joannopoulos, J. D. Adiabatic theorem and continuous coupled-mode theory for efficient taper transitions in photonic crystals. *Phys. Rev. E* 66, 1–15 (2002).
- Alasaarela, T., Korn, D., Alloatti, L., Säynätjoki, A., Tervonen, A., Palmer, R., Leuthold, J., Freude, W. & Honkanen, S. Reduced propagation loss in silicon strip and slot waveguides coated by atomic layer deposition. *Opt. Express* 19, 11529–38 (June 2011).
- Söllner, I., Mahmoodian, S., Hansen, S. L., Midolo, L., Javadi, A., Kiršanske, G., Pregnolato, T., El-Ella, H., Lee, E. H., Song, J. D., Stobbe, S. & Lodahl, P. Deterministic photon–emitter coupling in chiral photonic circuits. *Nat. Nanotechnol.* 10, 775–778 (2015).
- 53. Markoš, P. & Soukoulis, C. M. *Wave Propagation: From Electrons to Pho*tonic Crystals and Left-Handed Materials. (Princeton University Press, 2008).

- 54. Novotny, L. & Hecht, B. *Principles of Nano-Optics*. (Cambridge University Press, 2008).
- 55. Born, M. & Wolf, E. *Principles of Optics*. Seventh (Cambridge University Press, 2005).
- 56. Bendickson, J., Dowling, J. & Scalora, M. Analytic expressions for the electromagnetic mode density in finite, one-dimensional, photonic band-gap structures. *Phys. Rev. E* **53**, 4107–4121 (Apr. 1996).
- 57. Mutapcic, A., Boyd, S., Farjadpour, A., Johnson, S. G. & Avniel, Y. Robust design of slow-light tapers in periodic waveguides. *Eng. Optimiz.* **41**, 365–384 (2009).
- 58. Luke, K., Poitras, C. B. & Lipson, M. Overcoming SiN film stress limitations for high quality factor ring resonators. *Opt. Express* **21**, 22829 (2013).
- 59. Marcuse, D. *Light Transmission Optics*. Second (Van Nostrand Reinhold Company Inc., 1982).
- 60. Snyder, A. W. & Love, J. D. Optical Waveguide Theory. (Springer, 1983).
- 61. Marcuse, D. Loss Analysis of Single-Mode Fiber Splices. *Bell System Technical Journal* **56**, 703 (1977).
- 62. Norte, R. A. Nanofabrication for On-Chip Optical Levitation, Atom-Trapping, and Superconducting Quantum Circuits. PhD thesis (California Institute of Technology, 2015).
- 63. Dobbs, D. A., Bergman, R. G. & Theopold, K. H. Piranha solution explosion. *Chem. Eng. News* 68, 2–3 (17 1990).
- 64. *Technical Data Sheet: AZ 5214 E Photoresist*. Merck Performance Materials GmbH (Rheingaustrasse 190–196, D-65203 Wiesbaden, Germany).
- 65. *EBPG5000+/5200 Operator Manual*. Vistec Lithography, B.V. (De Dintel 27a, 5684PS Best, The Netherlands, 2012).
- 66. Reimer, L. Scanning Electron Microscopy: Physics of Image Formation and Microanalysis. Second (Springer, 1998).
- 67. BEAMER v4.7.0 Manual. GenISys GmbH (2014).
- 68. Topolancik, J., Ilic, B. & Vollmer, F. Experimental Observation of Strong Photon Localization in Disordered Photonic Crystal Waveguides. *Phys. Rev. Lett.* **99**, 253901 (Dec. 2007).
- 69. García, P. D., Smolka, S., Stobbe, S. & Lodahl, P. Density of states controls Anderson localization in disordered photonic crystal waveguides. *Phys. Rev. B* 82, 165103 (Oct. 2010).
- Seidel, H., Csepregi, L., Heuberger, A. & Baumgärtel, H. Anisotropic Etching of Crystalline Silicon in Alkaline Solutions. *J. Electrochem. Soc.* 137, 249– 260 (1990).

- Okada, T., Fujimori, J., Aida, M., Fujimura, M., Yoshizawa, T., Katsumura, M. & Iida, T. Enhanced resolution and groove-width simulation in cold development of ZEP520A. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures* 29, 021604 (2011).
- Eichfeld, C. M. & Lopez, G. G. Temperature dependent effective process blur and its impact on exposure latitude and lithographic targets using ebeam simulation and proximity effect correction. *J. Vac. Sci. Technol. B* 32, 06F503 (Nov. 2014).
- 73. Salvat, F., Fernández-Varea, J. & Sempau, J. *PENELOPE-2011: A Code* System for Monte Carlo Simulation of Electron and Photon Transport. in Workshop Proceedings (2011).
- 74. Schneider, C. A., Rasband, W. S. & Eliceiri, K. W. NIH Image to ImageJ: 25 years of image analysis. *Nat. Methods*, 671–675 (2012).
- 75. Henry, M. D. *ICP etching of silicon for micro and nanoscale devices*. PhD thesis (California Institute of Technology, 2010).
- 76. Jansen, H. V., de Boer, M. J., Unnikrishnan, S., Louwerse, M. C. & Elwenspoek, M. C. 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).
- 77. Williams, K. R., Gupta, K. & Wasilik, M. Etch rates for micromachining processing–part II. *J. Microelectromech. S.* **12**, 761–778 (Dec. 2003).
- 78. Ritala, M. & Niinistö, J. in *Chemical Vapour Deposition: Precursors, Processes and Applications* (Royal Society of Chemistry, 2009).
- Saleem, M. R., Ali, R., Khan, M. B., Honkanen, S. & Turunen, J. Impact of Atomic Layer Deposition to Nanophotonic Structures and Devices. *Frontiers in Materials* 1, 1–15 (Oct. 2014).
- Graugnard, E., Gaillot, D. P., Dunham, S. N., Neff, C. W., Yamashita, T. & Summers, C. J. Photonic band tuning in two-dimensional photonic crystal slab waveguides by atomic layer deposition. *Appl. Phys. Lett.* 89, 181108 (2006).
- Yang, X., Chen, C. J., Husko, C. A. & Wong, C. W. Digital resonance tuning of high-Q/V_m silicon photonic crystal nanocavities by atomic layer deposition. *Appl. Phys. Lett.* **91**, 161114 (2007).
- Kiravittaya, S., Lee, H. S., Balet, L., Li, L. H., Francardi, M., Gerardino, A., Fiore, A., Rastelli, A. & Schmidt, O. G. Tuning optical modes in slab photonic crystal by atomic layer deposition and laser-assisted oxidation. *J. Appl. Phys.* 109, 053115 (2011).

- Bera, A., Häyrinen, M., Kuittinen, M., Honkanen, S. & Roussey, M. Parabolic opening in atomic layer deposited TiO₂ nanobeam operating in visible wavelengths. *Opt. Express* 23, 14973 (2015).
- Häyrinen, M., Roussey, M., Säynätjoki, A., Kuittinen, M. & Honkanen, S. Titanium dioxide slot waveguides for visible wavelengths. *Appl. Optics* 54, 2653–2657 (2015).
- 85. Khanna, A., Subramanian, A. Z., Häyrinen, M., Selvaraja, S., Verheyen, P., Van Thourhout, D., Honkanen, S., Lipsanen, H. & Baets, R. Impact of ALD grown passivation layers on silicon nitride based integrated optic devices for very-near-infrared wavelengths. *Opt. Express* 22, 5684–92 (Mar. 2014).
- Gehl, M., Gibson, R., Hendrickson, J., Homyk, A., Säynätjoki, A., Alasaarela, T., Karvonen, L., Tervonen, A., Honkanen, S., Zandbergen, S., Richards, B. C., Olitzky, J. D., Scherer, A., Khitrova, G., Gibbs, H. M., Kim, J.-Y. & Lee, Y.-H. Effect of atomic layer deposition on the quality factor of silicon nanobeam cavities. J. Opt. Soc. Am. B 29, A55 (2012).
- Du, X., Zhang, K., Holland, K., Tombler, T. & Moskovits, M. Chemical corrosion protection of optical components using atomic layer deposition. *Appl. Optics* 48, 6470–4 (Nov. 2009).
- Woetzel, S., Talkenberg, F., Scholtes, T., IJsselsteijn, R., Schultze, V. & Meyer, H.-G. Lifetime improvement of micro-fabricated alkali vapor cells by atomic layer deposited wall coatings. *Surf. Coat. Technol.* 221, 158–162 (Apr. 2013).
- Weckman, T. & Laasonen, K. First principles study of the atomic layer deposition of alumina by TMA-H₂O-process. *Phys. Chem. Chem. Phys.* 17, 17322–17334 (2015).
- 90. Eickhoff, W., Yen, Y. & Ulrich, R. Wavelength dependence of birefringence in single-mode fiber. *Appl. Optics* **20**, 3428–3435 (1981).
- 91. Siegman, A. E. Lasers. (University Science Books, 1986).
- 92. Riley, P. E. Comparison of etch rates of silicon nitride, silicon dioxide, and polycrystalline silicon upon O2 dilution of CF4 plasmas. *J. Vac. Sci. Technol. B* **7**, 1352 (Nov. 1989).
- 93. Muniz, J. A. *Nanoscopic atomic lattices with light-mediated interactions*. PhD thesis (California Institute of Technology, 2017).
- 94. *Scanning Electron Microscope A To Z.* JEOL Ltd. (1-2 Musashino 3-chome Akishima Tokyo 196-8558 Japan, 2009).
- 95. *Macor*®. Corning Incorporated (Houghton Park CB-08, Corning, New York, 14831, 2012).
- 96. Hidnert, P. & Krider, H. Thermal expansion of aluminum and some aluminum alloys. *J. Res. Nat. Bur. Stand.* **48**, 209 (1952).

- Middelmann, T., Walkov, A., Bartl, G. & Schödel, R. Thermal expansion coefficient of single-crystal silicon from 7 K to 293 K. *Phys Rev. B* 92, 174113 (Nov. 2015).
- 98. Connolly, E. W. *Experiments with Toroidal Microresonators in Cavity QED*. PhD thesis (California Institute of Technology, 2009).
- 99. Abraham, E. R. & Cornell, E. A. Teflon feedthrough for coupling optical fibers into ultrahigh vacuum systems. *Appl. Optics* **37**, 1762–1763 (1998).
- Aschroft, N. W. & Mermin, N. D. Solid State Physics. (Thomson Learning, Inc., 1976).
- Ritter, R., Gruhler, N., Pernice, W., Kübler, H., Pfau, T. & Löw, R. Atomic vapor spectroscopy in integrated photonic structures. *Appl. Phys. Lett.* 107, 041101 (2015).
- 102. Ritter, R., Gruhler, N., Pernice, W. H. P., Kübler, H., Pfau, T. & Löw, R. Coupling thermal atomic vapor to an integrated ring resonator. *New J. Phys.* 18, 103031 (2016).
- 103. Hatakeyama, A., Wilde, M. & Fukutani, K. Classification of Light-Induced Desorption of Alkali Atoms in Glass Cells Used in Atomic Physics Experiments. *e.-J. Surf. Sci.* 4, 63–68 (2006).
- 104. Stephens, M., Rhodes, R. & Wieman, C. Study of wall coatings for vapor-cell laser traps. *J. Appl. Phys.* **76**, 3479–3488 (1994).
- 105. Hartung, J. *Optical microfibre surface protection for caesium spectroscopy*. MA thesis (Friedrich-Wilhelms-Universität Bonn, 2012).
- Seltzer, S. J., Rampulla, D. M., Rivillon-Amy, S., Chabal, Y. J., Bernasek, S. L. & Romalis, M. V. Testing the effect of surface coatings on alkali atom polarization lifetimes. *J. Appl. Phys.* **104** (2008).
- 107. Blasius, T. D. *Optomechanical Inertial Sensors and Feedback Cooling*. PhD thesis (California Institute of Technology, 2016).
- 108. Foot, C. J. Atomic Physics. First (Oxford University Press, 2005).
- 109. Wohlleben, W., Chevy, F., Madison, K. & Dalibard, J. An atom faucet. *Eur. Phys. J. D* **15**, 237–244 (2001).
- Dimova, E., Morizot, O., Stern, G., Garrido Alzar, C. L., Fioretti, A., Lorent, V., Comparat, D., Perrin, H. & Pillet, P. Continuous transfer and laser guiding between two cold atom traps. *Eur. Phys. J. D* 42, 299–308. ISSN: 1434-6060 (Feb. 2007).
- 111. Shang, S.-Q. and Sheehy, B. and Metcalf, H. and van der Straten, P. and Nienhuis, G. Velocity-selective resonances and sub-Doppler laser cooling. *Phys. Rev. Lett.* 1094–1097 (1991).

- 112. Steane, A. M., Chowdhury, M. & Foot, C. J. Radiation force in the magnetooptical trap. *J. Opt. Soc. Am. B* **9**, 2142 (Dec. 1992).
- 113. Goban, A. Strong atom-light interactions along nanostructures: Transition from free-space to nanophotonic interfaces. PhD thesis (California Institute of Technology, 2015).
- 114. Ketterle, W., Durfee, S. & Stamper-Kurn, D. M. Making, probing and understanding Bose-Einstein condensates. arXiv: 9904034v2. http://arxiv. org/abs/cond-mat/9904034> (1999).
- 115. Alessi, P. J., Carter, E. C., Fairchild, M. D., Hunt, R. W. G., McCamy, C. S., Kranicz, B., Moore, J. R., Morren, L., Nobbs, J. H., Ohno, Y., Pointer, M. R., Rich, D. C., Robertson, A. R., Schanda, J. D., Seve, R., Trezona, P. W., Witt, K. & H, Y. *CIE Technical Report: Colorimetry*. tech. rep. (International Commission on Illumination, 2004).

Appendix A

TIMELINE OF CHIPS IN LAB 2

Figure A.1 shows a history of chips installed in the Lab 2 chamber, with events with personal or project-level significance indicated. In mid-2014 we made radical changes to atom delivery. Prior to this point atoms were transferred to the chip via a MOT "toss and catch" method, in which the magnetic field zero was moved from the science MOT to a location in the chip. The direct digital synthesis (DDS) necessary for conveying atoms is achieved using an Analog Devices AD9854 chip. The synthesizer is interfaced with an Arduino Due using the Serial Peripheral Interface (SPI); my Arduino code is based on code Mike Martin brought us from CU for single-tone synthesis. This was my last major project in Lab 2 before becoming active in device fabrication. My first solo operation of the EBPG5000+ occurred in early 2015. In June 2015 the Ultratech Fiji G2 ALD system was purchased and installed; all science chips after this have been coated with ALD-alumina.



Figure A.1: **Timeline of chips in Lab 2 vacuum chamber.** Red line indicates the beginning of conveyor delivery. Green line indicates the date I was signed off to operate the EBPG. Blue line indicates the beginning of the "ALD" era.

MPB SCRIPTS

MIT Photonic Bands (MPB) is a eigenmode solver for Maxwell's equations in periodic dielectrics [46].ⁱ There are many free and commercial tools which can perform these types of calculation, each with particular advantages; in Quantum Optics Group, most have elected to use COMSOL. I chose to use MPB instead for a few reasons: as software without a license, it is portable, and there is no limit on how many instances can be run concurrently; it has a text-based, command-line interface, simplifying version control, scripting, and remote execution; it is a domain-specific tool, as opposed to COMSOL, resulting in useful, built-in functionality; and, unlike COMSOL 3.5a, it could be installed on my personal computer.

These virtues aside, the interface can be quite frustrating, if Scheme (and the functional programming paradigm more generally) is unfamiliar. I hope the scripts provided can supplement the examples on the MPB wikiⁱⁱ to accelerate the learning process.

B.1 APCW Unit Cell

It is straightforward to create photonic crystal geometries consisting of rectangular parallelpipeds and cylinders in the MPB interface language, but the APCW unit cell is a more sophisticated shape. The code below provides a minimal working example, which draws the unit cell using tail recursion and calculates band edges for an APCW.

```
;;; CONSTANTS ;;;
(define pi (* 2 (acos 0))) ; pi
(define ix_SiN 1.9935) ; index of Si3N4
;;; PHC DIMENSIONS ;;;
(define-param amp 120) ; nominal amplitude (nm)
(define-param wid 280) ; nominal width (nm)
(define-param thk 200) ; nominal thickness (nm)
(define-param gap 238) ; nominal gap (nm)
```

ⁱWaveguides with constant cross-section are a special case of periodic dielectrics, having zero period.

iihttp://ab-initio.mit.edu/wiki/index.php/MIT_Photonic_Bands

```
(define-param a 370) ; lattice constant (nm)
;;; AUX SETTINGS ;;;
(set-param! resolution 37) ; number of blocks per
   lattice constant
(define dx (/ 1 resolution 10)) ; thickness of blocks;
   sub-resolution
                       ; filename prefix
(set! filename-prefix
     (string-append
      "w" (number->string wid)
      "A" (number->string amp)
      "g" (number->string gap)
      )
     )
;;; CELL DIMENSIONS ;;;
(define sx 1) ; lattice constant in lattice constants
(define-param sy 15) ; y dimension in lattice constants
(define-param sz 15) ; z dimension in lattice constants
(set! geometry-lattice (make lattice (size sx sy sz)))
;;; DEFINE UNIT CONVERSION FUNCTIONS ;;
(define (simUnits x) (/ x a))
;;; CREATE GEOMETRY ;;;
(define (drawBlock x)
  (let (
       (dy
        (+ wid (* amp (cos (* 2 pi x)))); block width
        )
       )
   (let (
         (cy
          (/ (+ dy gap) 2)
                                ; block center
          )
         )
     (list
      (make block
                                         ; top
```

```
(center x (simUnits cy) ∅)
         (size (* 1.1 dx) (simUnits dy) (simUnits thk))
         (material (make dielectric (index ix_SiN)))
         )
       (make block
                                          ; bottom
         (center x (* −1 (simUnits cy)) 0)
         (size (* 1.1 dx) (simUnits dy) (simUnits thk))
         (material (make dielectric (index ix_SiN)))
         )
      )
     )
   )
 )
(set! geometry '()) ; initialize empty geometry
(define (makeGeom x)
  (if (>= x 1))
      '()
                       ; exit loop
      (begin
       (set! geometry (append geometry (drawBlock x)))
       (makeGeom (+ x dx))
       )
     )
 )
(makeGeom 0) ; recursively fill geometry
;;; RUN
;; set k-points, num-bands
(set! k-points (list (vector3 0.5 0 0))); only calculate
   band edge
(set! num-bands 2)
;;; RUN ;;;
(run-yodd-zeven) ; TE
```

B.2 Mode Overlap Calculation

This script uses MPB to calculate the TE modes for a Nufern 780HP optical fiber and for a rectangular coupler, and then performs the mode overlap calculation described in sec. 2.5.

```
;;; COMMAND LINE PARAMETERS ;;;
(define-param frq 343)
                              ; frequency in THz
(set-param! resolution 10)
(define-param wid 120)
                               ; waveguide width in nm
(define-param thk 200)
                            ; waveguide thickness in nm
;;; DO FIBER CALCULATION ;;;
;; define dimensions
(define core-diameter 4400)
                              ; 4.4 um core diameter
(define core-radius (/ core-diameter 2))
(define core-index 1.4628)
;; indices given by Nufern (private correspondence)
(define cladding-index 1.4570)
(define a 100)
                                ; mpb unit, here 100 nm
(define c (* 299792458e9 1e-12)); speed of light (THz nm)
;; define helper functions
(define (toMPBunits x) (/ x a))
;; convert to mpb units
(set! frq (/ frq (/ c a)))
(set! core-radius (toMPBunits core-radius))
;; define parameters
(define sx 120)
                              ; size of simulation region
;; set up space
(set!
geometry-lattice
(make lattice (size no-size sx sx))
)
;; define geometry
(set! geometry (list
```

```
(make block ; cladding
                  (center 0 0)
                  (size infinity sx sx)
                  (material (make dielectric (index
                     cladding-index)))
                  )
                (make cylinder ; core
                  (radius core-radius)
                  (center 0 0)
                  (height infinity)
                  (axis (vector3 1 0 0))
                  (material (make dielectric (index
                     core-index)))
                  )
                )
      )
;; call find-k
(find-k TE frq 1 1 (vector3 1 0 0) 1e-4 0.1 0 1)
(define fiber-frq (list-ref freqs 0))
;; store fiber fields
(get-hfield 1)
(define fiber-h-field (field-copy cur-field))
(get-dfield 1)
(get-efield-from-dfield)
(define fiber-e-field (field-copy cur-field))
  ;;; DO WAVEGUIDE CALCULATION ;;;
;; define dimensions
(define ix 1.99)
                               : SiN refractive index
;; convert to mpb units
(set! wid (toMPBunits wid))
(set! thk (toMPBunits thk))
;; define geometry
(set! geometry (list
```

```
(make block
                  (center 0 0)
                  (size infinity wid thk)
                  (material (make dielectric (index ix)))
                  )
                )
     )
;; call find-k
(find-k TE frq 1 1 (vector3 1 0 0) 1e-4 0.1 0 1)
(define waveguide-frq (list-ref freqs 0))
;; store waveguide fields
(get-hfield 1)
(define waveguide-h-field (field-copy cur-field))
(get-dfield 1)
(get-efield-from-dfield)
(define waveguide-e-field (field-copy cur-field))
  ;;; DO MODE OVERLAP CALCULATION ;;;
(define (field-calc e h)
  (field-map! cur-field
                               ; put in cur-field...
              (lambda (a b)
                (vector3-cross a (vector3-conj b))
                )
                                ; ...this vector function...
              e h)
                                ; ...mapped over these
                 fields
  (integrate-fields (lambda (r f) (vector3-x f)) cur-field)
 )
(print "***\n")
(print "performing mode overlap calculation...\n")
(print "***\n")
(define A (field-calc fiber-e-field fiber-h-field))
(define B (field-calc fiber-e-field waveguide-h-field))
(define C (field-calc waveguide-e-field fiber-h-field))
(define D (field-calc waveguide-e-field waveguide-h-field))
```

129

```
(define efficiency
 (/
  (real-part (/ (* B C) D))
  (real-part A)
  )
 )
(print "overlap: " efficiency "\n")
(print "fiber frequency: " (* fiber-frq (/ c a)) "\n")
(print "waveguide frequency: " (* waveguide-frq (/ c a))
  "\n")
```

Appendix C

FILM THICKNESSES

The apparent color of thin films changes depending on film thickness as a result of interference effects. This allows us to measure film thicknesses using the spectral reflectance method (see fig. 3.9). These changes are also apparent by eye. Figure C.1 shows approximate colors for different nitride films at normal incidence under white light illumination. These colors were calculated using refractive index data for silicon nitride and silicon, and published tristimulus values [115]. While these colors should be rendered correctly by a calibrated monitor or printer, they are provided only as a qualitative guide.



Figure C.1: **Silicon nitride film color.** Approximate color for different thicknesses on silicon substrate.

Appendix D

COLD DEVELOP

It is known that reducing the temperature of the developer can improve process latitude in *e*-beam lithography [72]. Transitioning our process to cold develop would require some calibration, but may well be worth it in the end. Here are initial efforts to determine the dose-to-clear for a cold develop process.

I exposed a pattern of 120 μ m swatches, with dose ranging 120–720 μ C/cm² in steps of 25 μ C/cm². The pattern was developed in ZED-N50 chilled in ice water to ca. 0 °C, and compared to a pattern developed with our standard, room-temperature process. Figure D.1 shows microscope images of the swatches. All of the swatches developed at room temperature (23 °C) cleared, while on the sample developed at 0 °C only those dosed above 345 μ C/cm² cleared. The increased dose-to-clear is consistent with published cold-develop dose-to-clear curves [71, 72]. Qualitatively it is also noteworthy that the edge effects apparent in the room-temperature sample are not present in the cold develop sample, consistent with a steeper dose-to-clear curve.



Figure D.1: Cold develop test. 120 μ m squares dosed 120–720 μ C/cm² (upper right to bottom left, raster scan; step 25 μ C/cm²). Developed (ZED-N50 150s, MIBK 30s) at (a) 25 °C and (b) 0 °C.