

Nano-Photonic Platform for Atom-Light Interaction

Thesis by
Su-Peng Yu

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



CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2017
Defended May 26th, 2017

© 2017

Su-Peng Yu

ORCID: 0000-0003-1348-7447

All rights reserved

ACKNOWLEDGEMENTS

I would like to acknowledge with sincere gratitude the collaboration and support from my friends and colleagues that had made this project possible: My advisor, Prof. Jeff Kimble, has always been supportive to my work and my various crazy ideas. I'd also like to acknowledge Prof. Oskar Painter for the resources and effort he provided for this collaboration between atoms and photonics. Chen-Lung Hung, who motivated many of our efforts; Michael Martin, who is great with both optics and electronics; Jae-Hoon Lee, with his attention to the details; Mingwu Lu, who is knowledgeable and never hesitant to share; Alex Burgers, with whom working is never boring; Dalziel Wilson, who inspired my interest in machining; Akihisa Goban, who seems to get anything working in a breeze; Jonathan Hood and Andrew McClung, with whom I cooked various chemicals together that were often both poisonous and corrosive; Juan Muniz, who traps atoms and also makes the best barbecue; Lucas Peng, whose programming skill allowed our project better analysis; and Xingsheng Luan, whose photonic designs are even crazier than mine. I would also like to acknowledge the people who helped me a lot in the Painter group: Richard Norte and Justin Cohen, who taught me how to do nano-fab; and Sean Meenehan, who is known for great simulation skills and appetite for coffee. I'd like to acknowledge the work of our theoretical collaborators, Prof. Darrick Chang, who answers every question with clarity and patience; and Ana Asenjo-Garcia, who is always around for a fruitful discussion. Finally, I'd like to offer my thanks to Barry Baker for keeping the cleanroom clean and well-supplied, Guy DeRose for keeping the lithography machine always up and running, and Scott Curtis for keeping track of our immense amount of accounts and purchases.

ABSTRACT

Development of technology can allow access to new regimes in science and creation of new fields of research. The development of ultra-high finesse mirror technology enabled the development of the field of cavity quantum electrodynamics, and an abundance of wonderful physics experiments soon followed. The sophistication of the field of nano-fabrication and nano-photonics would allow unprecedented capability to mold the shape and flow of light, and provide a novel platform for efficient and hopefully integrable quantum systems. In this project, we hope to interface cold atoms, perhaps the most quintessential of quantum systems, to the engineering power of nano-photonics. We believe this field of study will not only lead to the demonstration of new physics in the quantum regime, but work toward building a network with quantum capabilities mediated by optical channels.

In this project, we develop a nano-fabricated platform capable of interfacing nano-photonic devices with cold Cesium atoms in free-space. Nano-photonic waveguide devices are fabricated in a Silicon Nitride device layer on Silicon substrate. The fabrication is compatible with conventional semiconductor fabrication processes, and the chip design has been adapted to allow incorporation with free-space optics to support cold Cesium atom cloud around the waveguides. An ultra-high vacuum system that is compatible to the chip and its supporting structures was constructed to perform experiments.

With our system, we were able to fabricate and characterize nano-photonic structures, including 1D photonic crystal waveguides, cavities, and 2D photonic crystal slabs. For the 1D photonic crystal waveguide devices, enhanced atom-light coupling between localized Cesium atoms in the vicinity of the devices, and also atom-atom interaction between Cesium atoms mediated by the guided mode of the photonic crystal waveguide, has been observed. The 2D photonic devices allow us many capabilities beyond that of the 1D waveguide. Demonstration of exotic optical properties including natural decay rate suppression and circular polarization engineering, should be within reach in the near future.

PUBLISHED CONTENT AND CONTRIBUTIONS

S.-P. Yu et al. “Nanowire Photonic Crystal Waveguides for Single-Atom Trapping and Strong Light-Matter Interactions”. In: *Applied Physics Letters* 104 (2014), p. 111103. DOI: 10.1063/1.4868975

A. Goban et al. “Atom–Light Interactions in Photonic Crystals”. In: *Nature Communications* 5 (2014), p. 3808. DOI: 10.1038/ncomms4808.

A. Goban et al. “Superradiance for Atoms Trapped along a Photonic Crystal Waveguide”. In: *Physical Review Letters* 115 (2015), p. 063601. DOI: 10.1103/PhysRevLett.115.063601.

J. D. Hood et al. “Atom–Atom Interactions Around the Band Edge of a Photonic Crystal Waveguide”. In: *Proceedings to National Academy of Science* 113.38 (2016), pp. 10507–10512. DOI: 10.1073/pnas.1603788113.

The project of atom-light interaction with nanophotonic devices was made possible only by the combined effort of all members of Prof. Kimble’s group, and the collaboration with Prof. Painter’s group. My role in this project includes: simulation and device design, in collaboration with Juan Muniz, and Sean Meenehan; the developing of fabrication process and the fabrication of the devices, in collaboration with Jonathan Hood, Andrew McClung, Richard Norte, and Justin Cohen; the optical characterization of the devices and the developing of the chip-fiber assembly design, which was a joint effort with Jonathan Hood and Michael Martin. All these made the results in Yu et al. (2014) possible. I worked intermittently with the atom measurement lab, with Akihisa Goban, Jonathan Hood, Chen-Lung Hung, and Mingwu Lu. It was primarily their work that produced the results in Goban et al. (2014), Goban et al. (2015), and Hood et al. (2016). My work for these experiments were mostly with the building of the devices, and perhaps some minor contribution to the day-to-day work of the optics lab.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	iv
Published Content and Contributions	v
Table of Contents	vi
List of Illustrations	x
List of Tables	xiii
Chapter I: Motivation	1
1.1 Nanophotnics and Optical Property Engineering	1
1.2 Light Interaction with Quantum Emitters	1
Achieving strong coupling to quantum emitter	2
Light-mediated emitter interactions	2
Emitter-mediated light interactions	3
1.3 Interfacing Hybrid Systems	3
Quantum networks	3
1.4 Thesis Outline	4
Chapter II: Support Structure Design	6
2.1 Design Targets for Atom-Nanophotonics Hybrid System	6
Chip layout to allow optical access	6
Power handling	7
Vacuum compatibility	7
2.2 High-Stress Silicon Nitride and Mechanical Stress Handling	8
Avoiding stress concentration	8
Stress balancing to prevent deformation	9
2.3 Thermal Conductance Structures	9
2.4 Connecting Waveguides to Outside World	11
FEM calculation and mode projection	12
Direct time-domain simulation	13
2.5 Anti-Reflection Designs on Highly Evanescent Waveguides	14
'Round-about' anti-reflection tether junction	15
Y-shaped free-space coupling tether design	17
2.6 Assembly of Fiber-Coupled Chip	17
Notes on using glue	18
Fiber feedthrough and strain handling	19
2.7 Coupling into Free-Space Lens-Coupled Chip	20
Chip configuration adaption to allow free-space coupling	21
Vacuum and optics system consideration	22
Chapter III: Fabrication of Devices	24
3.1 Fabrication Requirements	24
Precision of device dimensions	25

Device footprint and throughput	25
Issues associated with wet chemistry processes	25
3.2 Substrate Selection and Preparation	27
Pre-patterning the backside for through-window etch	28
3.3 Electron Beam Lithography with Large Device Footprint	28
Coarse-fine double exposure	29
Back-scattering electron	29
Main-field stitching	31
3.4 Etching	32
Plasma etching for pattern transfer	32
Directional wet etching of Silicon substrate	33
3.5 Applications of Silicon Anisotropic Etching	34
Etch-in of positive corners	35
'Fuse' timing method for residual Silicon control	35
Silicon half-aperture for stray-light protection	37
Etched precision cleave grooves	38
Substrate rotation check patterns	39
3.6 Wet-Chemistry Processing	40
Chemical processing holder design	41
Chip installation onto the holder	42
Etching of the through-window	42
Device chemical cleaning	43
Device drying and surface tension	44
Manipulation of devices using SEM charging effect	45
3.7 Post-Release Plasma Cleaning	46
3.8 Characterization using photonic crystal cavities	47
3.9 Metal Deposition for Potential Electrostatic Tuning	48
3.10 Device Fine-Tuning Using Atomic Layer Deposition	49
3.11 Experiments with High-Temperature Annealing	50
Thermal treatment to improve yield	50
Post-release annealing	50
Chapter IV: Engineering Optical Properties with APCW	52
4.1 Engineering Atom-Light Coupling Using Photonic Crystal	52
Mode profile of the photonic crystal waveguide	53
Coupling enhancement on band-edges	53
4.2 Placement of Photonic Crystal Frequency	56
FDTD simulation and Green's function	56
4.3 Tapering Photonic Crystal Waveguides	57
Envelop from residual taper reflection	58
Phase-matched taper design	58
4.4 Cesium Dual-Frequency Operation	60
Alignment of PhC frequencies to Cs D1 and D2 lines	60
Two-mode trap-and-probe scheme	60
Side-illumination trapping scheme	61
4.5 Additional Engineering Options	62

'TM' modes	62
Anti-symmetric modes	63
Circular polarization engineering	63
Chapter V: Optical Characterization of Devices	68
5.1 Reflection and Transmission Measurement	68
Light sources and analysis tool	68
Polarization alignment	69
Measurement using cleaved fiber-end	70
Measurement using free-space optics with cleaved chips	71
5.2 Features of Measured Spectra	73
Power level and band-gap alignment	73
Band-edge resonance regularity and quality	73
Residual transmission through the band-gap	74
5.3 Interaction Enhancement Estimation	74
Using free spectral range of resonances	75
Using scattering measurement	76
5.4 Power Handling Capability Test in Vacuum	76
5.5 Precision shifting to Cs frequencies using optical measurement feed-back	77
5.6 Characterization of Single-Beam Fabry-Perot-Type Cavity	77
Separation of propagation and mirror loss	78
Reachable parameters and potential for cavity QED experiments	78
Chapter VI: Atom-Light Interaction Measurement	85
6.1 Vacuum System Adaption for Photonic Devices	85
Background pressure of Cesium	85
Limitation associated with interfacing chip to outside optics	86
Replacement of chips in vacuum chamber	86
6.2 Vacuum System Design and Atom Delivery	86
Dual-chamber design with atom 'faucet'	87
Atom collecting and delivery in the 'Science' chamber	88
6.3 Methods for Measurement of Atom-Light Interaction	89
Frequency domain spectroscopy measurement	89
Time-domain decay rate measurement	90
6.4 Observation of Enhanced Single-Atom Decay Rate	91
Results	91
Discussion	92
6.5 Observation of Light-Mediated Atom-Atom Interaction	92
Result: Superradiance	93
Result: Band-gap mediated interaction	93
Discussion	95
Chapter VII: Engineering Optical Properties with 2D Photonic Crystal Slabs	97
7.1 Adaptation of Support Structures to Photonic Crystal Slabs	97
Mechanical deformation of the 2D structure	97
Coupling 2D structures to input waveguide	98
7.2 Optical Properties of 2D PhC Modes	98

TE bandgap and free-space decay suppression	98
Dual-frequency directional emission	99
Circular modes at K point air band	100
Atom trapping in the 2D lattice hole	100
7.3 Optical Properties of 2D PhC with Defects	101
Point-defect super-suppressor	101
1D defect slot waveguide	101
7.4 Atom-Light Interaction Parameters of the Slot Waveguide	102
Tapering Into the Slot Waveguide	102
Anti-symmetric mode	103
Green's function and state-dependent atom response	103
7.5 Characterization and Tuning of 2D-Based Devices	104
Chapter VIII: Outlook	116
8.1 Migration Toward Free-Space Coupled Devices	116
8.2 Deterministic Loading and Positioning with Optical Tweezers	117
8.3 Electro-capacitive dynamic tuning of slot waveguide	117
8.4 Potential Future Development	117
Circular Polarization Engineering	118
Multi-port devices	119
Potential for 'random-access' quantum system	120
Bibliography	122
Appendix A: FEM Simulation Methods	128
A.1 Optical Band Structures	129
Geometry and boundary conditions	129
Meshing and solving	129
Information output	130
A.2 Mechanical stress and resonance frequency	130
Geometry and boundary conditions	130
Meshing and solving	131
Information output	131
Appendix B: FDTD Simulation Methods	132
B.1 Transmission and Reflection	133
Source placement	133
Monitor placement	133
Field projection into guided mode	133
Data analysis and application	134
B.2 Green's Function Calculation	134
Source placement	134
Monitor placement	134
Data analysis and application	135
Appendix C: Dimensions for Chemistry Processing Holder	136
Appendix D: Typical Design and Fabrication Process Flow	137

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
1.1 Conceptual illustration of the Quantum Internet	4
2.1 Suspension of large, hundred- μm scale structures	8
2.2 Geometries that create local heightened stress	10
2.3 Tensile versus compressive corners	10
2.4 Mechanical design to handle deformation	11
2.5 Angled tether array for force balancing	12
2.6 Waveguide and cooling rail	13
2.7 Coupling waveguide mode projection into Gaussian beam profile . .	14
2.8 FDTD simulation plot of free-standing coupling waveguide	14
2.9 Reflection from coupling waveguide support tethers	15
2.10 Concept of hole-based dipole cancelation tether	16
2.11 Dipole cancellation by tilted tethers	18
2.12 Angled tether design field plot	19
2.13 Fiber placement and automated glue application setup	21
2.14 Chip holder assembly with fibers	22
2.15 Free-space direct-coupling test setup	23
3.1 Preliminary exploratory photonic crystals	26
3.2 Stitching error	27
3.3 Illustration of coarse-and-fine exposure with overlap area	30
3.4 Slice-grid method	31
3.5 Exposure and etch test patterns	34
3.6 Illustration of anisotropic etching of Silicon substrate	35
3.7 Full view of through-window chip and V-grooves	36
3.8 Fuse-protected positive angles	37
3.9 Silicon half-barrier formed by timed etching	38
3.10 Illustration of chip cleaving by fabricated cleave groove	40
3.11 Rotation check pattern	41
3.12 Teflon chip holder for wet chemistry processing	43
3.13 Congregate of solid contaminant removed by Hydrofluoric acid . . .	45
3.14 Large, multi- μm scale motion of double beam devices due to SEM charging	46

3.15	Re-condensation of organic material due to environmental contamination	47
3.16	Single-beam photonic crystal mirror	48
3.17	Signs of potential thermal damage to silicon substrate	51
4.1	Mode profile and band structure of the 'Alligator' photonic crystal waveguide	54
4.2	Formation of 1D band structure from waveguide dispersion curves . .	55
4.3	FDTD calculation of field profile and Green's function	59
4.4	Residual cavity induced resonances near band-edge	64
4.5	Dipole trapping potential and Casimir-Polder forces	65
4.6	Atom trapping using side-illumination	65
4.7	TM mode profile at band-edge	66
4.8	Shifted APCW to create very flat dispersion curve	66
4.9	Systematic engineering of circularly polarized 'pockets'	67
5.1	Residual reflection and interference fringes of different coupler designs	69
5.2	Optical image of coupling waveguide tether junction scattering	70
5.3	Illustration of the characterization setup	71
5.4	Optical image of coupling of device using free-space optics	72
5.5	Features of a device in reflection and transmission measurement . . .	75
5.6	Analysis of band-edge resonances using simple cavity model with quadratic dispersion	80
5.7	Features of a device in reflection and transmission measurement . . .	81
5.8	Device damaged in vacuum chamber	81
5.9	Shifting the band edges using ALD coating	82
5.10	Shifting the Nitride coupling waveguide optimal for 10nm ALD coating	83
5.11	Optical properties of the single-beam Fabry-Perot-type cavity	84
6.1	Illustration of vacuum system	87
6.2	Illustration of Experiment Cycle	90
6.3	Simulated distribution of Cesium atom near device	92
6.4	Measurements of superradiance in number-dependent decay	94
6.5	Measurements of Off-Resonance and Band-Gap Spectra	95
7.1	Mechanical support frame and tension redistribution tethering	105
7.2	Band structure and mode profile for TE modes of 2D photonic crystal	106
7.3	Dipole emission suppression with uniform 2D photonic crystal	107
7.4	Directional emission guided by photonic crystal dispersion	108
7.5	Circular polarization of the K-air modes	109

7.6	Moving blue-detuned lattice atom loading and trapping scheme . . .	109
7.7	Suppression of in-plane dipole emission with point defect	110
7.8	Embedded 1D defect 'slot' waveguide bands and mode profiles	111
7.9	Tapering slot waveguide	112
7.10	SEM image of slot waveguide taper	113
7.11	Atom decay rate calculation for the slot waveguide	114
7.12	Measured optical spectra of 2D and slot waveguide devices	115
8.1	Vibration modes of nano-string photonic crystal	118
8.2	Capacitive actuation tuning of double-beam device	119
8.3	Preliminary experiment of multi-port device	120
8.4	Potential development of devices with random access capability . . .	121
C.1	Dimensions for Chemistry Processing Holder	136
D.1	Flow chart of a typical design-to-device process	149

LIST OF TABLES

<i>Number</i>	<i>Page</i>
2.1 Summery of optical performance of coupling waveguide tethering designs	17
2.2 Outgassing properties of thermal and UV-cure glue	20
2.3 Lenses used for testing current generation free-space chips	23
4.1 Tuning rates of various tuning methods in THz/nm dimension	56
7.1 Geometric tuning rates for the 2D and slot waveguide devices	104

Chapter 1

MOTIVATION

This chapter will describe the background for development of our project. We hope to leverage the optical engineering capabilities of nanophotonics to create light-matter interactions beyond what was achievable in conventional macroscopic optics. The goal of this project is to provide a versatile and reliable platform to interface cold atoms with nano-photonics, and through that, with each other. Our goal is to achieve scalability of our system that will allow exploration of exotic quantum optical physics, and eventually provide a useful component toward a future of large-scale quantum networks.

1.1 Nanophotonics and Optical Property Engineering

The unique advantage of nanophotonics is its capability of manipulating light with sub-wavelength structures. The optical properties, such as mode profile and dispersion, can be engineered by patterning of material, allowing significantly more freedom than relying on bulk material and conventional optics elements. Efficient optical circuits can be formed by photonic waveguides, which are compact and alignment-free [1]. Strong frequency dependence can be engineered using micro-cavities or photonic crystal structures [2]. Even unconventional properties, such as negative refractive index, can potentially be engineered [3]. Each of these technologies was developed for their respective original application, but their capabilities can be leveraged readily into fields of research [4].

1.2 Light Interaction with Quantum Emitters

As science advances, we have access to many physical systems where quantum phenomena manifest naturally, from alkali atoms [5] to solid-state emitters such as quantum dots [6] and color centers [7]. An efficient way to probe these systems is to allow them to interact with light, converting the information within these systems into states of photons. The photons can then propagate out of the system in question, and be measured by suitable experiment apparatus, such as an Avalanche Photodiode (APD). The strength of the photon to carry information lies in its low loss properties preventing undesired interaction with the environment, such as the lenses or optical fibers that guides the photons in an optical setup. This isolating property

photons, however, creates difficulty in creating interaction between photonics and the quantum emitter in question.

Achieving strong coupling to quantum emitter

Various systems have been developed to enhance interaction strength between light and quantum emitters. One of the most successful methods was based on high-finesse optical cavities. In an intuitive picture, incoming light is reflected inside the cavity repeatedly, allowing the interaction between light and the quantum object to accumulate. In a more rigorous description, the sharp resonances of a cavity produce a high concentration of Local Density of State (LDOS) of light to which the emitter can interact with. Strong coupling, as typically characterized by Cooperativity $C = \frac{(2g)^2}{2\kappa\Gamma} \gg 1$ in cavity systems, can be achieved, as reported in [8]. The parameter g stands for the interaction strength between the quantum emitter and the cavity optical mode, where κ, Γ are the total decay rate of the cavity and the emitter itself, respectively. The coupling strength g increases as the mode volume of the cavity mode decreases $C \propto \frac{1}{\text{sqrt}V}$, favoring short cavities with concentrated mode profile. With the tools of nanophotonics, light can be tightly confined in small ($< \lambda^3$) volume, allowing enhanced coupling. Systems providing strong coupling between free-space atoms and light are actively studied [9]. Several devices with nanophotonic cavities strongly coupled to solid state quantum emitters have also been explored in literature [10].

Light-mediated emitter interactions

With sufficiently strong light-emitter interaction strength, it becomes possible to mediate interaction between multiple quantum emitters by exchanging photons. This can be done using cavity systems, where the quantum emitters collectively interact with the shared cavity mode, with the interaction strength varied by the local intensity determined by the cavity mode profile, as reported in [11]. A different set of physics can be observed in waveguide systems, such as in [12] and [13], where the continuum formed by the guided propagating modes become accessible to the quantum emitters. With systems such as atoms trapped near a waveguide, the motion of the atoms can also be manipulated by the guided mode through the atoms' optical response [14].

Emitter-mediated light interactions

The flip-side of this light-emitter coupling system is to mediate interaction between channels of light with strongly interacting emitters. Conventional nonlinear optical effects are usually driven at high power level [15], but single-photon level nonlinearity can be achieved [9], as observed as photon blockade in the cavity system, see [16]. Using one light channel to modulate another via a cavity-and-atom based system, effectively functioning as an 'optical transistor', has been demonstrated [17].

1.3 Interfacing Hybrid Systems

A constant struggle persists in the seemingly contradicting task of achieving stronger interaction strength and maintaining longer coherence time. It is then potentially advantageous to have 'hybrid' systems, where the strength of one system can be leveraged to cover the weakness of another, such as mapping and retrieval of quantum states between electronic states and nuclear spins [18]. As mentioned in section 1.2, it is advantageous to have the multiple quantum sub-systems be able to confer their information onto an optical channel. The channel can then be piped into a different sub-system, or can be directly measured. Extensive effort has been ongoing to connect radio-frequency waves to optical modes by mean of optomechanic systems [19, 20], which will eventually allow superconducting quantum circuits to be connected to neutral atom systems or color centers. A collection of single photons can be generated with a quantum dot system [21], for example, and then be injected into and processed by superconducting circuitry, and then stored in long-lifetime states in neutral atom systems. The combined system will provide functions beyond the capability of its parts.

Quantum networks

Beyond connecting sub-systems for their capabilities, we can imagine further interconnection between functional local systems into a larger network. Information can then be exchanged through efficient optical channels, and teleportation of quantum states [22] mediated by the channels. Technological application would naturally emerge at this stage of development, such as quantum communication and cryptography [23, 24]. With the rapid development of quantum information science, such a network can potentially allow future users to remotely access computation capability of 'servers' in a similar manner to accessing super-computing centers with the classical internet. Development of such 'Quantum Internet' would bring about a revolution in the sciences, and create significant technological impact to the

world.




Figure 1.1: Conceptual illustration of the Quantum Internet

Similar to the structure of a classical network, the quantum internet would be composed of nodes with quantum information processing and capability, as illustrated as boxes on this image. Each quantum node can have its strength, such as a high interaction strength $\chi \sim \frac{H_{int}}{\hbar}$ for processing, and low loss rate γ for storage in quantum memory. The different functional nodes are then interconnected by quantum communication channels. Image courtesy of H. J. Kimble [25].

1.4 Thesis Outline

This thesis will be arranged as follow: in Chapter 2, we describe the general structure of our photonic chip, from the macroscopic chip layout to methods for interfacing free-space optics to the nanophotonic devices. In Chapter 3, we describe the adaption of nano-fabrication and cleaning processes to fabricate our devices. The discussion will then focus on the quasi-1D 'Alligator' photonic crystal waveguide. Chapter 4 will describe the design of this waveguide and its optical properties. Chapter 5 will focus on the methods for optically testing our devices, and finally in Chapter 6 atom-light interaction measurements will be covered. Beginning from Chapter 7, the discussion will focus on suspended 2D photonic crystal slabs, and defect structures that can be embedded in them. Chapter 8 will provide an outlook from the immediate next steps to projecting what would be possible for the atom-nanophotonic platforms further on in the future. This thesis is structured in a way that each section is mostly self-contained. We hope this allows for ease of referencing

for future readers. Finally, step-by-step instructions for numerical simulation and design-to-device process will be included in the Appendices.

Chapter 2

SUPPORT STRUCTURE DESIGN

The photonic platform developed in my thesis research comprises a set of support structures, performing functions including mechanical support, thermal dissipation, and interfacing the device mode to external optics. The support structures are versatile in that the same set of structures can be used, while only the sections of photonic structure for atom interaction are swapped for multiple experiments. This chapter will cover the design principles, resulting structure geometries, and their simulated optical performance.

For reference regarding electromagnetic waveguides, see [26] and [27]. The mode projection and adiabatic tapering formalism can be found in [28]. The Knowledge Base website from the *Lumerical FDTD* program [29] also provides useful information regarding the material presented in this chapter.

2.1 Design Targets for Atom-Nanophotonics Hybrid System

There are several design goals for the photonics component of a system that need to be achieved in order to accommodate elements traditionally associated with cold atom systems as follow:

- Optical access to allow integration with free-space optics
- Power handling capability for atom trapping near nanostructure
- Ultra-high vacuum compatibility

Chip layout to allow optical access

The technology available to manipulate cold atom systems often utilizes multiple free-space laser beams, typically with spot sizes in the hundred- μm to few millimeter range. Allowing optical access requires the laser beams propagate through the vicinity of the photonic devices with minimal obstruction. The ideal case for optical access would be for the nanophotonic device to be suspended in open space, with good clearance from any objects that could interfere with these beams. This motivates our fully-suspended structure design, with substrate near and under the photonic device entirely etched away. The requirement that the free-space laser

beams are to bypass our chip with minimal clipping also sets the size scale of the through-window, hence a lower limit for our device footprint.

Power handling

The aforementioned clearance requirement resulted in the connection of devices to Silicon substrate pushed far away from the optically active part of the devices. The resulting structures have small ($\lesssim \lambda^2$) cross-section area, but millimeters in length scale. The microscopic cross-section area creates a very high optical power density that is partially absorbed by the Silicon Nitride waveguide. With the system suspended in vacuum, the primary thermal dissipation channel is thermal conduction that eventually dissipates the heat into the Silicon substrate, and then into the vacuum chamber wall. (Note that microscopic objects emit modified black-body radiation; see [30].) With the small waveguide cross-section area, thermal conduction structures become necessary to run alongside the photonic waveguide device to conduct away heat generated by absorption of light by the device. The maximum power we are able to put into the device guided modes without damage eventually determines the available optical atoms trapping schemes in our experiment.

Vacuum compatibility

The high vacuum requirement for atomic physics experiments places extra constrain on the photonic system in question. The materials used to build the photonics chips, the processing of the chips, and the assembly that eventually contains the chip inside the vacuum chamber must all be compatible to the required vacuum standard ($10^{-9} \sim 10^{-11}$ Torr). The requirement is two-fold. Firstly, the achievable vacuum level is determined by a dynamic balance between the rate of introduction of gas by leaking and material out-gassing, and the rate of removal of gas by pumping. The materials associated with the chips placed into a vacuum chamber must be sufficiently low out-gassing to allow the vacuum to be sufficient. Secondly, the photonics and support structure must tolerate other handling process used in vacuum systems. For example, baking of a vacuum system is a typical procedure to improve vacuum level. As the system is heated to the baking temperature, differential thermal expansion of the holder materials would become an issue. The geometric constrains of vacuum chambers should also be considered. The placement and alignment mechanism of the chip must work in the usually confined space inside a vacuum chamber.

2.2 High-Stress Silicon Nitride and Mechanical Stress Handling

The material of choice for our experiment is stoichiometric Silicon Nitride (Si_3N_4) on Silicon substrate, chosen for its low optical absorption for our wavelengths [31], and also its good physical and chemical resistance. This material retains a very high ($\sim 1\text{GPa}$) intrinsic tensile stress upon deposition on the Silicon substrate. The high tensile stress provides us with the possibility to make long, suspended devices spanning millimeter distances. Upon being released from the Silicon substrate, the Nitride layer will deform under this residual stress until force balance is reached for all points in the device layer. Care must be taken in designing these devices to avoid issues associated with this stress.




Figure 2.1: Suspension of large, hundred- μm scale structures

Large, hundred- μm scale structures can be suspended by sub- μm scale Nitride beams, millimeters away from attachment point to substrate. This figure depicts a defect-mode photonic crystal waveguide device, embedded in a large mechanical support frame. The structure is suspended only by the waveguides (object A and B in this figure), connecting eventually to the Silicon substrate (far left and right beyond the edges of this image).

Avoiding stress concentration

Silicon Nitride is a mechanically resilient material with reported tensile yield strength in the $\sim O(\text{GPa})$ range[32], but its high intrinsic tensile stress ($\approx 1\text{GPa}$ [33]) can cause various mechanical issues. Improper geometry design can create high-stress points in the structure, where the local stress in the Silicon Nitride

device layer after being released from the Silicon substrate can be significantly higher than the initial tensile stress. Sharp corners should be avoided, as they create significant concentrated stress and can initiate a propagating crack. In case sharp corner geometries are required due to desired optical properties, suitable support structures such as tethers should be added to redistribute the mechanical stress, or make the corner into a compressive corner. The other common situation to cause stress concentration is when a thick section of Silicon Nitride is connected to a thin one. Like a heavy spring connected to a thinner spring under load, the larger structure will release its stress and shrink in size, redistributing all strain onto the thin section and create high stress. This should also be avoided in design. See Figures 2.2 through 2.3.

Stress balancing to prevent deformation

Fully releasing the initial ~ 1 GPa tensile stress of the material produces a 0.5% strain change. This influence becomes amplified by the large aspect ratio of our devices to create large deformation. The waveguides must terminate into free space to couple to external optics, creating a zero-stress boundary that is free to move. For our typical 2mm devices and without stress balancing, a $\sim 5\mu\text{m}$ pull-in deformation is observed at the each of the waveguide terminations. Additionally, stress imbalance could lead to significant deformation of the released structure as compared to the lithographically defined shapes on substrate. This can create a systematic error for parameters such as the size of the vacuum gap in a double-beam structure; see Figure 2.4.

We try to tackle this by arranging the device in a geometric configuration that balances tension force at all junctions, with designs such as in Figure 2.5, whenever possible. In case the deformation is inevitable, their influence must be taken into consideration and compensation prior to lithography steps.

2.3 Thermal Conductance Structures

In order to provide thermal conduction for our structures, we make use of the anchoring structure as described in [34]. The anchoring junction comprises a widened waveguide supported by a set of thin periodic tethers. The tether array with small (220nm) lattice constant behaves as an effective uniform but lower index material, allowing the guided mode to be index-guided in the widened center waveguide and preventing scattering into the tethers. With each device, we build two large ($5 \sim 10\mu\text{m}$) suspended Nitride rails alongside the main waveguide to serve as




Figure 2.2: Geometries that create local heightened stress

This plot shows simulation of mechanical equilibrium stress distribution of released structures. The color plot is the resulting Von Mises stress of the material. (A) Connecting a larger (left half) piece of stressed material to a smaller one (right half) results in the strain of the system to accumulate on the smaller one, hence resulting in heightened stress compared to the larger one. (B) Sharp corners, where the center horizontal beam connects to the left side of the plot, produces very large local stress. Making a fillet of finite radius alleviates this issue, as shown in the right half of the plot.




Figure 2.3: Tensile versus compressive corners

This figure depicts a square hole with sharp corners embedded in a material stretched horizontally. The thin red arrows indicate deformation vector, and the color plot shows the Von Mises stress of the material. The up and down corners are 'tensile corners' (red solid arrows), meaning they are being pulled apart and hence produce high stress concentration in their vicinity. The left and right corners are 'compressive corners' (green solid arrows). We can observe the relatively lower stress associated with the compressive corners on this plot, compared to the tensile ones.




Figure 2.4: Mechanical design to handle deformation

As an example for the deformation caused by stress, consider building a structure to split a single-beam waveguide into double-beam. (A) A naive design of direct tapering would deform under sufficient tensile force on the single beam, due to the finite rigidity of the material. (B) Properly shaped splitter design creates significantly less deformation than the previous design under the same tensile force. The deformation is magnified by a factor of 5 on this plot for visibility. The color plot indicates total displacement.

heat-sinks. The tether arrays conduct heat from the main waveguide onto the rails, then further dissipate this heat into the Silicon substrate via the larger cross-section area and hence thermal conductivity of the rails.

The efficiency of thermal dissipation can be characterized by the effective thermal resistance $Z_T = \frac{L}{A\sigma_{SiN}}$, where $\sigma_{SiN} = 30W \cdot m/^\circ C$ is the thermal conductance for Silicon Nitride. For simplicity of estimation, we will assume the heat from the optical mode is being deposited at the center of the device. A typical 2mm length suspended device would have a thermal resistance of $Z_T = 0.5 \cdot \frac{1mm}{200nm \cdot 500nm \cdot \sigma_{SiN}} = 170^\circ C/\mu W$ absorbed power, compared to the devices with heat sink tethers and rails where $Z_T = 0.25 \cdot (Z_{tethers} + Z_{tail}) = 0.25 \cdot (\frac{75\mu m}{200nm \cdot 15 \cdot 100nm \cdot \sigma_{SiN}} + \frac{1mm}{200nm \cdot 5\mu m \cdot \sigma_{SiN}}) = 10^\circ C/\mu W$. We note the factors of 0.5 and 0.25 are associated with that the device is two-sided in its propagation direction, and each side has two rails. The factor of 15 in $Z_{tethers}$ stands for the typical number of tethers used per array.

2.4 Connecting Waveguides to Outside World

We have explored two methods for coupling light into our devices. We initially adapted the fiber-based adiabatic tapering technique in [34] to our wavelength and the optical fiber mode profile. With this tapering technique, the interfacing mode




Figure 2.5: Angled tether array for force balancing

The terminated tip of the coupling waveguide of a typical device. The main waveguide to the left of the image carries strong tensile stress. A set of tether array anchors the tensile force of the waveguide into the substrate, preventing pull-in of the thin waveguide termination to prevent damage. The tether tilt angle θ satisfies $2n_{tether} \cdot w_{tether} \sin(\theta) = w_{waveguide}$ to cancel the total force on the junction on the left.

profile at the end of the device has both small (NA=0.11) numerical aperture and low ($n_{eff} - 1 = 0.003$) effective mode index. The small divergence angle and the low residual reflection at the waveguide termination further allow us to develop a direct coupling technique using aspheric lenses.

FEM calculation and mode projection

Various methods [34, 35, 36] have been developed to match the guided mode profile of a photonic waveguide with sub- μm thickness to the mode of a conventional optical fiber, which is typically few μm in mode-field diameter. For our case, we taper the waveguide to make its guided mode highly evanescent in order to match the larger fiber mode profile. A Finite-Element Method (FEM) calculation using *COMSOL Multiphysics* was carried out to calculate the guided mode profile for given waveguide width and optical frequency. We then select the waveguide geometry that produces the highest mode overlap to the fiber mode. For our specific application, we match the mode of a rectangular waveguide of 140nm by 200nm cross-section to a $5.5\mu\text{m } e^{-2}$ diameter Gaussian in order to couple to the Nufern 780HP fiber.




Figure 2.6: Waveguide and cooling rail

The suspended waveguide (lower horizontal structure) is connected to the cooling rail (upper horizontal structure) by a tether array.

The calculation result using a mode-projection calculation as described in [28] is plotted in Figure 2.7. The two polarization modes of the waveguide generally have different efficiency for a given waveguide width, unless we operate at 200nm-square geometry.

Direct time-domain simulation

We make use of Finite-Difference Time-Domain (FDTD) methods to directly verify the coupling efficiency at interfaces between various optical components on our waveguide, and also between the waveguide and an external coupling channel. FDTD computational time scales unfavorably with total simulation length in the propagation direction, as the enlarged simulation volume compounds with the longer run-time required for the injected excitation to propagate through the length of the simulation volume. The simulation is typically done component-by-component as opposed to the entire device, due to computational resource limitations.

Here we use the commercial FDTD package *Lumerical FDTD* to simulate launching a wave packet of $\sim O(30\text{THz})$ bandwidth from either a cleaved fiber-end or a Gaussian beam generated by the program into a highly-evanescent nanobeam waveguide. The program then Fourier transforms the transmitted field at the other end of the simulation volume to obtain the transmitted power of each frequency. We then

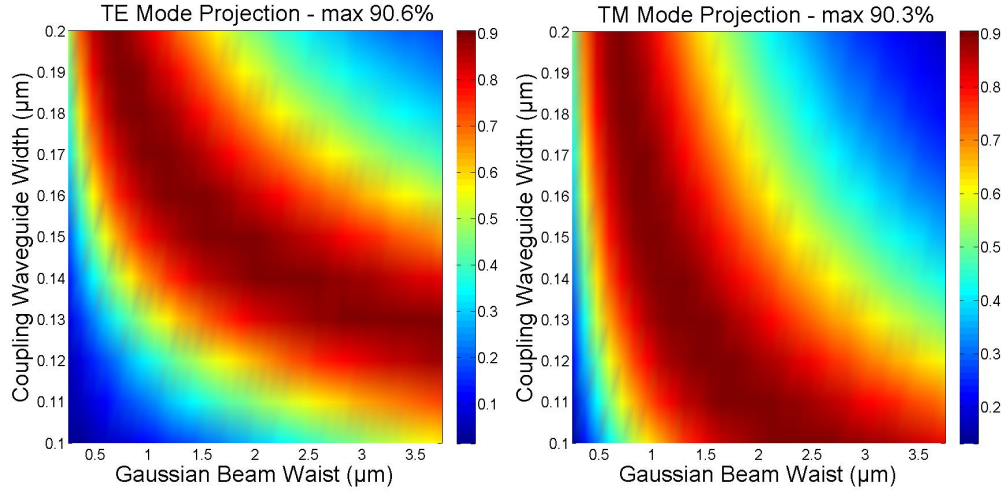


Figure 2.7: Coupling waveguide mode projection into Gaussian beam profile

The guided modes of the coupling waveguide for our devices are highly evanescent and can be matched to a Gaussian mode profile with reasonable efficiency [34]. The width of the waveguide determines the optimum beam waist for the Gaussian beam.

project this power into the guided mode profile of interest using the method described in previous subsection, 2.4. See Figure 2.8 for a simulated mode profile.

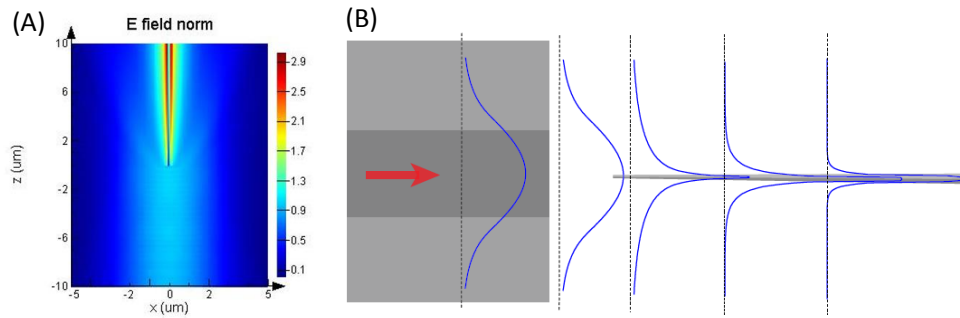


Figure 2.8: FDTD simulation plot of free-standing coupling waveguide

(A) A Gaussian beam source is launched upward from the bottom of this plot, which then couples into the guided mode of the Silicon Nitride waveguide with optimally matched field profile. (B) Once the initial fiber mode is guided on the coupling waveguide, it is gradually transformed into a more localized waveguide mode by an adiabatic taper.

2.5 Anti-Reflection Designs on Highly Evanescent Waveguides

To provide structural support and thermal dissipation channel for the waveguide, tethering structures were added to connect the waveguide to the substrate. These

structural components need to be designed to suppress residual reflection and loss. This is particularly difficult at the coupling waveguide section, where the guided mode becomes highly evanescent to match the fiber mode, as shown in Figure 2.8. The formerly used 'tether array on wide waveguide' method as described in section 2.3 can no longer be applied. A pair of 100nm tethers attached to the end of the waveguide would produce approximately 1.0% reflection for the TE guided mode; see Figure 2.9. In the coupling waveguide section, we are also restricted by the minimum tether size of $\sim 80\text{nm}$ before fabrication yield becomes an issue. In this section, we develop anti-reflection designs specifically for this evanescent waveguide.




Figure 2.9: Reflection from coupling waveguide support tethers

This figure depicts the input of a Gaussian beam onto the coupling waveguide, similar to 2.8, but with a set of tethers for mechanically support. The tethers create reflection, as visible in the interference fringes in the lower half of the plot. The reflection from the tether junction would then back-reflect into the input source and interfere with the reflected signal from the device, creating difficulty in signal analysis.

'Round-about' anti-reflection tether junction

One method we have experimented with to suppress reflection of the tether pair was the idea of dipolar cancellation. We begin with the assumption that an infinite

and uniform waveguide does not create scattering. It is the polarization of the additional material, namely the tethers, that creates a dipole moment and generates the reflection and scattering at the junction. The intuition for this design is that removing material from where the uniform waveguide would originally occupy is effectively the same as adding material of negative polarizability to the total scattering material, in hope to cancel the total scattering dipole moment. Therefore we hope that by opening a hole in the path of the original unpatterned waveguide, the dipole moment associated with the tethers can be partially cancel.

The geometry design is based on fillets on the tether-waveguide junction corners to provide mechanical support, then opening a circular hole on the path of the waveguide, see Figure 2.10. This has the advantage in mechanical consideration as well, eliminating the sharp corners formed by the tether-waveguide junction with fillet structure. We calculated using FDTD method that adding a 300nm radius fillet reduces the transmission of the junction to 72%, but further opening a hole of 270nm diameter restores it to 83%. Note that the numbers quoted here are from the waveguide guided mode before and after passing the tether junction. With both modifications, the reflection can be suppressed from 1.0% to 0.1%.




Figure 2.10: Concept of hole-based dipole cancellation tether

The removal of material from the originally unpatterend waveguide effectively creates a 'negative-polarizability' material to partially cancel the total scattering dipole moment of the tether junction, hence suppressing reflection and partially suppressing scattering.

Y-shaped free-space coupling tether design

Devices designed for free-space coupling allow more available tethering geometries, due to the releasing of the half-plane of space that was occupied by the cleaved fiber-end. We experimented with a design where the tethers are tilted from the perpendicular direction to the waveguide by a sharp angle. This design makes use of the property of sub-wavelength objects with sharp aspect ratio, that they tend to have anisotropic polarizability that favors their long direction as opposed to the thin direction. The suppression effect is two-fold. Firstly, the projection of the waveguide electric field onto the long direction of the tether is reduced, hence reducing the polarization of each tether. Secondly, the two non-parallel tethers partially cancel the total polarization of the system in a vectorial addition manner; see Figure 2.11.

With this concept in mind, numerical optimization using the FDTD program was carried out. We eventually settled on a tilt angle of 60° from perpendicular. The FDTD simulations suggest a Gaussian-beam-to-waveguide transmission of 87% for a $5\mu\text{m}$ beam spot and our frequency range of interest (Cs D1 and D2, 335.1THz to 351.7THz). This method produces an overall simpler design and better efficiency, at the cost of no longer being able to directly couple to the cleaved end of an optical fiber. A summary of optical performance of these devices can be found in Table 2.1; also see Figure 2.12.

Type	Transmission		Reflection	
	TE	TM	TE	TM
Gaussian to 140nm waveguide	90%	66%	<0.1%	<0.1%
with 100nm tether	79%	65%	2.7%	0.6%
with Y-shape tether	87%	56%	<0.1%	<0.1%
Guided mode through 100nm tether	86%	98%	1.2%	0.5%
300nm fillet 270nm hole	83%	97%	0.2%	0.1%

Table 2.1: Summary of optical performance of coupling waveguide tethering designs

2.6 Assembly of Fiber-Coupled Chip

In the fiber-based optics-to-chip coupling scheme, optical fibers are directly attached to the photonics chip. The core of the optical fiber must then be precisely aligned to the terminated coupling waveguide. The transverse alignment of the fiber core to the fabricated waveguide is enforced by the physical contact of the fiber and the V-shaped grooves etched into the chip substrate. The distance of the fiber end surface to the waveguide end is controlled using high-precision micrometers, with

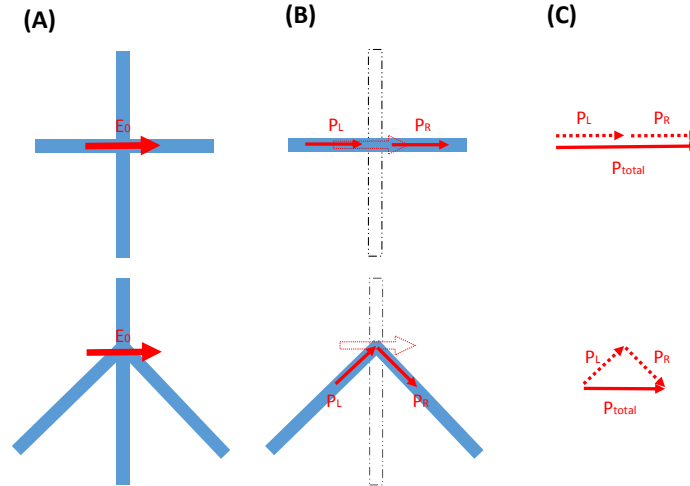


Figure 2.11: Dipole cancellation by tilted tethers

- (A) We conceptually treat the tether junction as addition of dielectric as a perturbation being polarized by the guided mode electric field. (B) Due to the shape-dependent anisotropy of polarizability of deeply sub-wavelength objects, the tethers primarily polarize along their long direction, with the angled tethers effectively only polarized by the parallel projection of the polarizing electric field. (C) The total polarized dipole moment of the tethers is further reduced due to vectorial partial cancellation between the dipole moment of each tether.

real-time feedback of the optical signal from the fiber to be attached. Once aligned, we affix the fiber to the chip using glue. The assembly process can take several hours, and should be carried out in a properly clean environment to prevent chip contamination.

Notes on using glue

We have experimented with the UV-cure glue as in [34] and also thermal-cured glue for preferable higher-temperature performance and out-gassing properties; see Table 2.2. A small droplet of glue is applied to the back end of the mounted fiber using a suitable fine tool, in our case, a short length of stripped optical fiber. Upon application, the glue seeps into the space between optical fiber and the Silicon groove surfaces due to surface tension. It is critical that the glue must not be allowed to overflow beyond the end of the cleaved fiber end. Should the coupling waveguide become encapsulated by glue, it will no longer provide the correct mode profile, and could potentially suffer mechanical damage from the surface tension of the glue.




Figure 2.12: Angled tether design field plot

This is a FDTD field plot of an angled tether junction, where the optimal transmission tether angle from orthogonal was found to be $\frac{\pi}{3}$. We see that the interference fringe as seen in Figure 2.9 is strongly suppressed.

The UV-cure glue can be cured in sufficiently short time (flowing stops within <1 s exposure, full cure typically <1 minute) upon being irradiated by a UV lamp or UV LED. The thermal-cure glue we used typically takes up to 10 minutes to fully cure at elevated ($\sim 100^\circ\text{C}$) temperature, during which time the glue would flow until it reaches equilibrium. The total volume of glue drop determines whether the thermal glue would overflow the device, and the slower cure time allows us to micro-adjust the mounted fiber position during curing to eliminate any motion associated with glue surface tension. Care must also be taken with the differential thermal expansion between chip and holder, which could cause cracking due to thermal cycling. For types and properties of the glues used in our project, see Table 2.2. A step-by-step instruction for gluing can be found in Appendix D.

Fiber feedthrough and strain handling

We found that the adhesion strength between the glue and the Silicon substrate is relatively low, compared to adhesion of the same glue to aluminum or Macor. This is likely a combined result of the smoothness of the Silicon and Nitride surfaces, and

the limited amount of glue we could safely apply without damaging the waveguide. The low adhesion strength makes the fibers susceptible to tearing off the chip under mechanical tension, particularly, if the tension force is perpendicular to the chip surface. We did not notice difference in adhesion strength for the ALD coated chips versus Nitride only chips. We partially alleviate the adhesion issue by putting additional gluing point for the fiber further away from the chip and directly onto the chip holder. At this position, more glue volume can be used without concern of flooding the devices. The adhesion to the material of the chip holder (aluminum or Macor) is also significantly stronger than to the Silicon chip.

We then run the fibers through the vacuum chamber by using Teflon feedthrough as described in [37]. This feedthrough consists of a bulk piece of Teflon with small through-holes drilled in to feed the fibers through. The feedthrough and fibers are then compressed using a *Swagelok* tube fitting, relying on the mechanical compliance of Teflon, which is sufficiently soft to be able to conform to the fiber jacket and the interior of the Swagelok piece, forming a vacuum seal. Multiple fibers can be put through one feedthrough; we have experimented up to four fibers per feedthrough. The highest number of devices we could connect this way for our vacuum chamber is 8 devices with transmission and reflection channels (16 fibers total). We try to minimize the residual length of fiber inside the vacuum chamber, as their acrylate jacket remains a potential source of outgassing. Note that the feedthroughs should be tested with a spare fiber prior to use, to confirm that the holes are through. Forcing the device-coupled fiber could result in breakage. Also, the tightening of the Swagelok to form the seal will cause it to attach tightly to the fibers and is irreversible.

2.7 Coupling into Free-Space Lens-Coupled Chip

Migrating toward free-space direct coupling could provide several advantages for our system, but also comes with its share of technical issues. Elimination of glue and fiber would simplify the vacuum system, and allow for higher temperature baking.

Glue	Curing Flow Stop	Full Cure	Outgassing (Total Material Loss)
Torr-Seal		1hr, 90°C	0.63%
Epo-Tek 353ND	5min at 120°C	1hr, 150°C	0.22%
Dymax OP-67	<1s	1min, UV	3.33%

Table 2.2: Outgassing properties of thermal and UV-cure glue




Figure 2.13: Fiber placement and automated glue application setup

An optical fiber is placed onto a five-axis stage actuated by high-precision differential screw micrometers (A), allowing fibers to be placed safely into the V-shaped grooves etched into the chip substrate. A set of automated stages (B) can then apply small droplets of glue in a controlled manner to attach the fibers. The process is monitored using a long working distance microscope (C).

Removal of mounted fibers on the chips can also reduce device footprint and increase the number of usable devices per chip. Issues including optics working distance, vibration, and lens and window aberration need to be taken into consideration in designing a free-space based system.

Chip configuration adaption to allow free-space coupling

The typical coupling schemes to nanophotonic devices in a non-contact manner are grating coupler [36] and butt-coupling [34, 35]. We opted for the latter, due to our device layer configuration. The mirror symmetry plane in the device layer suggests that out-coming light from a grating coupler would emit equally upward and downward from the device plane, imposing a maximum 50% one-way coupling efficiency limit. Conveniently, since our devices are initially designed to match the mode of conventional optical fiber, the alignment requirements and hence the techniques to use are similar also to conventional fiber-coupling. The challenge for the chip configuration would be to allow sufficient clearance for the free-space beam, whose divergence is associated with the coupling waveguide numerical aperture. We eliminate the obstruction of optical access to the coupling waveguide by terminating




Figure 2.14: Chip holder assembly with fibers

This figure depicts the full chip-holder-fiber assembly ready for vacuum chamber. The chip is mounted on a Macor holder (A), then screwed onto an aluminum arm (B). Then, fibers are glued onto the selected devices, as shown in this figure that a total of three devices are connected fully with transmission and reflection, and two devices are connected only with one fiber, hence only allowing the reflection channel from one side. The fibers are reinforced by putting a second gluing spot directly onto the holder (C), and then attached directly to the aluminum arm (D) to allow easy handling and installation into the vacuum chamber.

the waveguide at the edge of the chip. A guided precision cleaving technique was developed to achieve this, as will be covered in the following chapter.

Vacuum and optics system consideration

The challenge on the vacuum system and optics side is then to provide the correct beam spot and alignment from outside the vacuum chamber. It is also desirable to have long working distance, to allow freedom of positioning the chip, and the capability to manipulate the optical access by placing fixed mirrors in the vacuum chamber. With the numerical aperture of our devices ($NA=0.11$), we found that a working distance of up to 40mm can be achieved with commercial one-inch aspheric lenses; see Table 2.3. Sending a non-collimated beam through the window of a vacuum chamber would lead to aberration of the beam, but correction optics for glass thickness is well understood and commercially available. Thickness-compensated multi-element objectives can also be used [38]. With these free-space devices,

the mode profile of the device mode can be directly imaged for the free-space devices, and advanced free-space optics such as deformable mirrors or Spatial-Light Modulators (SLM) can be applied to map the device mode into a desired free-space or fiber mode with high efficiency.




Figure 2.15: Free-space direct-coupling test setup

The free-space coupling test setup comprises of (A) a five-axis stage with goniometer and rotation axis centered at the chip position (B). The translational stages are actuated with (C) both micrometer screws and piezoelectric stacks for fine adjustment. Then, two high-angular-precision mirror pairs (D) couple light into the device's either ends using the same technique as typically used in conventional fiber coupling. We found that once the the mirror pairs are aligned to the first device on the chip, the translation degrees of freedom on the five-axis stage are generally sufficient to translate from one device to the next while maintaining sufficient coupling efficiency for successive devices.

Lens	Diameter	EFL	WD	NA
Device				0.11
Design target	25mm	40mm	35mm	0.34
Thorlabs A397-B	7.2mm	11mm	9.64mm	0.30
Thorlabs AL2550H-B	25mm	50mm	46mm	0.20

Table 2.3: Lenses used for testing current generation free-space chips

Chapter 3

FABRICATION OF DEVICES

This chapter describes the fabrication process for our project. The process includes electron-beam lithography to define the nanophotonics and their support structures, plasma etching for pattern transfer, the etching of the substrate to release the devices, and post-release processing. The majority of these techniques are standard in semiconductor processing, so we were able to integrate them into our project with varying degrees of adaption. With the fully developed process, we can typically process 6 chips per fabrication cycle with 16 devices each, at a 90% yield rate.

For references on plasma etching, please see the thesis [39]. The *Oxford Instruments* website also provides relevant information on their website page [40]. Information regarding the critical point drying can be found in [41] or other thermodynamics textbooks. The manufacturer's website of our Atomic Layer Deposition (ALD) machine provides some fundamental information on their website page [42]. A step-by-step description of our fabrication recipe can be found in Appendix D. Another useful resource would be the article [43], which includes comprehensive tables of etch rates of various etchants and materials.

3.1 Fabrication Requirements

The fabrication of Silicon Nitride structures is reasonably well understood among the semiconductor and nanophotonics communities. Some adaptations are made to accommodate the properties of our devices, such as the aspect ratio and the device footprint due to fiber and optical beam clearance. We tailor our fabrication process to achieve the following targets:

- Precision in fabrication to allow alignment of device frequency responses to atomic resonances
- Sufficient throughput and yield rate for millimeter-scale footprint devices
- Control of surface tension, bubbling, and cleanness during wet chemistry processes

Precision of device dimensions

The optical resonances of free-space Cesium atoms of interest for this project consist of two invariant and spectrally narrow frequencies: the D1 (335.12THz) and D2 (351.73THz) manifolds. We then need to align the frequency responses of the devices to that of the atom, while being constrained by finite fabrication precision. This is achieved in two primary directions. Firstly, in the design phase, the device geometry should be designed to have low sensitivity to dimension variation, so that the same fabrication precision in dimensions corresponds to a finer sweep of frequencies between devices. The geometry thus designed also reduces sensitivity to imperfection and fluctuation. Secondly, the actual fabrication process should be optimized for producing faithfully the dimensions that critically determines optical properties, such as the photonic crystal parameters. Design paradigms that are robust against small variation are preferred. For example, coupling between different waveguide modes can be done using adiabatic tapering, but also phase-matching methods. The former is efficient provided the taper is sufficiently long, and hence robust against small variation in taper parameters, while the latter require exact matching between interfering channels.

Device footprint and throughput

The physical footprint of the functional section of our waveguides, typically the photonic crystal section, are on the order of $\sim 100\mu\text{m}$ in length, and up to $30\mu\text{m}$ in width for the larger 2D-based devices. The few-mm length of our devices is necessitated by spanning the through-window for compatibility with free-space optics and laser cooling techniques, such as Magneto-Optical Trap (MOT)[44] or Far-Off-Resonance Trap (FORT)[45]. For fiber-based devices, the width of each device can be no smaller than the fiber diameter ($125\mu\text{m}$ cladding without jacket). Having large footprint devices built to high precision resulted in long lithography exposure time and limited the number of devices available per chip. The device length also makes it susceptible to local errors, such as large point scatterers or lithography machine stage stitching failure (see Figure 3.2). A single local error anywhere on the waveguide length can completely ruin the waveguide device's performance, and in some cases cause mechanical breakage.

Issues associated with wet chemistry processes

Several properties of our chips make them particularly susceptible to damage during the wet chemistry processes. Firstly would be their susceptibility to liquid surface




Figure 3.1: Preliminary exploratory photonic crystals

(A) Early attempt to build double-beam photonic crystal using conventional hole-based design. In order to produce the small 15THz photonic band-gap (4.3% gap-to-midgap ratio), very small holes are required, and become difficult to fabricate consistently. (B) Tapering the hole-based photonic crystals is ultimately limited by the smallest hole size allowed by fabrication. Holes too small become only partially etched. (C) Inward-facing modulation in double-beam design. (D) Early corrugated photonic crystal design with outward-facing modulation. The sharp triangular shapes resulted from poor lithography exposure techniques prior to the application of the 'grid' method.

tension. Device designs with small ($\sim O(200nm)$) mechanical gaps could collapse from liquid surface tension. Also, the influence of some mechanical damage can become amplified. An example to that would be broken support structures like the side rails. A device with one of the two side rails broken would typically maintain its optical properties, but liquid surface tension would cause it to tangle with the waveguide and completely ruin the device. Having the devices pass through liquid-air interface during the wet chemistry processes should be avoided. Secondly, in order to etch the through-window, we simultaneously etch the Silicon substrate from both sides. The chip should be placed in a configuration that avoids accumulation of bubbles on both sides of the chip edges to avoid local masking. Accumulation of bubbles can prevent contact of the etch solution to the Silicon substrate, which




Figure 3.2: Stitching error

Stitching errors between main-fields are infrequent but can lead to reduced device yield. Transverse misalignment with small amplitude, as shown in this figure, typically creates an unacceptable level of local scattering loss and residual reflection.

can occur on a microscopic scale (see [46]) and result in undesired features on the etch surface, or in a macroscopic scale where large patches of the pattern become covered and fail to etch. Finally, in order to allow the devices to tolerate the high power density needed for atom trapping, aggressive cleaning procedures become necessary.

3.2 Substrate Selection and Preparation

Our chips are fabricated from substrates of $200\mu\text{m}$ thickness Silicon. We obtain our substrates polished on both sides and coated with Si_3N_4 using the Low-Pressure Chemical Vapor Deposition (LPCVD) process, which is carried out at high (900°C) temperature. This was preferred over the lower-temperature Plasma-Enhanced Chemical Vapor Deposition (PECVD) process due to the reported (see [47]) lower optical absorption for our wavelength range of interest. The $200\mu\text{m}$ substrate thickness was a compromise between the structural strength of the chip and the feasibility of etching the through-window. We have experimented briefly with the type of Silicon, such as P- or N-doped ($50\Omega\text{-cm}$) or high-resistivity ($>5\text{k}\Omega\text{-cm}$), which seem to produce no difference for our processes. We requested the manufacturer to dice the wafer into 10mm by 10mm square chips along the $\langle 1,1,0 \rangle$

crystalline directions. The importance of precise alignment to the crystal axes will be described in detail in section 3.4. We typically pre-clean the chips with solvent chain and Piranha solution (see Section 3.6) before application of UV or electron beam resists.

Pre-patterning the backside for through-window etch

We pattern a 2mm x 6mm rectangle on the backside of the chips, to allow the through-window etching to progress from both sides. Since the optically critical structures are typically suspended in the center of the through-window with large clearance from the edges, the positioning of the back-window relative to the front side patterns is non-critical, and a precision of few tens of μm would suffice and can easily be achieved with UV lithography.

A UV lithography was carried out to define the back-window. We spin-coat the UV resist *Microchemicals GmbH AZ5214E* on both sides of the chip at 4000rpm and then carry out a 1.5 minutes bake at 100°C to form a $1 \sim 2\mu\text{m}$ film. The front side is first coated this way as a protection layer, and then the chips are flipped and coated on the backside for UV exposure. The UV exposure is carried out using a Chromium-on-glass mask, referencing the edges of the pre-diced chips. We use this resist as a positive tone resist, and hence no post-bake is carried out, and the resist is developed directly with *Dow Microposit MF-CD-26* developer. The developed rectangle pattern is transferred to the Nitride layer by plasma etching, then the remaining resist can be removed using solvents (acetone, isopropyl alcohol, then deionized water).

3.3 Electron Beam Lithography with Large Device Footprint

A pattern can be exposed in lithography using a positive tone resist (exposed area is developed), or alternatively, have its inverse be exposed using a negative tone resist (exposed area remains after developing). The pattern of our device contains both large empty areas such as inside the grooves and window, and large filled areas that are the remaining frame part of the substrate. For a single-exposure fabrication process, large-area exposure will be unavoidable for either positive or negative tone resists. An alternative method is devised to avoid large area exposure.

We expose our chip using a *Raith EBPG 5000+* lithography machine on the *Zeon Chemicals ZEP520A* positive tone resist. The resist is spin-coated onto the substrate at 4000rpm, then hotplate-baked at 180°C for 2 minutes, resulting in a 350nm resist layer. Early brute-force approach attempts exposure of the entire device area at the

same high precision required for the critical photonic sections, and had a throughput of approximately one hour exposure time per device. With proper optimization, we eventually were able to speed up our process to approximately 45 minutes per 16-device chip with sufficient precision.

Coarse-fine double exposure

The precision for exposure and the exposure time are interlinked. The limitation of our lithography machine speed manifests itself as a maximum frequency of the electronics that steers the beam from one pixel to the next:

$$f_{DAC} = \frac{1}{t_{Dwell}} = \frac{I_{Beam}}{(x_{Pixel}^2 \cdot C_{Dose})}$$

where t_{Dwell} is the dwell time of the beam spot on a pixel, I_{Beam} is the beam current being used, and C_{Dose} is the electron dose density required to clear the resist of choice. We make use of an exposure scheme that comprises of a fine exposure with highest precision (pixel size $x_{Pixel}=2.5\text{nm}$ using $300\sim 500\text{pA}$) surrounding the waveguide. Then, we fill in the remaining empty areas with coarse exposure using high ($5\text{nA}\sim 10\text{nA}$) beam current and large pixel size. The machine we use seems to typically have an overlay precision of $\sim O(100\text{nm})$ between two exposures consecutively done on the same chip without alignment markers. We eventually allow the coarse and fine exposure to overlap by 500nm to prevent stitching issues between the two layers; see Figure 3.3. Note that this overlap region is dosed twice, and should be made as uniform around the device as possible, in order to reduce artifacts created by uneven proximity effect on the fine exposure.

Back-scattering electron

Flood-exposing large area on a chip results in significant proximity effect. In our case of filling in the grooves and window, the background exposure is approximately uniform. The primary result of the proximity effect in our system manifests as a constant background dose over the entire pattern, effectively shifting the dose-to-clear to lower value. For comparison, exposing a single free-standing $1\mu\text{m}$ line with our resist requires typically $200\mu\text{C}/\text{cm}^2$, while with the flood-exposure from the coarse layer, the fine layer patterns of similar geometries would clear at dose below $150\mu\text{C}/\text{cm}^2$. We noticed significant roughening of the structures from this background dose, which could have resulted from a combination of fluctuation within the back-scattered background dose, and the reduction of the dose contrast between exposed pattern total dose versus the background dose itself; see panel (D)




Figure 3.3: Illustration of coarse-and-fine exposure with overlap area

The lithography exposure contains two overlapping layers. The fine exposure (blue) traces around the device area at high precision, then the coarse exposure (red) fills in the space, with some overlapping area (purple) to prevent stitching problems. The waveguide and 'tiles' (white) are not exposed.

in Figure 3.1.

We countered this issue by using a slice-grid method as described in [48], where instead of filling any large area, we only expose a $1\mu\text{m}$ grid that covers that area. Please see Figure 3.3 for the tile exposure schematics, and Figure 3.4 for an SEM image. The grids separate the Silicon Nitride layer into free-standing 'tiles' detached from the waveguide structures or the groove or window edges. We find that this method reduces the background electron dose sufficiently and both improved our

device quality and reduced beamwrite time. The existence of the tiles seems to slightly slow down the KOH etch, lengthening the etch time by ~ 10 min for the $30\mu\text{m}$ tile size we used. After the tiles are released from the substrate, they become free-floating and disperse in the processing liquid. We occasionally observe tiles landing back on the released devices, with a rate of approximately one occurrence per 10 chips (160 devices).




Figure 3.4: Slice-grid method

Exposing only $1\text{-}\mu\text{m}$ -wide grid as opposed to filling in the entire area of the groove improves exposure quality and also reduces exposure time.

Main-field stitching

The area where the lithography machine is able to expose by steering the beam is called a 'main-field'. The size scales in proportion to the maximum number of pixels allowed by its electronics times the size of each pixel. At the precision setting we use to define our waveguide, the main-field available to us is $160\mu\text{m}$ by $160\mu\text{m}$. This is sufficient to cover the width of a fiber-based device. However, the total length of our device is longer than the field size. The lithography machine is equipped with laser interferometer stabilized motorized stages that can translate the chip between main-fields to 'stitch' them together. We occasionally observe misalignment occurring at points where the field stitching occurs; see Figure 3.2. This could cause scattering of the guided mode, or mechanical breaking of the waveguide.

Care should be taken during pattern file generation regarding main-field placement.

The program to convert our CAD file into pattern file for our lithography machine (*GenISys BEAMER*) allows for control of the main-field placement. We instruct the program to place main-fields only on a rectangular grid, and then place our devices in a coordinate system commensurate with said grid. This ensures each device is covered by one single column of main-fields, and that every device has consistent main-field placement. In this configuration, the machine exposes adjacent main-fields along a direction parallel to the waveguide, minimizing stage drift between two fields. The typical stage drift values from our lithography machine, as automatically measured during the beamwrite using a set of holder markers, were on the order of 10nm per 10 minutes. Finally, main-field boundaries should not overlap optically critical structures, such as photonic crystals. Strong scattering points can be observed consistently at the same position of the photonic crystal for each device, where the main-field boundary crosses the crystal.

3.4 Etching

Here we describe two etching processes we use to fabricate our device. We use a Fluorine-based plasma etch to transfer the lithography-generated pattern into the Silicon Nitride layer. Once the pattern is transferred to the Nitride, we can inspect the resulting dimensions of the devices using a Scanning Electron Microscope (SEM) with minimal charging effect or stress-induced deformation. We then dissolve away the Silicon substrate using Potassium Hydroxide (KOH) to create suspended devices. This step also produces the through-window. Please see the article [43] for comprehensive tables of etch rates of various etchants and materials.

Plasma etching for pattern transfer

A plasma process creates ionized particles by pumping the process vacuum chamber with strong Radio Frequency (RF, $\sim 10\text{MHz}$) power. Due to higher mobility of electrons in the plasma, electrons are accelerated and absorbed into the pumping RF capacitor plates, creating a positive charge in the remaining plasma. This is referred to as the 'DC bias', which will then accelerate the positively charged ions in the plasma toward the same capacitor plate, where the sample to be etched is placed. This ion acceleration by the DC bias provides a strong stream of ion flow in the direction perpendicular to the chip substrate. Generally, the DC bias (and hence the ion kinetic energy) increases with the RF pump power. Coupling a separate RF source inductively (Inductively Coupled Plasma, ICP) allows for control of density of ionized plasma independent of the DC bias.

We transfer our pattern using a C_4F_8 / SF_6 chemistry reactive-ion etching process. This chemistry creates a competition between a near-isotropic growth chemistry by the C_4F_8 and the combination of physical milling and chemical etching of the Fluorine ion and radicals. Tuning the rate of these two processes allows cancellation of the sideways etching rate of the combined plasma. The downward etching rate remains high due to the directionality of the physical milling created by the DC bias at the cost of higher resist layer erosion due to physical milling. We optimize the parameters of the etching process by preparing etch-test patterns, and then cleave the chip in an orthogonal direction to the patterns to obtain a cross-section in an SEM; see Figure 3.5. We include patterns of thin tether arrays, double-beams, and holes to simulate the various parts of our devices. Note that in narrow structures, such as the space between tethers in Panel (B) and in the holes in Panel (C) of Figure 3.5, the plasma etch rate is lower and less uniform than cases where the geometries are more open, like the single-beam in Panel (A) and the space outside the tether array in Panel (B). The difficulty of etching through narrow slots or small holes using the plasma process places a limit on the smallest closed geometries our process can deliver.

Directional wet etching of Silicon substrate

We use a Potassium Hydroxide (KOH) etch to dissolve away exposed areas of the Silicon substrate, with the Silicon Nitride layers serving as etch masks. Our etch recipe of 30% KOH solution in water at 73°C provides us with a $\langle 1,0,0 \rangle$ etch rate of approximately $1\mu\text{m}$ per minute. The etching of Silicon using KOH generates Hydrogen bubbles that need to be taken into consideration for the holder design. It also has a small but finite etch rate on SiO_2 .

The etching of crystalline Silicon is highly anisotropic. In particular, the etching is significantly slower on the $\langle 1,1,1 \rangle$ crystal plane. We make use of the smoothness and precise angle of this plane to position optical fibers as described in [48]. The crystalline directions of our substrate are specified such that the wafer surfaces are in the $\langle 1,0,0 \rangle$ direction, and the wafer is then diced into square pieces in the two orthogonal $\langle 1,1,0 \rangle$ directions. In this configuration, any patterns etched into the Silicon Nitride layer will expose fast-etching directions of the Silicon substrate, and the substrate will etch until only the $\langle 1,1,1 \rangle$ planes remain. Practically, this means any exposed shapes in the Nitride layer will eventually create an 'inverse pyramid' shape that forms the smallest rectangle that contains the exposed shape. The angle between the chip normal and the $\langle 1,1,1 \rangle$ direction of the substrate provides a way




Figure 3.5: Exposure and etch test patterns

Chips are plasma-etched, and then cleaved to expose the side-wall profile of the etch. This allows us to inspect the effectiveness of an etch recipe, such as straightness of the side-walls, and confirm sufficient etching depth into the silicon substrate. We prepare test patterns for cleaving including (A) single beams, (B) array of thin beams, and (C) array of holes. We also prepare test patterns to optimize our lithography condition and dose (D).

to convert the in-plane pattern generated in lithography into a 3-dimensional, V-shaped groove whose depth we can control. We note that, however, the $\langle 1,1,1 \rangle$ plane etch rate is small but finite, and the $\sim O(2\mu\text{m})$ etch into the slow direction must be pre-compensated during the patterning stage to achieve correct fiber placement. Please see Figure 3.6 for an illustration of such an anisotropic etch, and Figure. 3.7 for an SEM image of the etched chip.

3.5 Applications of Silicon Anisotropic Etching

Controlling the KOH etching rate of different crystalline directions allows for the fabrication of several 3D structures in the Silicon substrate from 2D patterning in the Nitride layer. Many of these methods require precise temperature control and timing of the etching. The precision of these etched 3D structures also strongly depends on the angular alignment between the lithography-written patterns and the true Silicon crystal axes.

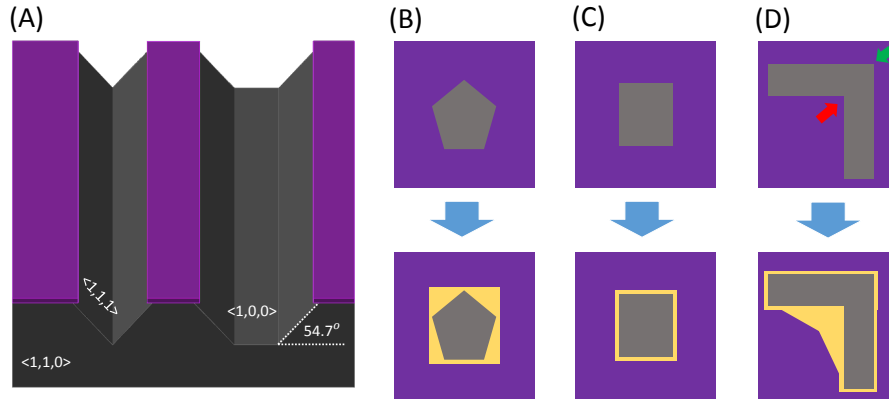


Figure 3.6: Illustration of anisotropic etching of Silicon substrate

(A) Illustration of the lattice planes, and their responding shapes after Potassium Hydroxide etching. The purple layer depicts the masking Silicon Nitride layer. (B) Any sufficiently small exposed geometry on the mask layer (gray) will become contained in a square-shape etched hole (yellow). (C) A square shape mostly maintains its shape and size, enlarging only slightly from the finite etch rate of the $\langle 1,1,1 \rangle$ plane. (D) A 'positive' corner, pointed out by the red arrow, would become attacked by the KOH. A 'negative' corner, pointed out by the green arrow, would be protected.

Etch-in of positive corners

By 'positive corners' we refer to shapes that protrude into un-masked space, such as the corners of a rectangle Silicon Nitride block that is free-standing in exposed Silicon; see Figure 3.6. Positive corners produce exposed $\langle 1,1,0 \rangle$ direction and will become attacked by the KOH etch. This can create undesirable large swathes of suspended Silicon Nitride membrane to cause structural problems, such as stress concentration and cracking described in Chapter 2. This is also the mechanism by which the 'tiles' from our lithography step are removed, since their shapes are that of free-standing squares and their corners are attacked. Alternatively, if a faster tile removal is desired and thus tilting the tile by 45 degrees, making its edge no longer parallel to the $\langle 1,1,1 \rangle$ planes, could potentially be helpful.

'Fuse' timing method for residual Silicon control

The cooling side-rails running along our waveguide serve the double-function of protecting the positive corners formed by the V-grooves terminating into the through-window. By connecting rails across the through-window, the grooves and the spaces between grooves in the through-window area are separated into closed rectangles with no positive corner. However, these rails also masked the etching for the




Figure 3.7: Full view of through-window chip and V-grooves

This figure is an SEM image of a fully etched chip with through-window. The devices, here being platforms to support defect waveguides in 2D crystals, are suspended in the middle of the through-window. The V-grooves and various other features at the edge of the window associated with the anisotropic Potassium Hydroxide etch are also visible. We note that the positive corners near the edge of the window have been protected by the 'fuse' method for this chip, as described in Section 3.5. The background of the through-window is an aluminum holder for the SEM.

Silicon directly under themselves, and created 'bridges' of Silicon spanning across the through-window that interfere with optical access. In order to remove these residual Silicon, we pattern holes on these rails to expose $\langle 1,0,0 \rangle$ planes at regular interval ($1\mu\text{m}$ holes at $3\mu\text{m}$ interval). The holes are only placed starting from $120\mu\text{m}$ away from the edge of the window, so that the initiation of etch under the rails begin at a distance away from the V-groove termination corners, and hence the corners are not attacked by KOH. When the etching is timed correctly, the etch-front generated by the holes would propagate through the $120\mu\text{m}$ 'fuse' length, clearing the through-window of Silicon, but not etching into the corner; see Figure 3.8.




Figure 3.8: Fuse-protected positive angles

The chip depicted in this image shows the etching of silicon substrate being generated by the holes on the safety rail and gradually advance toward the positive corners on the right. To make a through-window chip with protected corners, the positioning of the rail holes is adjusted so that the etch front reaches the corners at the time the through-window clears.

Silicon half-aperture for stray-light protection

In our fiber-in-groove configuration, light launched from the fiber that failed to couple into the waveguide guided mode can propagate in free-space toward the through-window, potentially interfering with the atom experiment. One method we experimented with to partially suppress this stray light was making a Silicon barrier using the anisotropic etch. The idea is to mask the Silicon etch with a wide ($8\mu\text{m}$) Silicon Nitride cross-beam across the v-shaped groove to create a Silicon barrier across the groove, and then use some of the timing control methods in etching to selectively open a small aperture directly under the waveguide to allow its guided mode to propagate through. We found that high ($>97\%$, simulated) transmission can be achieved by using a 1:10 long taper to convert the 500nm nanobeam guided mode to a $6\mu\text{m}$ beam, and then penetrate an unpatterend $8\mu\text{m}$ thick Nitride cross-beam. The resulting Silicon and groove geometry is shown in Figure 3.9. This method can also potentially provide a rigid anchoring point for mechanical considerations.




Figure 3.9: Silicon half-barrier formed by timed etching

A wide beam of Silicon Nitride (A) protects a ledge of Silicon (B) underneath, then, a through-hole (C) is created by placing two holes near the center of this beam. By changing the length of the holes in the direction perpendicular to the beam (parallel to the waveguide), the timing of the hole opening during the etch, and hence the size of the hole in the Silicon, can be controlled.

Etched precision cleave grooves

We develop a precision cleave-guiding groove technique integrated into our chip pattern. This is done to allow direct coupling into our device using external lenses. The cleaving groove allows us to position the end of the waveguide sufficiently close to the edge of the cleaved chip, and to allow the free-space beam path to clear the Silicon substrate. Please see Figure 3.10. A direct patterning of the waveguide near the chip edge is unfavorable due to the uneven resist thickness near the chip edge. The thicker resist at the edge of chips creates variation in dose-to-clear, and can create physical cracks in the resist layer during baking. The combination of the resist cracks transferred into the Nitride layer along with the device pattern, and also that the chip dicing edge direction is $\langle 1,1,0 \rangle$, creates an uneven chip edge during KOH etching that is difficult to control.

We fabricate a set of orthogonal V-grooves that run across the width of the chip, based on the same lithography exposure and KOH etch steps as the main pattern. The sharp lines where the $\langle 1,1,1 \rangle$ surfaces of the cleave grooves meet provide a precise sharp corner in the substrate to initiate a crack under applied torque, which then

propagates down the $\langle 1,1,0 \rangle$ crystal plane to form a clean cleave through the chip substrate. This process is sufficiently well-controlled that the waveguides remain undamaged during cleaving. Since these grooves are lithographically defined by the same step with the waveguide patterns, they are aligned to each other at high precision. The cleaved surfaces can then serve as positioning reference surface for ease of later alignment to holder and optics.

The cleaving of a free-space chip is typically done as follow:

1. Prepare a clean microscope slide and two pairs of carbon-tipped tweezers.
2. Place the chip on the glass slide, then move it with tweezers to the edge of the slide. Make sure the edge aligns well with the cleaving groove on the chip. The two lines should overlap on top of each other, leaving the unpatterend section of the chip suspended.
3. Apply force on the chip substrate away from the device patterns, on both sides of the chip, with a pair of tweezers. This is done to hold the chip firmly on the glass slide. On the clean and flat surface of a glass slide, this should not cause damage to the chip.
4. Use the other pair of tweezers to apply pressure to the suspended chip. It is advised that this to be done at two points on the chip using the two tips of the tweezers. Apply downward force between the two points in a uniform manner. The cleaving should be induced by the torque vector parallel to the cleave groove. If this application of force is not uniform, it could cause the cleave to occur along a device groove, or the corner of the through-window.
5. Remove the cleaved chips for storage or testing. For fragile chips, it is possible to use a small piece of Kepton tape to handle it, to avoid applying force to the substrate with tweezers.

Substrate rotation check patterns

The angle alignment precision between the patterning directions and the substrate crystal plane direction is critical for the anisotropic etching methods described above. We align the lithography pattern axes based on the dicing edges of the chips, and can only confirm the orientation of the true lattice planes after wet etching. We found that a precision of $<0.25^\circ$ angle in rotation is required for the etch-controlling




Figure 3.10: Illustration of chip cleaving by fabricated cleave groove

(A) Optical image of a cleaved chip with through-window. (B) A cleave-chip is fabricated with horizontal grooves that run across the entire chip. Upon applying pressure to torque the substrate in the vertical orientation, the chip cleaves from these horizontal grooves, exposing the ends of the devices to allow free-space coupling.

techniques to work, which can be smaller than the tolerance quoted by our provider ($<1^\circ$). To counter this, we put test patterns on each of our chips to measure the relative rotation between the written pattern and the substrate crystal planes (see Figure 3.11), so that we can compensate for the rotation during lithography step for the remaining chips from the same wafer.

3.6 Wet-Chemistry Processing

This section will cover the process through which the chips go through after the plasma pattern transferring. This will include resist residue cleaning after plasma etching, the operation steps of the Potassium Hydroxide etch of Silicon substrate,




Figure 3.11: Rotation check pattern

(A) A set of rotation test patterns is placed on each chip to monitor potential differential rotation between the written pattern (gray rectangles) and the underlying Silicon crystal planes (bright yellow rectangles). (B) The misalignment angle θ between lithography and the true crystalline axes of the substrate can be measured by the angle between the two rectangles.

followed by Nanostrip and Hydrofluoric acid cleaning, and the critical-point drying.

Chemical processing holder design

A chip holder is built specifically for wet chemistry processing of our chips. We design our holder to satisfy several constraints. Firstly, we chose an upright orientation for chip placement, mainly to prevent bubble accumulation on the back side of the chips, which can cause masking and prevent proper etching. The upright chip placement also allows us to view the chip during the process more easily to determine whether the through-window is completely cleared. Secondly, we make the attempt when possible to make our holder out of a monolithic piece of Teflon, to ensure the holder not become damaged by the aggressive process chemicals. We also minimize potential places that could trap residual processing chemical between steps, such as vertical blind holes or threaded holes. Finally, but most importantly, we make use of a base-and-ring design as used in [48]. While moving the holder from one processing fluid to the next, the holder is surrounded by a ring that contains the fluid around the chips, to prevent premature drying. Once placed in processing fluid, the ring can be removed and replaced in the fluid, allowing the processing fluid to react with the chips.

To form a seal between the holder base and the ring, we found that putting flat glass against a Teflon surface or placing two Teflon surfaces together is sufficient to

prevent leakage for all processing fluids we used. The seal is likely due to Teflon's hydrophobic surface, allowing the process fluid's surface tension to prevent it from leaking. We would advise against making a physically snug fit between Teflon pieces, due to its tendency to deform during thermal cycling that could cause the ring to get stuck on the holder. Should a new Teflon holder with ring be built, it may be helpful to put the pieces together and put them through thermal cycling without actual chips, to allow them to conform to each other and fit properly.

Chip installation onto the holder

Due to the large amount of bubbles generated by the KOH etch, it is necessary to hold the chip firmly in the holder to prevent it from floating out. We note that a free-floating chip will travel around the KOH solution randomly, carried by varying buoyancy from the attached bubbles, and can potentially damage other chips. After the chips are placed into the pockets on the holder, Teflon spacers are wedged between the chip and the holder. The softness of Teflon allows it to deform slightly around the chip, applying force to hold the chip in place. The holder is designed in such a way as to hold the chip from its edges, in order to avoid the Teflon pieces trapping bubbles around it due to its hydrophobic surface property. This side-holding configuration instead serves to wick away the bubble generated by the unmasked sides of the chip substrate. Our current generation of holder has a maximum capacity of six chips. Detailed dimensions of our holder can be found in Appendix C.

Etching of the through-window

We carry out the Potassium Hydroxide etch at 30% concentration at 73°C, while stirred using a magnetic stir rod. The temperature of the mixture is monitored by a thermometer and controlled by an electronic PID loop. With our setup, we observe the variation in our etch time, as defined by the time it takes for the through-window to clear for all chips, to be on the order of 15 minutes in our 2.5 hours process. Near the end of the etch, we observe visually to confirm that the through-window is clear for all chips in the holder before terminating the process. To terminate the KOH etch, we transfer the holder into deionized water that has been heated to equal or higher temperature than the etch solution, as a precaution to prevent potential generation of sediment from cooling and lowering of saturation concentration. The Teflon ring is lifted in the rinse water, allowing the heavier etch solution to drain out and clean water to fill in.

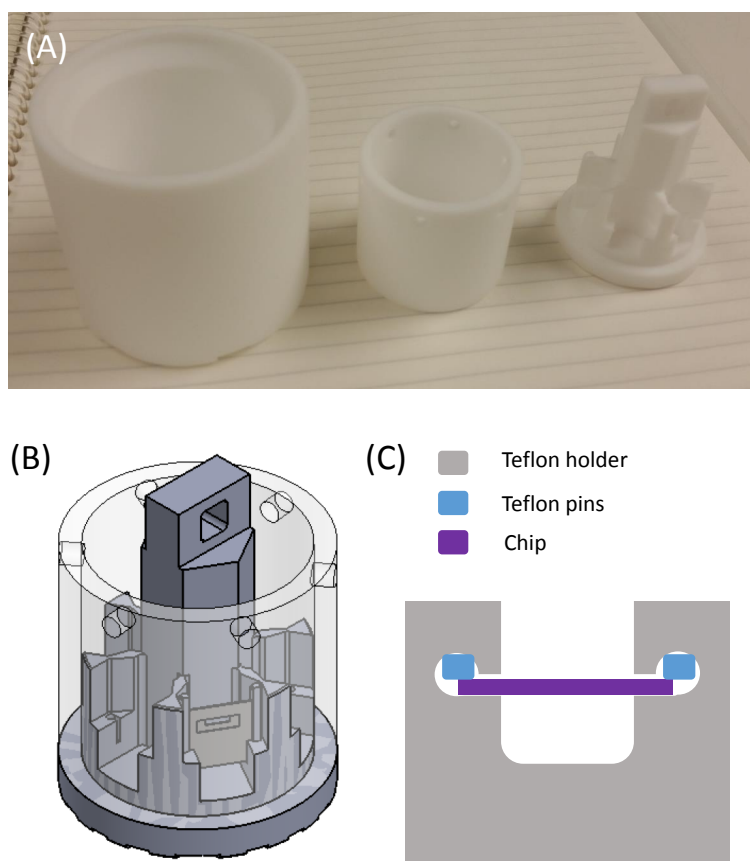


Figure 3.12: Teflon chip holder for wet chemistry processing

(A) Left to right: filler piece for critical point dryer, Teflon ring, Teflon holder. (B) Illustration of assembled holder, Teflon ring, pins, and chip. (C) Top view of chip-holding pocket. Note that there is sufficient clearance on both front and back of the chip to avoid bubble accumulation.

Device chemical cleaning

Our process makes use of two primary sets of chemical cleaning processes. The first set is carried out after the plasma pattern transfer step. We remove the residual resist after plasma etching with the series of solvents: acetone, isopropyl alcohol, and then deionized water. We found that drying chips from isopropyl alcohol sometimes leaves a thin layer of residue that interferes particularly with visual contrast in the SEM. Rinsing the isopropyl alcohol off with deionized water seems to alleviate this issue. We typically carry out a Piranha solution (heated 2:1 mixture of Sulfuric acid and Hydrogen Peroxide) cleaning afterward; the water rinse also prevents potential residue of isopropyl alcohol from violently react with this solution. The Piranha

solution is one of our two acid-oxidant cleaning processes, which we mostly restrict to cleaning only chips prior to KOH etch due to the turbulent nature of the solution. The addition of Hydrogen Peroxide to Sulfuric acid is highly exothermic and creates a large amount of Oxygen bubbles and acidic fume. Care must be taken in using this cleaning recipe.

The second set of cleaning is carried out after the conclusion of the KOH etch. At this point the devices are already released, and the acid-oxidant cleaning uses the Cyantek Nanostrip™ resist stripper instead. Both Piranha solution and Nanostrip are strong oxidizers that react with and remove organic material, with the active ingredient being Hydrogen Peroxide and Peroxymonosulfuric acid. The Potassium Hydroxide solution also weakly etches glass, potentially resulting in Silicate material dissolving into the solution, either from the glass beaker for the etching or from earlier holder designs with glass components. We suspect this Silicate material generates Silicon Dioxide upon contact with strong acid and deposit particles onto our structures; see Figure 3.13. Thus, we further clean our chip after the Nanostrip process with Hydrofluoric acid (commercially pre-mixed Buffered HF 10:1 for better selectivity of SiO_2 etch rate versus Si_3N_4 , then diluted 1:5 with water). This step removes the glass contaminant we observed and also other Oxide species that react with Hydrofluoric acid.

Device drying and surface tension

The conclusion of our wet-chemistry processing involves removing the chips from liquid, where the chips must avoid crossing the liquid-air boundary to avoid damage from surface tension. This is achieved by undergoing Critical Point Drying (CPD). At high temperature and pressure, the phase difference between liquid and gas phase vanishes, allowing us to avoid crossing liquid-gas surfaces. We transfer our chips into high-purity isopropyl alcohol (VWR *OmniSolv*, spectrophotometry-grade), and then place the holder into a critical point dryer. The isopropyl alcohol can be continuously replaced with liquid Carbon Dioxide, since the former dissolves in the latter. We note that geometries that can locally trap residual isopropyl alcohol should be avoided, which was taken into consideration in our holder and CPD chamber spacer designs. The CPD chamber is repeatedly purged with pure liquid CO_2 until the residual isopropyl alcohol is sufficiently low. Since the critical point for CO_2 is only 31.1°C , it is relatively simple to bring the liquid in the sealed high-pressure chamber above critical temperature, then vent back to room pressure at constant temperature. This process brings the chip into air without passing it through the




Figure 3.13: Congregate of solid contaminant removed by Hydrofluoric acid

Large congregate of solid material as depicted in this image was found on our devices after the wet process, particularly with chips held with glass spacer. Neither Nanostrip nor aggressive Oxygen plasma affect these materials, but Hydrofluoric acid seems to eliminate them efficiently.

liquid-air boundary. We found this process deposit small $\sim 5\text{nm}$ fragments on our devices, particularly in the section suspended across the through-window. This type of contaminant is typically removed using a post-drying Oxygen plasma process.

Manipulation of devices using SEM charging effect

Charging effects in Scanning Electron Microscopes (SEM) is a well-studied phenomenon. The SEM operates by sweeping a focused electron beam to image nano-scale structures. The irradiation of such energized electrons can cause the device to accumulate electric charge. Typically, this is associated with image artifacts, but for our suspended and highly electrically insulating structures, we found charging to result in significant but transient physical deformation.

At sufficiently high acceleration voltage ($>10\text{kV}$), we begin to observe large deformation of our structures, which can be useful in two ways. Firstly, it becomes possible when necessary to dislodge moderate-sized ($\sim 1\mu\text{m}$) dust particles or Nitride fragments from the waveguide by placing the SEM field of view directly onto the particle for a few seconds. Secondly, it becomes possible to repair double-beam structures that have stuck together due to temporary forces such as surface tension.

Please see Figure 3.14 for a demonstration of this charging-induced motion. We have experimented with bypassing the CPD, instead directly drying the chips in air. We then un-stick the double-beams by this charging effect. We found this to be applicable to our 1D waveguides. Larger, more rigid structures, such as the frame-supported defect waveguide as described in Chapter 7, appear no longer able to be moved by charging.




Figure 3.14: Large, multi- μm scale motion of double beam devices due to SEM charging

This image depicts the mechanical separation of an early hole-based double-beam device in the SEM, taken at 15kV acceleration voltage. We see that the vacuum gap separated to $d' \approx 3\mu\text{m}$ from an initial value of 250nm. Once the SEM field-of-view is moved away, or a lower voltage is used, the device relaxes back to its original geometry in a timescale of few seconds.

3.7 Post-Release Plasma Cleaning

We found that Oxygen plasma processing to be applicable to our device after release and critical point drying. The plasma processes we used were operated at 20mTorr and a DC bias of $\sim 140\text{V}$ in a capacitively coupled plasma etcher (*Plasma-Therm RIE*). We found that a short (4 minutes) process was sufficient to remove the small amount of residual fragments after the CPD process. We found no sign of device degradation for Oxygen processing up to 30 minutes using this recipe. This process was also used as attempts to re-clean chips from organic contamination, such as organic material evaporated from the environment and re-condensed onto the devices

during the application of thermal glue; see Figure 3.15. Plasma processes, since they are dry processes, can be carried out after initial optical test, and also after ALD coating. Suitable plasma etching recipes can be used to tune the devices with feedback from optical measurement; for more information please refer to the thesis [49].




Figure 3.15: Re-condensation of organic material due to environmental contamination

This image depicts one of the cooling tether array on our chip after baking on our thermal-cure glue processing stage. If the stage is not sufficiently clean, material can evaporate from the stage and condense onto our devices. We found that this influence is severe in the sections over the through-window but limited in the V-groove section. The contaminant is believed to have passed through the liquid phase due to its capability of sticking the double-beam and the tether arrays together, and the apparently surface-tension-induced shapes as observed on this image. We believe this material is organic, as it can be removed by Oxygen plasma processing, and the stuck devices then repaired using the SEM charging technique.

3.8 Characterization using photonic crystal cavities

We fabricated a set of Fabry-Perot-like cavities using a single suspended nanobeam cladded with photonic crystal mirrors on either end. This configuration is chosen for its ease of analysis and of finding a conventional cavity QED correspondence. The cavity devices then serve as a testing standard for device quality associated with different processing methods. The quality factors of these cavities provide a direct numeric comparison between results from sets of process parameters, such as the

choice of cleaning chemistry or the length of Oxygen plasma processes. The cavity system can also help us study loss mechanisms of our system. A set of cavities with nominally identical mirrors but different intra-cavity beam length can differentiate between the loss from propagation versus from mirror transmission and scattering. Further comparison of propagation loss as function of thickness of the intra-cavity beam, and hence the surface field strength, provides information regarding surface loss of devices. The difference in surface quality between the side surface created by reactive ion etching versus the top and bottom surfaces made from initial LPCVD process is studied by using different polarization modes.




Figure 3.16: Single-beam photonic crystal mirror

A pair of single-beam photonic crystal mirrors are built to create a Fabry-Perot-like cavity. A longer taper is built on the inside (right side on this image) of the cavity mirror to reduced scattering loss. The reflectivity of this mirror can be changed by varying the number of nominal lattice sites.

3.9 Metal Deposition for Potential Electrostatic Tuning

We explored preliminarily with metal deposition on our chips for applications in on-chip capacitors and wiring. We apply a second, aligned lithography step after the initial lithography, plasma etching, and cleaning steps to define areas to be coated with metal. An alignment electron beam lithography involves using the beam-writing machine temporarily as an SEM to look for pre-defined patterns ('marker') on a chip. The position of these markers, defined in a previous lithography writing step, can then be referenced to place the new layer of pattern precisely on top of

the previous one. While it is typically preferred to deposit the metal layer first and include metal markers, our resist thickness and Silicon-under-Nitride negative-tone marker seem to allow sufficient alignment precision for our exploratory capacitors.

The metal is deposited using electron beam evaporation, where a reservoir of pure metal is struck with a high-energy electron beam in vacuum, locally heating up the metal to create metal vapor emitted from the reservoir and deposited onto the chip. The metal pattern is defined by exposed areas on the chip that are not covered by the resist. Once the metal deposition is complete, the resist is lifted off from the substrate in solvents (room-temperature trichlorethylene) with the metal layer, leaving the metal pattern only on where it was directly deposited onto the chip itself. We experimented with gold wiring, with Chromium as adhesion layer between Nitride and gold (5nm Cr, then 100nm Au). This configuration has been tested to withstand the KOH etch and subsequent Nanostrip cleaning. We have not yet tested the metal-coated chips with buffered Hydrofluoric acid. With metal wiring, we have preliminary capacitive tuning results for the double-beam photonic crystal, which will be covered in Chapter 8. Also, capability to coat metal can provide future possibilities such as metalized thermal conduction wires or local magnetic field wiring for atoms.

3.10 Device Fine-Tuning Using Atomic Layer Deposition

The technique of Atomic Layer Deposition (ALD) is a chemical deposition technique that uses multiple self-limiting reactions. Here 'self-limiting' refers to that in suitable range of reactant pressure and substrate temperature, the material grown onto the substrate prevents further reaction with the depositing reactant. In an ideal condition, the reactant chemical will deposit one single atomic layer on the substrate so no further reaction can occur; then the reactant is pumped away, and a different chemical is used to grow the subsequent layer. The process we use is based on trimethylaluminium (deposits Aluminum) and water (deposits Oxygen) to grow Aluminum Oxide. This self-limiting, discretized process allows for addition of material precisely one atomic layer at a time. We find this to be most useful for precision tuning of frequencies of our photonic structures. The tuning rates of our typical 1D photonic crystal devices can be found in Table 4.1. More information can also be found in the thesis [49] and [50].

3.11 Experiments with High-Temperature Annealing

We have also experimented with some thermal treatments of our chips. For the earlier devices, we relied on a pre-release annealing step to improve device yield. We also briefly studied the influence of thermal processes on the optical properties of released devices.

Thermal treatment to improve yield

During the development of our process, we experimented with heat treatment to our chips in order to improve yield. On early devices prior to the development of stress balancing tether array, the releasing of stress on the waveguide resulted in an approximately 4 degrees tilting of the coupling waveguide tethers. This resulted in frequent breaking of the coupling waveguide tethers, resulting in a device yield rate of approximately 5%. We found that a N₂ atmosphere 900°C 1-hour heat treatment after the plasma etching step improves device yield for these early devices to around 90%.

The thermal treatment seems to induce some change to the Silicon substrate, resulting in enlarged and uneven V-grooves; see Figure 3.17. We suspect that this was a result of thermal shock due to rapid change of temperature as we translated the chips between the hot zone in our furnace (900°C) to the cold zone at room temperature. The furnace (*Lindberg Blue M Tube Furnace*) consists of a quartz tube that goes from a room temperature section for loading (cold zone) into a heated high temperature section for thermal processing (hot zone), with the chips carried on a quartz boat manually moved with a quartz rod. This substrate issue seems to be alleviated by a multiple-step sample transferring process where the boat is parked at intermediate positions between hot and cold zones to gradually thermalize. While the pre-release thermal process seems viable, it is no longer necessary after the stress-balanced design was introduced.

Post-release annealing

We have some limited experience with post-release annealing of our devices. For Nitride-only devices, we have annealed them up to 1,100°C for up to 1 hour. The devices survive the process with approximately 0.3% shrinkage in waveguide total length, as observed in tilting of the coupling waveguide tethers. We did not observe any significant change in optical properties of the devices. For ALD-coated devices with 10nm Al₂O₃, we observed a 40% reduction in device reflection level after a 15 minute annealing at 600°C, and a total loss of reflection signal after an 1 hour




Figure 3.17: Signs of potential thermal damage to silicon substrate

In the initial attempts of thermal annealing of our chips, we found the process to drastically increase groove size fluctuation. Cracks on the $\langle 1,1,1 \rangle$ surfaces, as shown in this image, and irregular shape of where these surface meet the edge of the V-grooves, were believed to be hidden damage to the crystalline structure of the substrate that is eventually exposed by wet etching.

annealing at 800°C . We suspect the potential causes are loss of stoichiometry of the ALD material similar to that observed in [51], or mechanical issues associated with increased stress from annealing of Al_2O_3 as reported in [52]. We have thus far not observed any difference between processing in N_2 or O_2 atmospheres.

Chapter 4

ENGINEERING OPTICAL PROPERTIES WITH APCW

This chapter will cover the design ideas of the quasi-1D 'Alligator' photonic crystal waveguide (APCW). This design is optimized for atom-light interaction with its guided mode, with capability to provide guided mode Far-Off-Resonance Trap (FORT) for atoms. During the development of this photonic crystal structure, we found several other important considerations. A practical device requires proper tapering to couple to unpatterend waveguides for input and output, and the finite length of actual devices imposes limitation on how close we can operate to the true band-edge of the infinite structure. In the later sections of this chapter, we also present some alternative engineering options that make use of different polarization or spatial symmetry, to potentially achieve stronger slow-light effect or create circular polarization.

For references on photonic crystals in general, the book [53] provides an intuitive overview. The formalism for classical Green's function can be found in electromagnetism textbooks such as [26]. For information regarding FORT, which is based on the AC Stark effect, see [54] and also [45]. Information for calculating Casimir-Polder force for atoms can be found in [55] and [56].

4.1 Engineering Atom-Light Coupling Using Photonic Crystal

The means provided by nanophotonics to engineer atom-light coupling are two-fold. Firstly, there is the capability to manipulate the spatial structure of the optical modes, such as mode profile, effective mode area, and polarization. Secondly, the engineering can be done in the frequency domain, where properties such as guided mode group velocity and photonic band-gap can be manipulated.

Photonic crystal structures are interesting in particular. In the spatial consideration, it allows us to break the continuous translational symmetry of a conventional waveguide without causing scattering loss. In the frequency domain, it provides interesting dispersion properties. The periodic modulation of photonic crystal can also provide discrete trapping sites for atoms. The discussion in this section will focus on the quasi-1D photonic crystal. In this system, the structure is periodic only in the propagation direction in consideration. The light is assumed to be con-

finned in the other directions orthogonal to propagation by the refractive index of the waveguide material.

Mode profile of the photonic crystal waveguide

We opted for a double-beam design for our waveguide as the first step in mode profile engineering. The vacuum gap between the two beams is 250nm in width, and the sub-wavelength size allows the guided mode field to span across the gap in an evanescent manner. This configuration creates strong field intensity in the vacuum gap where there is no dielectric material, to allow atoms to be placed. The mode profile is further manipulated in the longitudinal direction by properties of a photonic crystal near the edges of its Brillouin zone in reciprocal space.

For guided mode \vec{k} 's in the interior of the Brillouin zone, the periodic structure behaves like an ordinary waveguide with some finite, slowly varying effective group index. As the \vec{k} vector approaches the Brillouin zone edge, the guided mode profile begins to pick up a coherent backward-propagating component, which becomes a complete standing wave on the edges of the zone. The standing-wave modes on the edge of the Brillouin zone have the same \vec{k} , but can be in-phase or out-of-phase with the dielectric modulation. One mode occupies the higher average index half of the period, and hence have a lower frequency. We refer to this as the 'dielectric' mode. Then, another mode that is out-of-phase with the dielectric mode by half the lattice period occupies the lower average index section, and hence have a higher frequency. We refer to this as the 'air' mode. Along the propagation direction, these two modes each has intensity minimum at the other's intensity maximum position; see Figure 4.1. These modes, with different frequency separated by the band-gap but maintaining the same spatial periodicity, will be utilized for our various schemes for atom trapping and interacting.

Coupling enhancement on band-edges

We can phrase the argument regarding back-propagating mode and standing wave in Section 4.1 in the frequency domain as follow. The equivalent picture to having back-scattering in an intuitive sense is that the forward- and backward-propagating mode of an unpatterend waveguide can interact with each other through the addition of modulation of the photonic crystal. Therefore, an anti-crossing in the dispersion of these two modes is produced, mediated by the lattice momentum associated with the crystal lattice, $k_{Lattice} = \frac{2\pi}{a}$, where a is the lattice period; see Figure 4.2. Since the anti-crossing is formed from a guided mode crossing the backward-propagating

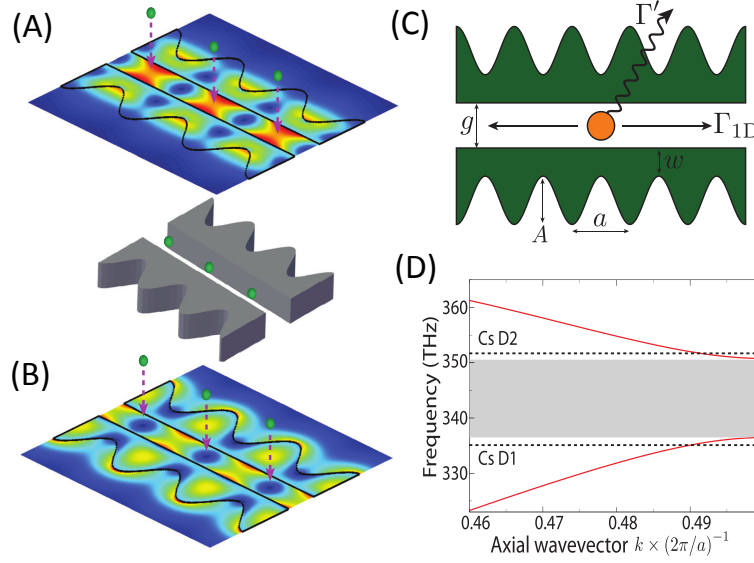


Figure 4.1: Mode profile and band structure of the 'Alligator' photonic crystal waveguide

(A) The 'air' mode and (B) the 'dielectric' mode of the 'Alligator' photonic crystal waveguide (APCW) structure. Note that the spatially alternating minimum and maximum of the two modes. The green spheres indicate placement of atoms, assuming the dielectric mode is used as blue-detuned trap mode, and the air mode as probe. (C) Illustration of the crystal geometry parameters. The optical properties of interest are the decay rate of an atom into the guided mode Γ_{1D} and the decay rate into spurious decay channels that are not collected by the waveguide Γ' . (D) The dispersion plot of the (infinite-length) waveguide, showing photonic band-gap, and the slow-group-velocity region on the edge of Brillouin zone. Image courtesy of M. J. Martin in the article [57].

version of itself, the resulting dispersion curve must be symmetric respected to the crossing point, and hence the slope $v_g = \frac{d\omega}{dk} = 0$ at this point. The guided mode has group velocity zero and effectively ceases propagating.

This conclusion can also be interpreted as time-inversion invariance of modes on the Brillouin zone edge. The guided mode profiles are eigenmodes of the Maxwell's equations subjecting to the Bloch periodic boundary condition. The Bloch boundary condition can be expressed as:

$$\psi(x + a) = \psi(x) \cdot e^{i\vec{k} \cdot \vec{a}}$$

where ψ stands generally for both the \vec{E} and \vec{H} fields of the guided modes. We note that this relation becomes a completely real-valued relation on the edge of the Brillouin zone, where $k_{BZ} = \frac{\pi}{a}$. The resulting mode profile eigenfunctions can

then be made real-valued, hence time-invariant. We also note that the reality of the eigenfunctions prevents circular polarization from being produced at the Brillouin zone edge.

The physical importance of the vanishing group velocity of guided modes near the edge of Brillouin zone lies in the local density of state. The decay rate of the quantum emitter, for example, can be related to the density of state using Fermi's Golden Rule:

$$\Gamma = \frac{2\pi}{\hbar} \cdot | \langle f | \hat{H}_{int} | i \rangle |^2 \cdot \rho(\omega)$$

where ρ is the density of state of the decay channel. For a quantum emitter interacting with a 1D guided mode, the corresponding density of state $\rho(\omega)$ is proportional to the group index $n_g = \frac{c}{v_g} = c \cdot \frac{dk}{d\omega}$. We see that this value diverges on band-edge, suggesting that we can make use of the slow-group-velocity property of photonic crystal to engineer atom interaction strength with the waveguide that is significantly enhanced relative to free-space.




Figure 4.2: Formation of 1D band structure from waveguide dispersion curves

The dispersion of an unpatterend waveguide comprises of forward- and backward-propagating modes (blue curves) that do not interact with each other. The introduction of periodic modulation of lattice constant a creates momentum components shifted by $k_{Lattice} = \frac{2\pi}{a}$ from the backward-propagating mode (green dashed curve), which can now interact with the forward-propagating mode to form an avoid crossing. The crossing generates the photonic band gap and slow-group-velocity region properties (red curves).

4.2 Placement of Photonic Crystal Frequency

There are various methods available to control the band structure of the photonic crystal, first by geometry design prior to fabrication, and then by methods in post-release modification as described in Chapter 3. We introduce dielectric modulation to the waveguide by making the outside edges sinusoidal in shape. This choice of shape pushes the dielectric modulation away from the center of the waveguide, where the electric field is the most intense, hence introducing less longitudinal index contrast to the guided mode per unit added dielectric. This choice allows reduced sensitivity of crystal frequency to geometry changes in modulation dimensions $\frac{d\omega(Bands)}{dx(Geometry)}$, hence also serving to reduce impact of fluctuation of geometry parameters due to fabrication imperfection.

We then carry out post-release modification using feedback from optical measurements. We experimented with plasma etching of suspended devices in [49], and eventually migrated to ALD-coating of devices. For tuning rates of various schemes, see table 4.1. The combination of these two methods provides an approximately $\sim O(1\text{nm})$ wavelength precision for devices after all fabrication and modification processes. The final step to align the bands to Cesium frequencies is done by heating the device in experiment vacuum chambers. For future development, capacitive dynamic tuning would allow active locking of the bands. We have explored this possibility with preliminary positive results; see Chapter 8.

Method	Dielectric		Air	
	Simulated	Measured	Simulated	Measured
Vary a	-0.62	-0.66	-0.66	-0.69
Vary A	-0.13	-0.16	-0.07	-0.11
Vary w	-0.18	-0.11	-0.27	-1.39
Vary g	0.028		0.041	
ALD Coating	-0.77	-0.72	-0.98	-0.94

Table 4.1: Tuning rates of various tuning methods in THz/nm dimension

The parameters are as follow: a stands for the lattice constant; A is the peak-to-peak amplitude of the modulation; w is the width of the double-beam at the thinnest positions; and g is the vacuum gap size between the two beams. See Figure 4.1 (C).

FDTD simulation and Green's function

Finite-Difference Time-Domain (FDTD) method is a valuable tool to calculate the various optical properties of our devices. This method solves Maxwell's equations by discretizing them in time-steps, and sequentially updates each frame in time.

We make use of this in two ways. Firstly, we simulate the response of our devices to excitation from a guided mode, and calculate the field intensity enhancement. Secondly, we calculate the classical electromagnetic Green's function by computing the complex field of the structure responding to the excitation of a classical dipole source.

The Green's function can be converted to the local density of state and hence atom decay rate, as in [58]. The FDTD method also allows separation of dipole decay into the structure guided mode from dipole power radiating into free-space. The Green's function calculation is computationally intensive, as every point in space requires a separate simulation run. In early designing phase, it is convenient to use the field enhancement as a proxy to the full Green's function calculation.

This approximation is justified under a weak-coupling assumption as follow: imagine the dipole excitation in the simulation volume is a classical dipole antenna fed by an imaginary feed-line; then, by scattering theory of time-inversion-symmetric systems, the transmission from one input port (the dipole) into the an output (guided mode of the structure) must equal the transmission of the reverse process; Calculating this reverse process can be approximated by simply calculating the electric field generated by the guided mode, assuming that the coupling to the dipole ports is weak, and hence the field with the dipole antenna present can be approximated by the field without.

The typical process of designing a photonic structure in our project is to first compute its field profile by injection power into the guided mode, optimize the field enhancement, then calculate Green's function only on the optimized structure. Once a device is designed and fabricated, its field enhancement can be optically measured by imaging the scattered light for a given power.

4.3 Tapering Photonic Crystal Waveguides

We create an adiabatic tapering section to transition the unpatterned double-beam waveguide into our photonic crystal. The tapering is done by gradually increasing the modulation amplitude of the crystal. The primary function of this taper is to suppress scattering into free-space during the mode conversion, which is particularly severe for our system due to its relatively low index contrast ($n_{SiN} = 2.0$ for our wavelength range, comparing to $n_{Si} = 4$ for typical Silicon photonics). The tapering also gradually converts the normal group velocity waveguide mode into the slow group velocity photonic crystal mode.

Care must be taken when tapering photonic crystals so that the local band structure throughout the taper does not create a band-gap in the frequency range of interest; see [59]. This can be achieved by symmetrically opening the band-gap by adjusting the width of the un-modulated waveguide. As the frequency range approaches the band-edge of the photonic crystal, the group velocity vanishes, and matching the group velocity to suppress reflection demands impractical length for tapering. Our resulting tapering strategy was to use a short taper to suppress scattering (FDTD calculation showed the scattering suppression plateaus at around 15 tapering sites), and then work with the residual reflection from group velocity mismatch.

Envelop from residual taper reflection

The residual reflection of the taper sections creates effectively a low-finesse cavity around the primary photonic crystal section. Inside this residual cavity, a set of discrete resonances 'super-modes' are formed, defined approximately by phase-matching conditions inside the cavity.

$$\beta := \frac{\pi}{a} - k(x, \omega), \omega_n \vdash \oint \beta(x, \omega_n) dx + 2\phi_{Taper}(\omega_n) = n\pi$$

Near the band-edge, the quantity β determines the residual phase propagation beside the π phase flip associated with the Brillouin zone edge $\frac{\pi}{a}$. The phase propagation β determines the 'envelop' of a super-mode inside the residual cavity. The n -th super-mode then has n lobes along the length of the photonic crystal, with the smaller n modes providing stronger atom-light interaction from their closer vicinity in frequency to the band-edge. Light transmits efficiently through the crystal on these super-mode resonances, and is primarily reflected otherwise. We see that while an infinite photonic crystal has vanishing group velocity at the band-edges, reachable group velocity for a finite-length crystal is limited by its distance to the band-edge, $\beta_{min} \simeq 1/L$, where L is the total length of the photonic crystal.

Phase-matched taper design

We experiment with possibilities to overcome the aforementioned limitation by allowing multiple sections of the photonic crystal. For this purpose we restrict ourselves to few discrete sections that are uniform within themselves, for ease of dimension calibration in fabrication. This is carried out by adding two short sections of 'phase matching' crystals between the scattering-suppression tapers and the main crystal, whose band structure is slightly detuned into the main crystal's band-gap region. The goal of this is to accumulate all phase wrapping into the matching




Figure 4.3: FDTD calculation of field profile and Green's function

This figure plots the excited mode electric field norm $|\vec{E}|$ in V/m of a photonic crystal waveguide, including tapering from unpatterned double-beams, versus the position on the photonic crystal in μm . (A) Excitation frequency on first band-edge resonance. Observe that the field is excited from the left (abrupt onset at $z = 0.5 \mu\text{m}$) and shows slight interference fringes on the reflected (left) side, but not on the transmitted (right) side. Also note that the rapid oscillation is associated with the Bloch mode profile of the photonic crystal with a period of the lattice constant. (B)

Similar calculation but with the injected frequency set in the band gap of the crystal. We see that there is significantly more reflection, and once inside the crystal the field profile decays exponentially. (C) Green's function calculation using FDTD program. A dipole excitation source is placed in the middle of the crystal. The field near the dipole is zeroed on this plot to avoid divergence. We see that the source excites the same resonance mode as in (A), but propagates symmetrically toward both directions. (D) Dipole excitation frequency in band-gap of photonic crystal. Observe that the out-going field on either end of the crystal is exponentially suppressed by the length of the crystal.

crystal section, so as to allow the main crystal to operate near $\beta = 0$. We observed in simulation that this modification allows transmission in frequencies closer to the band edge, and also creates a uniform field intensity inside the main crystal, suppressing the super-mode profile for the vicinity of the matched frequency range.

For a simulated transmission trace; see the lower left panel in Figure 4.4, also note the flattened super-mode in panel (D).

4.4 Cesium Dual-Frequency Operation

Here we describe the alignment of our devices simultaneously to Cesium D1 (335.116 THz) and D2 (351.726THz) resonances. Being able to address two branches of sufficiently different frequencies is critical to our trapping and probing schemes. Forming an optical trap near one atomic resonance while probing the atom states with the other allows high extinction ratio between the two, preventing the high power trap beams from flooding the single-photon-level detection electronics of the probe.

Alignment of PhC frequencies to Cs D1 and D2 lines

A 1D photonic crystal naturally has at least three parameters for tuning, the lattice constant, the average dielectric filling ratio, and the dielectric contrast that defines the modulation. The lattice constant influences the frequency strongly, while it can only be changed by integer multiples of the lithography machine step size. With the sinusoidal shape of our device, the amplitude A and the minimal beam w width vary the mid-band frequency and band-gap size at different rates, with the band-gap size more strongly related to the amplitude. Tuning rates as function of geometric parameters can be found in Table 4.1.

Interestingly, we found that sweeping the amplitude and width parameters in fabrication, even with steps between devices smaller than the nominal value for the lithography machine pixel size, still creates corresponding linear optical frequency response. We assumed this to have resulted from some averaging effect as the machine digitizes the curved shape of the modulation. Initially, a large number of chips and devices are fabricated to build devices within sufficient alignment with Cesium frequencies. Eventually, we settled on the process of building a set of devices with some target frequency spread, and then align the frequencies onto the Cesium frequencies using ALD.

Two-mode trap-and-probe scheme

We designed our photonic crystal with the trap-and-probe scheme in mind, using the half-lattice position difference of the dielectric and air mode of the photonic crystal. One of the two modes would be used as the 'trap' mode, which is blue-detuned from the Cesium transition line in its vicinity. The blue-detuned mode creates an

AC Stark shift that effectively forms a positive potential to repel the atoms toward its intensity minimum. The trap mode serves two functions, firstly to produce a potential barrier around the bulk of the device to prevent atoms from crashing into the material, and secondly to guide the free-flying atoms toward the section in the lattice where the trap mode has local minimum and the probe mode has maximum. The probe mode operates at the other edge of the photonic band-gap for maximal coupling.

This picture becomes realized when both bands are precisely aligned to the band-edge, which unfortunately was not achieved, due to both fabrication imprecision and also the finite-length device limitation as mentioned earlier in 4.3. The symptom of non-exact band-edge alignment manifests itself as non-zero intensity at the trap mode minimum location. This creates a potential bump, as shown in Figure 4.5. Adding red-detuned light to the probe mode to create attractive potential can partially alleviate this issue, at the cost of weakening the barrier near the device walls and potentially interfering with the out-coming atom signal. We only had limited success with this trap, due to the small volume of trapping sites, and the limited solid angle through which atoms can arrive at the device.

Side-illumination trapping scheme

We have also experimented with a 'side-illumination' method as used in [61]. In this scheme, an external red-detuned beam is irradiated onto our device. The reflection of such beam from the device itself interferes with the incoming beam to form a standing wave. This standing wave trap is localized to the vicinity of the device, as the sub-wavelength device scattering field decays rapidly with distance. This trapping site is positioned approximately 120nm above the top plane of the device surface, exposing it to more solid angle to accept atoms, at the cost of approximately a factor of five reduction in coupling strength to the guided mode as compared to trapping the atom in the center of the device layer; see Figure 4.6.

As the device is used mostly as a reflector, the properties of this trap is relatively independent of the device band placement. However, the macroscopic nature of the incoming beam creates a large trap parallel to the device, so only the averaged response of the atoms over both the device mode profile ($\sim \cos^2(\frac{\pi}{a} \cdot x)$) and the device super-mode ($\sim \cos^2(\beta \cdot x)$) can be observed. Techniques to attempt loading of this external trap into the guided mode trap are being developed.

4.5 Additional Engineering Options

The photonic crystal presented here actually guides several modes of different symmetries. The majority of properties we made use of are associated with the guided mode that has the same symmetry as the fiber mode, for ease of coupling from the fiber. That includes two modes associated with the two orthogonal polarization modes of the fiber. With the planar nature of our devices, we define the two modes using their symmetry property to the mirror plane parallel to the chip plane. We refer to the mode with anti-symmetric ($H_{\parallel} = 0, E_{\perp} = 0$ in the plane $z=0$) as the 'TE' mode, and the symmetric ($H_{\perp} = 0, E_{\parallel} = 0$) as the 'TM' mode. The naming is associated with the dominant transverse field in the device plane. We can further make use of the mode profile center-line symmetry to create symmetric and anti-symmetric profile modes, which compounds with the in-plane symmetry to form four families of modes.

'TM' modes

For the TM mode, the electric field is primarily perpendicular to the device plane. This direction of the electric field is typically in the thin direction of the device (thickness of the Nitride is 200nm, comparing to the 500nm to $1\mu\text{m}$ transverse size of the waveguides), and the TM mode typically has lower effective mode index and higher band-gap frequency than the TE modes of the same geometry. A dispersion plot that includes both the TE and TM modes for an APCW device can be found in Figure 1 (C) in [63].

The first application of the TM mode is using it as an alternative channel for the photonic crystal designed for TE mode. The TM mode behaves like an effective waveguide mode throughout the frequency range, where as the TE mode demonstrates its band-edge and band-gap properties. This allows a convenient 'baseline' channel for the atomic experiments to probe the atoms without group velocity or residual cavity enhancements. Alternatively, it is possible to create devices specifically designed for the TM modes, with similar slow group velocity and band-gap properties. Our experience with TM devices showed less irregularity in their optical spectra, perhaps due to the better surface quality of the LPCVD surfaces as opposed to the plasma-etched sidewalls. Due to the lower effective index contrast, however, a more aggressive geometry is required to form a sufficiently large band-gap.

Anti-symmetric modes

It is possible to map a symmetric guided mode into the anti-symmetric mode using suitable tapering methods. This allows an extra set of usable modes for use in our devices. Here we present a design that hybridizes the symmetric and anti-symmetric modes in our photonic crystal, as a tool to engineer exotic properties such as zero-group-velocity point inside the interior of Brillouin zone, and ultra-flat bands. By shifting the two beams of our double-beam sinusoidal crystal relative to each other in the propagation direction, we controllably break the left-right symmetry of the device, causing symmetric and anti-symmetric modes to hybridize. This creates an avoided crossing inside the Brillouin zone, whose frequency can be tuned by changing the relative strength of the sinusoidal modulation amplitude and the relative shift between the two beams.

We can potentially operate at this anti-crossing point to employ a guided mode that is both strongly enhanced by slow group velocity and having non-vanishing propagation constant. Alternatively, since the frequency of this anti-crossing point can be independently adjusted, it is possible to match the frequency of avoided crossing to the frequency of Brillouin-zone-edge to create a flat band for possible ultra-strong slow group velocity enhancement; see Figure 4.8 for the mode profiles and dispersion of such a device. We found it to be possible to taper into this hybridized photonic crystal by first tapering into the symmetric APCW using the standard taper as before, and then gradually taper into the shifted crystal by increasing the lattice constant of one beam while decreasing that of the other.

Circular polarization engineering

The capability to engineer circular polarization using nano-structure is valuable, as the coupling of propagation direction with polarization can potentially create interesting properties. It will allow the possibility of optically pumping of atoms in m_f sub-levels, or address closed transitions on the Cesium D2 line ($F=4$ to $F'=5$). Further properties such as directional emission have also been demonstrated in literature [64, 65]. For our photonic crystal waveguide, this can be engineered for devices operating in the interior of the Brillouin zone. For a symmetric crystal design like the APCW, it is difficult to create circular polarization on the band-edge modes, as claimed by the time-inversion argument in Section 4.1. However, the intra-Brillouin-Zone slow-light point as presented in the previous subsection may allow us to produce a slow group velocity device with circular polarization.

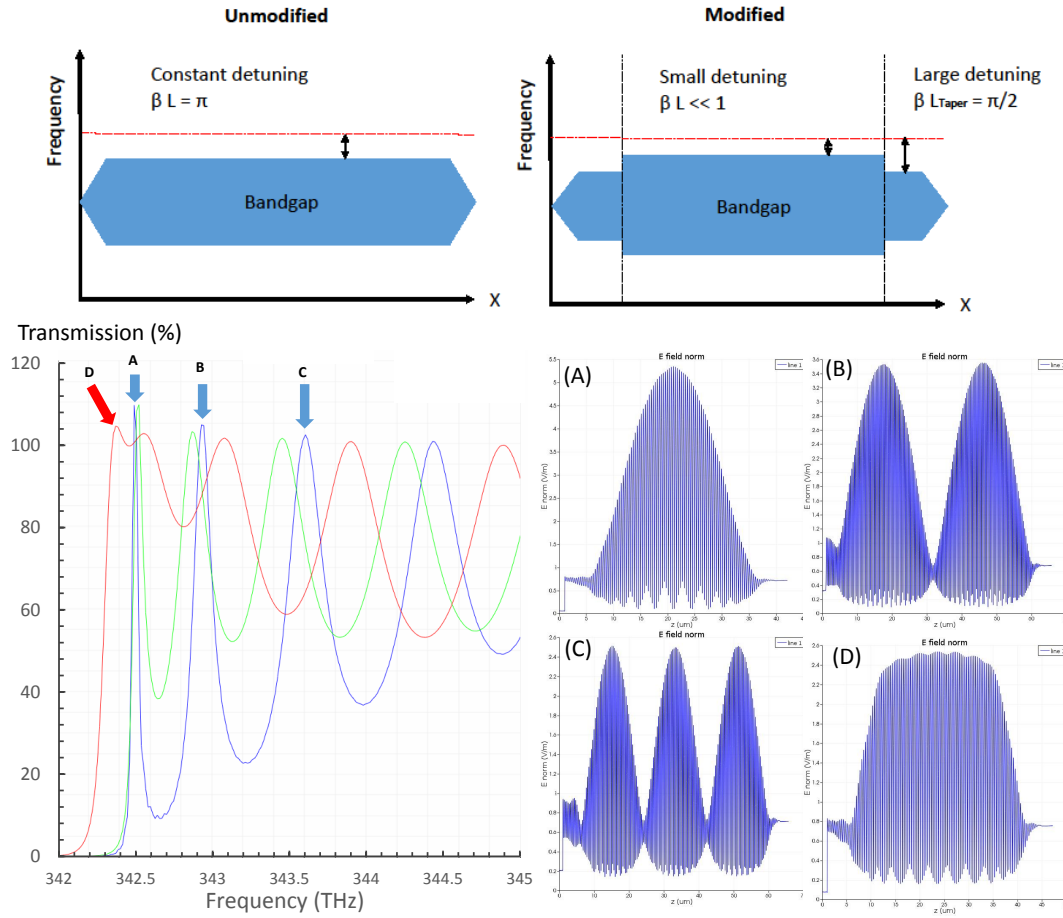


Figure 4.4: Residual cavity induced resonances near band-edge

Top: Illustration of band-gap opening by tapering in the propagation direction.

Left plot depicts simple linear tapering, right plot shows the phase-matching modification. Lower Left: Transmission through a photonic crystal device, including linear tapering. The color traces (blue, green, red) show gradual onset of the phase matching modification parameter, here a 5nm increase in minimal width w of the photonic crystal for the first and last 17 sites. With the phase-matching design, we observe that the inner-most resonance peak is moved closer toward the band-gap. Also note the decrease in frequency spacing between the inner-most resonance peaks as the phase matching onsets, indicating reduced group velocity. Lower Right: The band-edge resonances due to tapering reflection. (A)-(C) are the first, second, and third resonance super-mode profiles of simple linear tapered structure. (D) is the first resonance super-mode of the phase-matched design, where the middle section of the crystal has approximately uniform field intensity.




Figure 4.5: Dipole trapping potential and Casimir-Polder forces

This figure plots the potentials for atom near the device. The device propagation direction is along X, perpendicular to the plot axes, and the height is the potential at the 2D cross-section plane. (A) Casimir-Polder force associated with the dielectric. (B) Ideal case of blue-detuned trap with perfectly band-edge aligned device. Note that the raised potential barrier around the device that prevents atom from dropping into the Casimir-Polder potential wells and crashing into device side-walls. (C) Potential with slightly misaligned device, with $k_{Actual} = 0.99 \cdot k_{BZ}$. We see that the designed trapping site between device beams has become a saddle point. Image courtesy of C.-L. Hung [60].

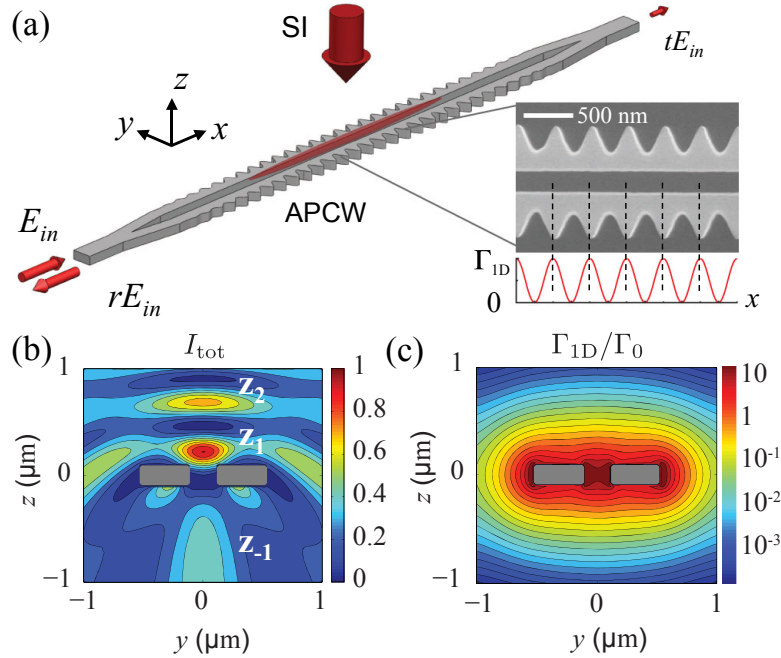


Figure 4.6: Atom trapping using side-illumination

(a) An external beam is pointed toward the device from perpendicular direction to the chip plane. A local trap is formed by the reflection off the device, and limited in transverse x direction by the spot size of the beam. (b) Multiple local trapping sites are formed due to nature of the standing wave. However, only site Z1 provides sufficiently strong coupling strength to produce a significant signal, as shown in (c). Image courtesy of A. Goban [62].




Figure 4.7: TM mode profile at band-edge

This figure shows the cross-section plots of the TM modes, with Y direction being the propagation direction. The TM modes extend out of the device plane further than the TE modes, due to the electric fields primarily polarized in the perpendicular direction to the device plane. At the Brillouin zone edge, the typical dielectric-mode-air-mode contrast manifests itself in the same manner as the TE modes.




Figure 4.8: Shifted APCW to create very flat dispersion curve

(A) For photonic crystals with center-line mirror symmetry, the symmetric modes (blue solid) and anti-symmetric modes (red dashed) do not interact. The dielectric band of the anti-symmetric mode, being of slightly higher frequency than the symmetric mode, can cross the air band of the symmetric mode. When the mirror symmetry is broken, this crossing becomes an avoided crossing. (B) The dispersion of shifted APCW optimized to flatten the third (red) band, for this plot, the relative shift between the beams is 0.3913 lattice constant. With this parameter set, the last 2% of the Brillouin zone spans only 200GHz in frequency. (C) The mode profiles for these modes demonstrate complicated mode patterns, but still with high intensity contrast in the gap between the beams.

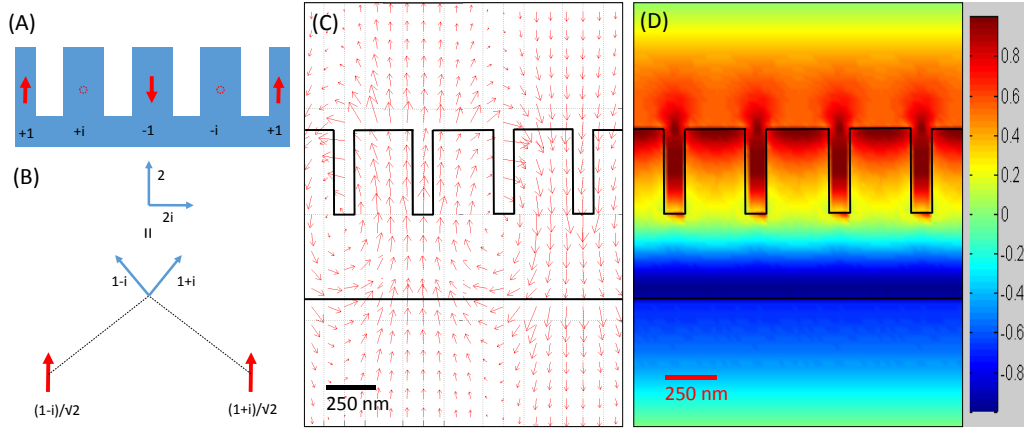


Figure 4.9: Systematic engineering of circularly polarized 'pockets'

(A) A photonic crystal can be intuitively pictured as an array of polarizable objects. If we choose to operate at, say, $k = 0.5 \cdot \frac{\pi}{a}$, then each subsequent dipole will have a $\exp \frac{\pi}{2} = i$ phase shift comparing to the previous one. (B) For convenience, we factor the i phase shift into $\frac{1+i}{\sqrt{2}} / \frac{1-i}{\sqrt{2}}$. The total field pattern of the crystal is contributed by the sum of each dipole. Due to the field pattern of electric dipole, we find points on the center plane between two dipoles, where the x component of the field is completely imaginary, while the y component is completely real, hence creating a circular polarization point. (C) FEM simulation using a test structure to verify the claim. Note how the field inside the gaps on the upper half of the waveguide rotates by 90° every period. (D) We can mathematically describe the field at any given point as linear combination of left- and right-circular modes. The color plot shows their intensity contrast, $C_I := \frac{I_{left} - I_{right}}{I_{left} + I_{right}}$.

Chapter 5

OPTICAL CHARACTERIZATION OF DEVICES

This chapter will describe the optical characterization methods for our devices. We optically characterize our devices in several steps after their fabrication to verify their properties, including coupling efficiency as measured in reflection and transmission levels, band alignment, quality of band-edge resonances, and power tolerance of the devices in vacuum. The information gathered during characterization helps estimate the total enhancement we could achieve with each device, and allows us to select which devices we would use for the atom experiments. The characterization methods are applicable to all photonic crystal types of our devices. For this chapter, we focus on the APCW. For optical measurement results regarding the 2D or slot waveguide devices, see Chapter 7.

5.1 Reflection and Transmission Measurement

The straight-forward measurements for the 1D devices are the reflection measurement from each of its ports, and the transmission from one port to another. The measurements are relatively simple and can be quite informative, provided that proper calibration procedure is devised, for both the power level and the absolute frequency. Ultimately, transmission and reflection measurements are employed in our experiment for atom-light interaction.

Light sources and analysis tool

We make use of a total of three sets of light sources. Firstly, for broad-band spectrum measurements, we make use of a superluminescent light-emitting diode (*Inphenix Inc IPSDM0822*) as light source, providing an approximately 100nm bandwidth around our wavelength of interest. This light is coupled into our devices, and then the out-coming light is analyzed using an optical spectrum analyzer (*Anritsu Co. MS9740A*). This configuration allows us rapid testing of devices to a precision of 50GHz as limited by the spectrum analyzer (1~5 minutes per device depending on resolution). Then, we have two external-cavity diode lasers that are designed to sweep approximately 20nm around Cesium D1 and D2 lines (*New Focus Velocity*). These single-frequency sources have better resolution (~10pm typical) and provide the capability to directly reference a Cs vapor cell, where saturation spectroscopy is




Figure 5.1: Residual reflection and interference fringes of different coupler designs

Broad-band reflection spectra of (A) conventional non-modified coupling-waveguide-tether junction (B) the 'roundabout' junction design (C) the free-space coupled 'Y' design. We see the very rapid oscillation on the band-gap reflection spectra, due to interference between photonic crystal reflection and device residual reflection. With development of anti-reflection coupling junction design and application of anti-reflection coated fibers, the reflected fringes are improvingly suppressed.

used to provide reference. Finally, it is possible to take broad-band, high-precision spectrum using a Ti:sapphire laser with feedback control from a high-performance spectrometer (*M-squared SolsTiS*). Our experience with this setup is that a full sweep over the bandwidth of interest can take up to an hour per device. All of these sources have varying spectral shape, and hence the power of the source independent of the device must be monitored and used for normalization to achieve proper calibration.

Polarization alignment

The optical response of our devices is strongly polarization dependent. We map a linear polarization source to the desired polarization using a zeroth-order half-wave plate, then through an additional zeroth-order quarter-wave plate before sending the light into single-mode optical fibers to the device. Here the quarter wave plate serves to cancel the birefringence associated with the bending of the optical fiber. We opted not to use a polarization-maintaining fiber for this purpose, in order to avoid using complicated stages to physically rotate fibers to align to the photonic devices.

We align the input polarization to our devices using several methods. Preliminary alignment is done by observing the scattering of light from the first tether junction

on the coupling waveguide. The scattering power for the TE (E field in-plane) mode for the first set of tethers is simulated to be approximately eight times stronger than the TM mode. Also, since an electric dipole does not radiate toward the axis of the dipole vector, the out-of-plane scattering of the TM mode is further reduced. For finer alignment, maximizing band-gap reflection is used for reflection measurements, and minimizing band-gap transmission can also be used by either imaging the waveguide downstream to the photonic crystal, or by measuring the transmitted power directly.




Figure 5.2: Optical image of coupling waveguide tether junction scattering

A cleaved optical fiber to the left of this image injects light into the coupling waveguide of our device to the right. The waveguide is supported mechanically by the thin tethers near its end, and this junction produces polarization-dependent scattering. Minimizing this scattering intensity in the optical microscope aligns the injected polarization to TM. The polarization is then rotated to TE mode by rotating the half-wave plate at the fiber input.

Measurement using cleaved fiber-end

The majority of our devices are designed with V-shaped grooves for placement of fibers. This offloads the effort of alignment of the fiber core to the coupling waveguide to the fabrication process of the chips. During testing, the fiber is pressed down into the groove using a five-axis stage. The elastic nature of the fiber allows it to conform to the alignment of the chip. The fiber position only needs to be adjusted in the propagation direction, with the optimal distance being as close as possible to the device. Care must be taken here due to the limited resolution and depth of field of our long-working-distance optical microscope. It is often

difficult to resolve the exact distance between the coupling waveguide and the fiber end. Collision between the two can cause irreversible damage to the waveguide. To apply this method in vacuum requires either motorized stages in vacuum, or gluing multiple fibers onto the chip before loading the chip into vacuum.

A bare cleaved fiber can be used for this measurement, but the refractive index mismatch between the fiber ($n=1.45$) and air or vacuum. This is a 3.5% reflection that can cause large amplitude fringes on reflected signal. We eventually migrated to using anti-reflection coated fibers (*OZ Optics*, AR coating designed for 790-890nm range), which maintains high transmission and suppress the reflection to $\sim O(0.1\%)$. Conveniently, we can make use of this index mismatch reflection as a reflection level calibration standard. For the anti-reflection fiber case, we cleave the fiber further upstream to calibrate away the efficiencies of fiber unions and other coupling loss in our setup. We found properly fusion-spliced fiber to have negligible loss for our purpose.




Figure 5.3: Illustration of the characterization setup

The reflection measurement is carried out using a 50/50 fiber beam splitter, by launching light from one input arm, then detecting the light returning to the other input arm. The calibration is achieved by cleaving the fiber after the output arm of the fiber beam splitter and measuring the index mismatch reflection. The spectral shape of the source is monitored by measuring the other output arm of the beam splitter. Image courtesy of J. D. Hood [63].

Measurement using free-space optics with cleaved chips

For the cleaved free-space chips, the alignment of input light into the device is off-loaded onto the testing setup, and no longer depends critically on the V-groove size. The free-space test setup comprises of one five-axis stage to place the chip, and two sets of mirrors and lenses on translational stages in a fiber-coupling configuration. Due to the larger number of degrees of freedom involved with the free-space testing setup, the optics on the coupling setup need to be sequentially pre-aligned prior to

measuring of a chip, typically by back-reflecting the light into the launching fiber. Then, the focusing lenses are translated toward the device to allow coupling.

The initial alignment to the chip can be done by watching the scattered light off the coupling waveguide tethers under a microscope. Then, several iterations of optimization of reflected signal using all degrees of freedom is carried out, in a similar manner to conventional fiber coupling. This is typically only needed for the first device on a chip. For the subsequent devices, the five-axes stage with their piezo-electric stack actuators are generally sufficient to align the device to obtain reasonable coupling. We note that it is generally easier to initially align the system using the broad-band source, as the single-frequency sources can suffer from strong fluctuation in reflection channel from residual reflection. Calibration of absolute level of reflection and transmission is more difficult for the free-space case, due to the larger number potential sources of loss.




Figure 5.4: Optical image of coupling of device using free-space optics

Light is injected using aspheric lens from the left (not shown) into a device on a cleaved chip. The scattering pattern of the device layer can then be used to determine the pointing of the incoming beam. In this image, we observe a slightly asymmetric scattering on both the coupling waveguide tether, and the first tether array. The location of the beam waist can be approximately determined by translating the lens, as if doing a conventional knife-edge technique by minimizing the scattered spot size at the device tether. The rough alignment typically is sufficient to obtain sufficient signal, which can then be optimized in a similar fashion to conventional fiber coupling.

5.2 Features of Measured Spectra

There are several features to look for on a measured spectrum, including the overall level, frequency alignment of the band-gap, band-edge resonance regularity and quality, residual transmission of the band-gap, and mixing between non-collected guided modes.

Power level and band-gap alignment

The first measurement of a fabricated device is the power reflection and transmission efficiency. This is typically done with the broad-band source and optical spectrum analyzer. The power level can be influenced by coupling waveguide dimension, fiber misalignment due to V-groove size, or device cleanness or damage. The first two can usually be parameter-swept on chips to find out the optimal values, so errors in those two situation can be diagnosed using SEM prior to optical measurement. Observing of scattering from devices during optical testing provides some information on the overall cleanness, but more quantitative measurement is done by building cavities. We check whether the band-edges are in the scan range of the scan-able diode lasers at this stage.

Band-edge resonance regularity and quality

With finite-length structures, a gradual onset of reflectivity with discrete resonances occurs as the device transits from waveguide-like mode into band-gap. Intuitively, we can picture the photonic crystal as an effective material with frequency-dependent group index, with 'mirrors' formed by the residual interface reflection between the crystal and the unpatterened waveguides outside, hence forming a cavity-like system. A simple model assuming the group index resulting from a quadratic dispersion near band-edge can be applied to fit the resonances using few parameters, which will help us characterize the Van Hove singularity at the band-edge.

The actual measured response of the devices would be influenced by the imperfection in fabrication, adding some level of disorder into the device dimensions, and hence the local band structure of the crystal. The disorder than manifests itself as non-ideal properties in the band-edge features, such as non-monotonic reflection level as the frequency approaches the band-edge, or the variation of frequency spacing between resonances that does not correspond to the predicted increase of group index. Propagation loss and scattering loss of the device would also result in non-ideal features of the band-edge resonances. In our intuitive cavity model, the loss would cause the cavity to go from over- or critical-coupling to under-coupling, and

cause the resonance dips to have non-unity contrast. This is especially severe with the resonances closer to the band-edge, since the loss would be enhanced by the (frequency dependent) slow group velocity.

Residual transmission through the band-gap

Theoretically, the transmission of light through the device when the frequency is in band-gap should be exponentially suppressed. In the transmission measurements of our devices, we often observe a level of residual transmitted light that limits our extinction ratio to around 25dB. We suspect that the light is bypassing the crystal through different symmetry or higher order modes. We noticed the extinction ratio is frequency dependent, and re-optimizing input light polarization at each frequency seems to allow better extinction ratio at that frequency. This could have resulted from residual frequency dependence of our polarization control setups. The other possibility is that the fabrication imperfection on the device, such as side-wall roughness or slight asymmetry of the device due to lithography precision limit or side-wall tilt due to plasma etching process, is scattering the injected mode into higher order modes.

We often observe dips in reflection of the band gap near the air band edge, which does not correspond to an increase in measured transmission at the same frequency. Simulation suggests that the anti-symmetric guided mode of our photonic crystal structure has its dielectric band-edge in the frequency range of these dips. We suspect these dips are cross-coupling of symmetric modes into the anti-symmetric modes. The anti-symmetric modes then become lost to free space as the waveguide tapers down to single-mode regime to couple to the fiber. For future development, mode-filtering structures may be considered.

5.3 Interaction Enhancement Estimation

We test our devices thoroughly before installation into the vacuum chamber for atom experiments. The most interesting property we look for are the potential to enhance atom-light coupling to the guided modes of the device. At this stage we approximate this property as the field enhancement of the device compared to non-group-velocity-enhanced waveguide. This is convenient as field enhancement can be directly measured.

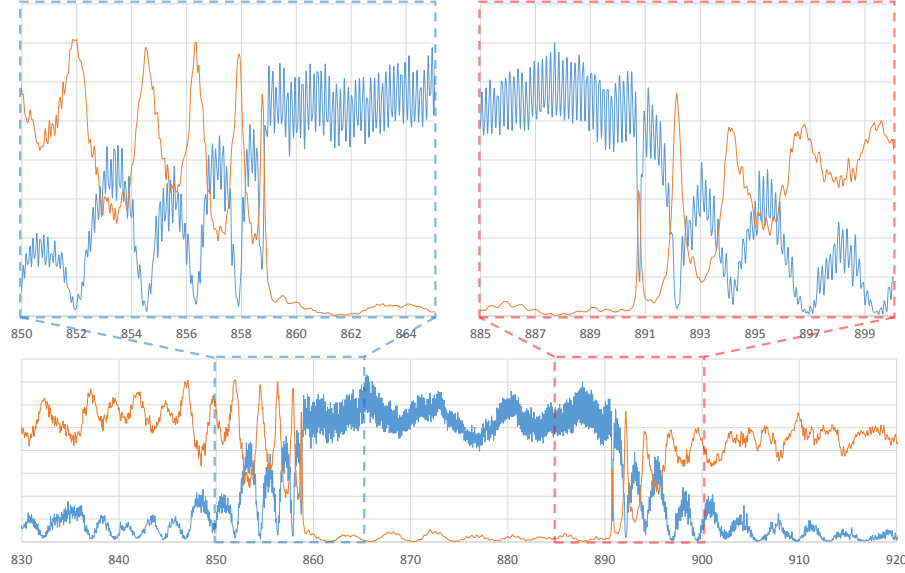


Figure 5.5: Features of a device in reflection and transmission measurement

The reflection and transmission spectra of a device shows many features. On the broad-band spectra (bottom panel) we observe the dip in reflection level in band-gap (blue trace), which does not match with the bumps in transmission of the band gap (orange trace). We also see some disorder in the band-edges resonances, such as the fluctuation in peak heights in transmission and non-monotonic change in reflection spectrum. The onset of decreased group velocity of light is visible from the spacing of the resonances, which reduces as the frequency comes closer to band-edge. Note that the transmission trace has been shifted blue by 0.4nm on this plot to compensate for some frequency drift on this chip between the two sets of measurements. Also note that the Y axes of the reflection and transmission traces are scaled relatively to allow ease of visual comparison.

Using free spectral range of resonances

Following our cavity model, we can estimate the field enhancement by first fitting out the cavity parameters. In the vicinity of the band-edge, we can assume a quadratic dependence of frequency on the residual phase propagation vector:

$$\omega = \omega_{band-edge} + \frac{1}{2}\alpha\beta^2$$

where α is the curvature of the dispersion at the band-edge. The group index then demonstrates a Van Hove singularity near the band-edge, $n_g = \frac{c}{d\omega/d\beta} = \frac{1}{\alpha\beta}$. Using this simple model and assuming the resonances correspond to $\beta_n = \frac{n\pi}{L}$, the group index can be estimated from the positioning of the band-edge resonances; see Figure 5.6.

We can then infer the reflectivity of the residual cavity by fitting the resonances for

their FWHM linewidth κ , and the finesse of these cavities can then be estimated as $F = \frac{c/2Ln_g}{\kappa}$. Finally, for a quick estimation of enhancement, we can assume a local effective free-spectral range exists for each resonance. Under this assumption, the enhancement can be estimated as

$$Enhancement = F \cdot n_g = \frac{c}{2L \cdot \kappa} = \frac{Q}{2} \cdot \frac{\lambda_{free-space}}{L}$$

This is useful for estimating the enhancement on-the-fly, in that this can be calculated for each resonance locally by identifying the cavity quality factor and device physical length. For a more thorough analysis using an optical transfer matrix method, which can include the behavior of the taper sites, please see the thesis [50] and [49].

Using scattering measurement

The more direct method for measurement of field enhancement is to measure the local intensity of the device modes by imaging their scattering. This scattering is created by randomly occurring scattering sources, likely from fabrication imperfection. The scatterer strength is generally not uniform in spatial distribution. The scatterer profile has to be calibrated away by using a baseline average of the device scattering far away from the band-edge. This method produces the additional information of the profiles of the super-mode as function of input light frequency, and some level of information regarding disorder. Up to this point, direct imaging of the scattered light provides us with the most detailed information regarding our devices. For future development, it may be possible to use a double-period modulation of the crystal lattice to engineer test devices with controlled scattering, as the double-period modulation maps the guided mode from band-edge to the Γ point and creates perpendicular, out-of-plane scattering.

5.4 Power Handling Capability Test in Vacuum

There are many aspects that can influence the power tolerance of our devices. The amount of optical power that can be put into the guided modes ultimately limits the types of optical trapping schemes our devices can support. The cooling structure design being of foremost importance, but it is also strongly influenced by processing steps such as cleaning and the existence of the ALD coating layer.

Nano-scale devices have high surface-to-bulk ratio, and tend to cool very efficiently in air. Thermal testing of our devices are typically carried out in moderate vacuum ($\sim O(10^{-6})$ Torr) that can be reached using a turbo pump. We calibrate for the power that arrives at the device by taking the absolute reflection power of the fibers and

chip in vacuum, and then taking the square root of such a value to estimate the one-way efficiency. This allows us to estimate the power that actually enters the device and reaches the photonic crystal.

The failure point for our nominal ALD-coated APCW devices are typically in the few-hundred- μW range, with the highest measured value being $375\mu\text{W}$ off-resonance power at the photonic crystal. The typical point of failure is the input coupling waveguide, perhaps due to its very thin dimension. We hope the smaller device size for the free-space coupling and the more robust structure of the Y-shaped coupler can improve this issue. In case of resonant driving on one of the band-edge resonances, the photonic crystal section can fail as well from the enhanced local field.

5.5 Precision shifting to Cs frequencies using optical measurement feedback

We make the last stage of device band adjustment using the atomic layer deposition technique; see Figure 5.9. Its advantages are two-fold. Firstly, its discrete-layer tuning makes it precise and reliable. Secondly, the ALD process is a dry process, and hence it can be applied to devices in multiple runs and optically test them in between with relatively low risk of damage or contamination. We measure the devices optically, and then compute the number of layers for ALD process that would optimally place the device band structure near the desired frequency range. This process has allowed us to at least align one band to the vicinity of a Cs resonance with sub-nm wavelength precision. We note that different bands of a photonic crystal can have very different tuning rates versus the number of ALD layers. Also, the coupling waveguide width providing optimal coupling efficiency would also shift as the ALD layers are deposited; see Figure 5.10. In addition, we found that post-processing the ALD coated chips with Oxygen plasma causes a slight red-shift of the device bands, and also seems to improve the device power tolerance.

5.6 Characterization of Single-Beam Fabry-Perot-Type Cavity

The single-beam Fabry-Perot type devices have the advantage of being simple to analyze. The single-beam cavity system produces sharp, well-defined resonances that are evenly spaced. The resonance peaks can be fitted directly with a Lorentzian lineshape to obtain their quality factor; see Figure 5.11. For the devices presented here, quality factors above 10^6 have been measured. These cavity devices can be used as 'standard candle' devices for testing the fabrication process, and they also provide potential to be used in cavity QED experiments.

Separation of propagation and mirror loss

The two elements of the nanobeam cavity are the mirrors (including the tapering, transmission, and scattering) and the nanobeam. Each would contribute to the total cavity loss, but since the propagation loss of the nanobeam scales with the length of the nanobeam, the two can be isolated. We fabricate an array of devices with identical mirror design and tapering scheme, and vary only the length of the single nanobeam. We then linear fit the one-way loss of the cavity with $\eta_{one-way} = \eta_{mirror} + \alpha \cdot L_{geom}$ to extract the mirror loss η_{mirror} and the propagation loss per unit physical length α .

We found that the propagation loss to be strongly dependent on the fabrication process. For example, devices made from critical point drying would initially have a propagation loss five times higher than what can be achieved after an Oxygen plasma cleaning. We confirmed the efficiency of our cleaning process steps using the cavities. We also found the propagation loss of the single nanobeam reduces with increasing width of the beam. This phenomenon is more pronounced in TE cavities than TM, which we suspect is due to the higher surface quality for the top-bottom surfaces formed by the LPCVD process compared to the side surfaces from the etching process. Increasing the beam width from 500nm to 600nm reduces the propagation loss for the TE mode by a factor of four. We observed the TM mode propagation loss to asymptote at around 400nm beam width to a value lower than the TE mode, despite its more evanescent mode profile.

Reachable parameters and potential for cavity QED experiments

With our current process, a minimum propagation loss of 0.08% per 100 μ m was achieved for the TE devices, and 0.05% per 100 μ m for TM. Mirror reflectivity above 99.7% was also achievable using our system. With these parameter ranges, we can estimate the cavity QED parameters of a cavity device, assuming the following: The physical length of the cavity is assumed to be $L=250\mu$ m, with a quality factor $Q = 10^6$. This gives a cavity decay rate of $\kappa = 2\pi \cdot 168MHz$. We define the effective mode area defined as [60]:

$$A_{eff} = \frac{\int \int \epsilon(\vec{r}) |\vec{E}(\vec{r})|^2 d^2\vec{r}}{\epsilon(\vec{r}_t) |\vec{E}(\vec{r}_t)|^2}$$

where \vec{E} is the guided mode profile \vec{r}_t is the trapping point. Assuming a hypothetical atom trapping scheme at 200nm above the top surface of a TM cavity device, we calculate the effective area to be $A_{eff} = 2\mu m^2$. Now, taking the atom decay rate to be the free-space decay rate of the Cesium D1 line $\gamma_{d1} = 2\pi \cdot 2.29MHz$, we estimate

the cooperativity of our cavity to be:

$$C = \frac{3}{8\pi^2} \frac{\lambda^3}{L \cdot A_{eff}} \cdot Q = \frac{3}{8\pi^2} \frac{(0.894\mu m)^3}{(250\mu m) \cdot (2\mu m^2)} \cdot 10^6 = 54$$

Or, in the form of interaction rates, $(g, \kappa, \gamma) = 2\pi \cdot (130, 170, 2.3)\text{MHz}$.

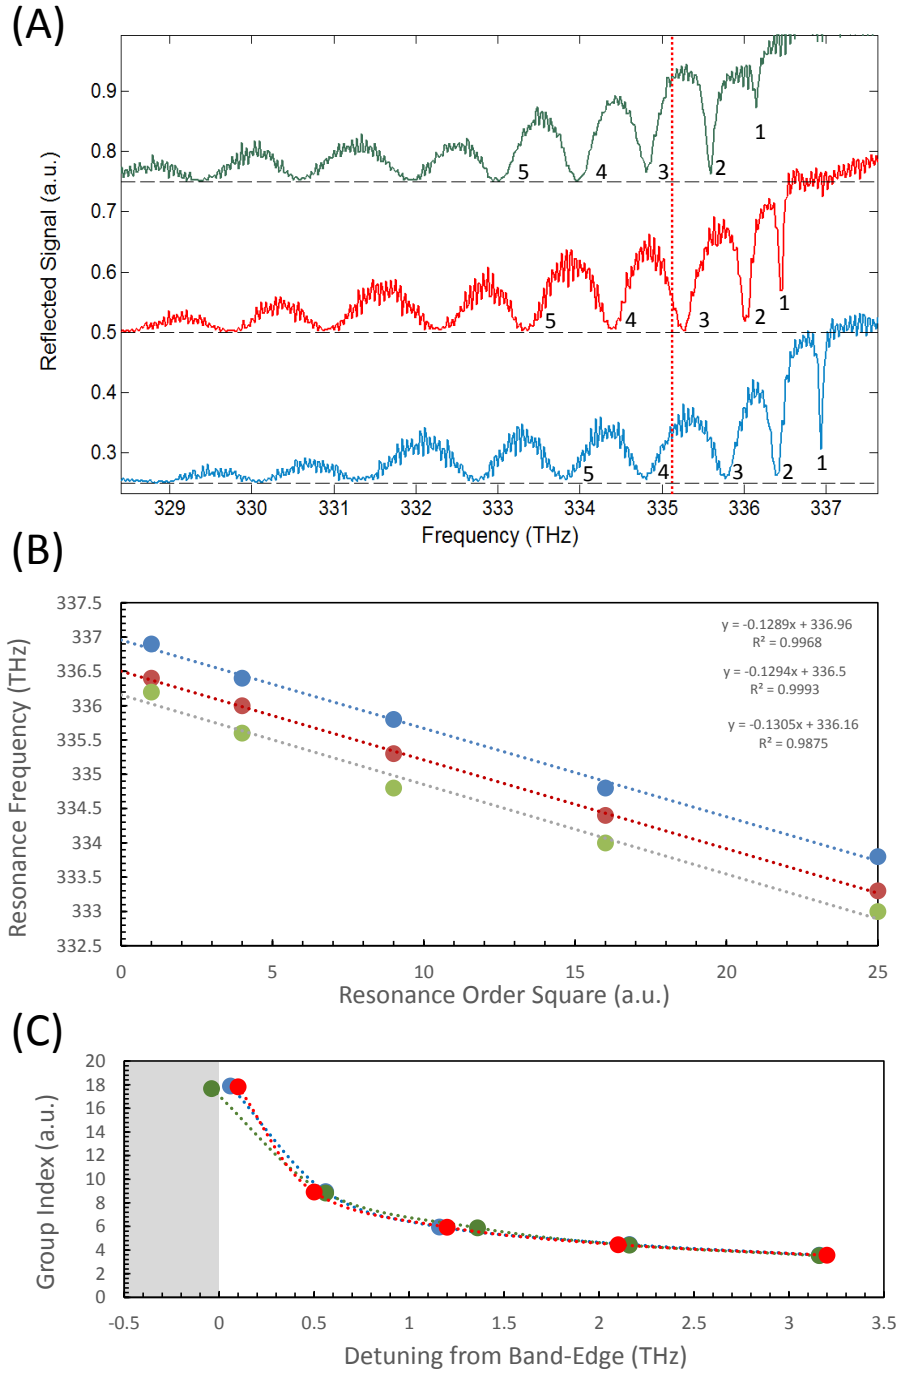


Figure 5.6: Analysis of band-edge resonances using simple cavity model with quadratic dispersion

(A) The frequencies of the band-edge resonances are recorded as a function of their resonance order. For this plot, the dielectric edge is analyzed. The vertical red dashed line indicates Cesium D1 frequency. (B) The resonance frequencies are plotted against the square of their resonance index. The true band-edge frequency $\omega_{band-edge}$ and the band-edge curvature α can then be extracted using linear fit. The R^2 values are shown on plot for reference. (C) With further information of the physical length of the devices, the group index $n_g = \frac{c}{v_g}$ can be estimated.




Figure 5.7: Features of a device in reflection and transmission measurement

Scattering light from a device as function of frequency, normalized using scattering level in frequency range far from the edge of the band. The x axis is the frequency detuning from the band-edge. The prominent features are the band-gap on the right, and the several visible resonances v_1 to v_5 . We see that the resonances (say, v_1) are indeed significantly brighter than the scattering further away from the band-edge (v_5 and beyond). More information can be found in [63] and the thesis [49]. Image courtesy of J. D. Hood.




Figure 5.8: Device damaged in vacuum chamber

This is a device that sustained damage during the experiment in the atom vacuum chamber. This device suffered abrupt failure, the waveguide ruptured and became tangled onto one of the side-rails. Its breaking point appears to have some form of continuous deformation almost as if it were melted at some point. We observed that the gradual deposition of Cesium onto devices during experiments degrade device quality and increase their optical absorption. The absorbed optical power eventually led to thermally-induced failure of the devices.




Figure 5.9: Shifting the band edges using ALD coating

In anticipation of ALD coating, the chip is fabricated with pre-adjusted parameters tuning its devices toward higher frequency by approximately 8THz. The adjusted frequencies are confirmed by measurement immediately after conclusion of chemical processing (left panel, blue and green). The target for the band placement was that the dielectric band is to be 1THz higher than Cs D1 line, and the air band is to be 1THz lower than the Cs D2 line, in order to allow transmission of both atomic frequencies. Since there is a parameter sweep creating a 2THz distribution of the band-edge frequencies. We can then select how far to operate from the band-edges by selecting the suitable devices. The tuning is complete after coating the devices with ALD for 95 cycles. A precision of 0.3THz was achieved for the 8THz shift, as plotted on the two panels to the right.

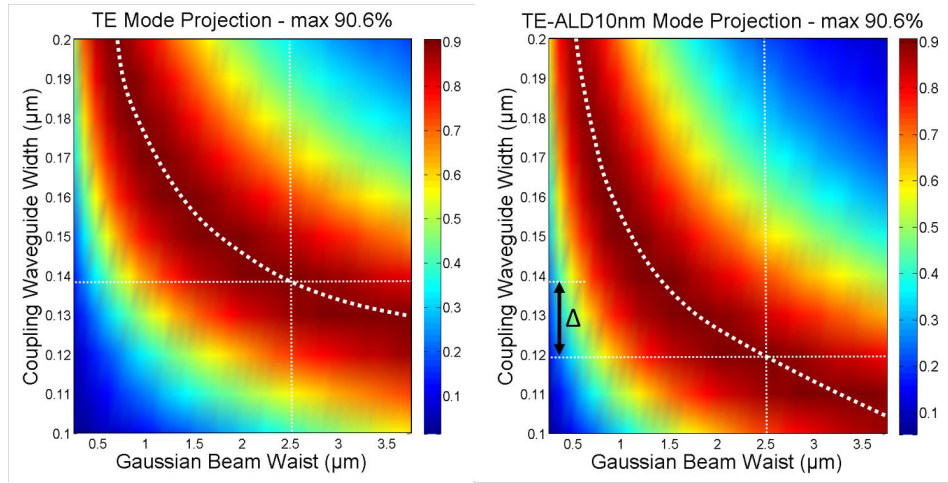


Figure 5.10: Shifting the Nitride coupling waveguide optimal for 10nm ALD coating

Using the same simulation method as in Figure 2.7, we calculate the mode projection of the guided mode of the coupling waveguide into various Gaussian beams with different beam waist size. These plots are calculated for the TE mode. Both plots calculated modes for the same initial Silicon Nitride waveguide, but the waveguide for the plot on the right is coated with 10nm conformal coating of Al_2O_3 . We observe a shift in optimal initial Nitride width of $\Delta = -19nm$ as respond to the ALD coating, which agrees with results from fabricated coupling waveguide sweep chips of $\Delta_{measured} = -18nm$.

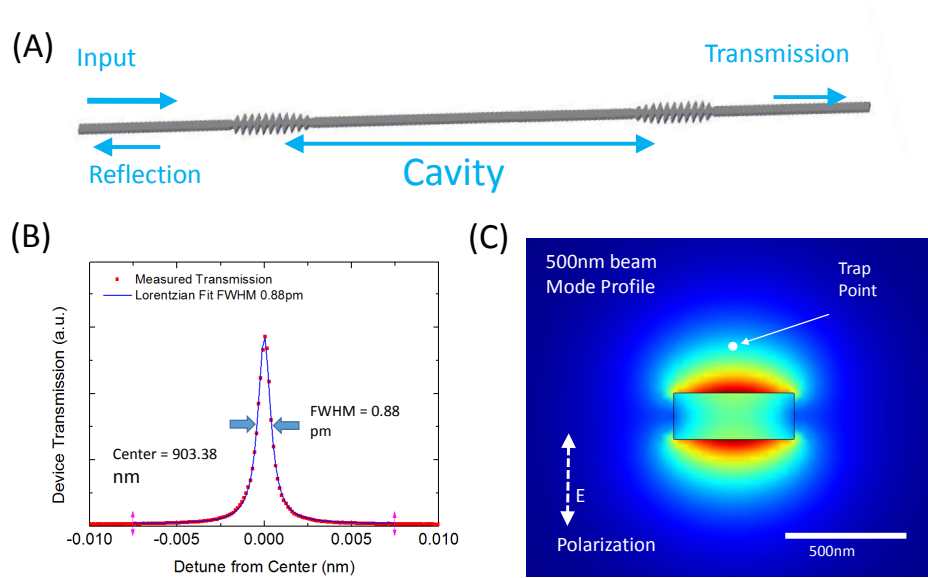


Figure 5.11: Optical properties of the single-beam Fabry-Perot-type cavity

(A) The geometry of the single-beam cavity is formed by cladding two ends of an unpatterned single-nanobeam waveguide with photonic crystal mirrors. These cavities can be measured using reflection and transmission methods. Their geometry also make it potentially possible to stack multiple cavities together on a single waveguide. (B) The transmission of these cavities contain resonance peaks that can be fitted with Lorentzian lineshape to extract their quality factor. For the spectrum plotted here, a quality factor $Q=903\text{nm}/0.88\text{pm}=1.0 \cdot 10^6$ was obtained. With the physical length of $230\mu\text{m}$ of this device, a mirror finesse of $F=1.3 \cdot 10^3$ was obtained. (C) For the TM polarization cavities, the guided mode profile extends suitably far from the top- and bottom-surface of the waveguide, allowing atom-light coupling at reasonable distances.

Chapter 6

ATOM-LIGHT INTERACTION MEASUREMENT

In this chapter, we describe experiments measuring the interaction of Cesium atoms with our devices. A vacuum system is specifically designed for working with the nanophotonic chips with optical fibers. The vacuum system is built with a two-chamber design to protect the devices from Cesium coating from background vapor pressure, where atoms are collected using a MOT in one chamber and then sent into the other chamber with the chip. We localize the atoms near our devices to observe atom-atom interaction mediated by the photonic crystal guided modes.

For reference on the Magneto-Optical Trap (MOT) and Polarization Gradient Cooling (PGC), information can be found in [66, 44, 67]. The 'atom faucet' technique for launching the atom across the two chambers are described in [68]. The vacuum component supplier *Agilent Technologies* provides application notes on vacuum technology in [69].

6.1 Vacuum System Adaption for Photonic Devices

Cold atom experiments place stringent demands of the vacuum system, both in vacuum quality ($< 10^{-10} \text{ Torr}$) and in the optical access for laser beams and imaging. Aside from these typical demands, the addition of nano-photonic structures creates the additional requirements to the atom and vacuum system as follow:

- Limiting Cesium background vapor pressure to prevent coating onto device
- Interfacing chip coupling schemes through vacuum chamber envelope
- Rapid replacement of chips in chamber

Background pressure of Cesium

The alkali atoms frequently used in cold atom experiments are highly reactive chemically, and will bind onto the surfaces of photonic devices. For our devices with both high surface-to-volume ratio and high surface field intensity, the deposition of Cesium would cause rapid frequency drift of the photonic properties, and also increase in device absorption loss. The chemical reactivity and heating associated with increased optical absorption eventually cause our devices to fail. The vacuum

system must then operate at sufficiently low Cesium background pressure to prevent device coating, while providing sufficient loading rate for the various atom traps.

Limitation associated with interfacing chip to outside optics

The vacuum system design must allow the input coupling methods for the chip, be they fiber based or free-space based, to enter the vacuum chamber with good efficiency. For the case of the fiber, this includes using proper feedthrough methods for fiber, avoiding designs that could cause damage to fiber during installation, and dealing with potentially undesirable outgassing properties of the fiber jacket. For the case of the free-space coupling, the primary concerns become the distance to the coupling viewport, and either the optical aberration associated with focusing beams through the chamber window thickness, or developing methods to place lenses and carry out alignment inside the vacuum chamber.

Replacement of chips in vacuum chamber

It is desirable to be able to replace the chip in the experiment apparatus rapidly, especially when our devices have finite lifetime associated with Cesium coating. This demands our vacuum system to be designed in a way to allow frequent venting and opening. It is also preferable to minimize the amount of supporting optics surrounding the chamber that need to be moved during chip replacement process. Both the materials associated with the fibers and the need to frequently open the chamber for chip replacement result in limitation to our ability to bake the vacuum chamber, and hence the ultimate level of vacuum that could be achieved.

6.2 Vacuum System Design and Atom Delivery

We design our vacuum system to comprise of two separated sub-chambers, connected with a tube in between. Each of the chambers is designed with a dedicated pumping setup, including an ion pump and an external port to which turbo and roughing pump can be connected. The three stages of pumps utilize different physical principles of pumping: the roughing pump works by expanding and collapsing of volumes; the turbo pumps use high-speed turbo blades to deflect gas molecules, hence working better in low-pressure regimes where gas molecule dynamics is ballistic; the ion pump uses electrons in high voltage to ionize gas molecules and accelerate them into the pump walls where they are chemically bound or physically embedded. The tube connecting the two chambers can be valved off with an inline valve, so we are able to vent and pump each chamber independently. We will refer

to the two chambers as 'Source' and 'Science' chamber, for one serves to collect and source cold Cs atom, and the other contains the chip. Each chamber has its own set of magnetic field control coils and sets of laser beams for forming a MOT.




Figure 6.1: Illustration of vacuum system

The vacuum system comprises two main chambers, which we refer to as 'Source' (S1) and 'Science' (S2), each with their designated ion pumps (I1 and I2) and ports to connect to external turbo-pump stations (V1 and V2). The 'Source' chamber has a liquid Cesium source attached (Cs), which feeds a MOT that is launched toward the 'Science' chamber by resonant laser (L) through the connecting pipe (P). The 'Science' chamber can be isolated from the 'Source' by closing the inline valve (V3), allowing it to be vented for chip exchange. The chip and fiber assembly is bolted to a rotation-and-translation arm (E) that places the chip in the center of the 'Science' chamber, and the fibers exit the chamber via a set of Teflon feedthroughs mounted to a multiplexer (F) on the 'Science' chamber. Image courtesy of A. Goban.

Dual-chamber design with atom 'faucet'

The 'Source' chamber is equipped with Cesium source, for our case an ampule with liquid Cesium in a sub-compartment that can be valved off to control the equilibrium Cesium pressure in the chamber, or isolate the source entirely. This section of the chamber operates at higher Cesium background pressure ($\sim 10^{-8}$ Torr) to allow quick loading of a large MOT in this chamber. The MOTs in our system comprise three pairs of mutually orthogonal counter-propagating beams that are red-detuned from the Cesium D2 resonance. In the reference frame of a moving atom, the Doppler effect causes the two counter-propagating laser beams in each pair of the MOT to become different in frequency, hence creating a preferential absorption of the beam in the opposite direction of the velocity of the atom. This way, a velocity-dependent optical force is created to slow down the atoms.

We then use the technique developed in [68] to deliver the atom cloud into the connected 'Science' chamber. A resonant beam is aimed at the source MOT to accelerate the atom cloud, which then ballistically travels through the connecting pipe. Another MOT in the 'Science' chamber catches and re-cools the atom cloud. This loading process is repeated many times until the longer-lifetime MOT in 'Science' chamber is fully loaded. The geometry of the pipe results in low gas conductivity and hence prevents the background Cesium vapor from flooding the 'Science' chamber, and hence the chip. The 'Science' chamber MOT loads from the pulses of atom cloud from the source chamber as opposed to its own background pressure. This is verified as when we turn off the resonant 'faucet' beam, the 'Science' chamber MOT will no longer load.

Atom collecting and delivery in the 'Science' chamber

The 'Science' chamber MOT catches the incoming pulses of atom cloud. Then, an additional Polarization-Gradient Cooling (PGC) step is carried out using this MOT setup. The PGC is one mechanism that makes use of the interplay of polarization-dependent optical potential and polarization-selective absorption of light by the atom to cool the atoms beyond the capability of the MOT. The general idea is that the set of beams forming the MOT, due to their vectorial nature, create spatially varying polarization that alternates between circular and linear polarization (hence the name 'polarization gradient'). The local circular polarization creates differential AC Stark shift for the m_f sub-levels of atoms. Since the optical transitions excited by the laser beams are also selected by the local polarization, a configuration can be found where a photon scattering event always maps the atom from a local higher energy m_f (optical-)ground state into a local lower energy one. This process then causes the atom to repeatedly climb the potential hills formed by the local polarization, hence cooling the atoms. This method selects atoms by their m_f levels, avoiding the natural linewidth limit of the optical transition that limits the MOT.

We need to compress the MOT cloud size to bring the few-millimeter scale cloud down to the order of few-hundred- μm , so to allow it be placed into the through-window of the chip. This can be done by increasing the beam detuning and magnetic field gradient, as the size of a MOT cloud is determined also by the intra-cloud light pressure (decreases with scattered light by atom, hence the use of larger detuning) and the confinement forces (which increases with the magnetic field gradient). The compressed atom cloud can then be accelerated by abruptly changing the position zero-magnetic-field point in the 'Science' chamber, causing the cloud to move

toward the new magnetic zero point, and then turn off the MOT laser beams once the desired velocity is reached. This 'throwing' technique is necessitated by the fact that the cloud will pass through the shadow of the chip substrate. If one side of the laser beams is obscured by the chip but the opposite propagating beam remains, the unbalanced light pressure will crash the cloud into the chip substrate.

A third set of millimeter-diameter laser beams form a small MOT inside the chip through-window to catch the thrown atoms ('mini MOT'). The cloud is decelerated and captured by the small MOT and confined in the vicinity of the devices. We note that due to the limit of atom scattering rate of photons, the force we can apply to decelerate the cloud, which controls the distance required to stop, is limited. This distance scale imposes a limit on the minimum size of the through-window of our chip. At this point, further cooling can be carried out, or a FORT trap can be loaded from this cloud. Care must be taken to avoid clipping of the MOT beams by the chip through-window edge, as scattering from the edges can locally influence the MOT cloud and cause it to move in an unpredictable manner.

6.3 Methods for Measurement of Atom-Light Interaction

The majority of our measurements made with atoms in the vicinity of devices have been done through the guided mode of the devices. We can launch near-resonant lasers through the device guided modes to measure the optical responses from the atoms, carrying out spectroscopy of the atoms to observe the influences they are under, such as broadened linewidth or resonance frequency shift. Alternatively, direct observation of atom emission of light can be achieved using a decay rate measurement in time domain.

Frequency domain spectroscopy measurement

We carry out spectroscopy for the atoms near the photonic device by sending in weak probe beams of varying frequencies, generated by diode lasers locked to Cs vapor cells, then shifted in frequency by using Acousto-Optic Modulators (AOM), where RF phonons are pumped into a solid material and scattered the photons of the laser light. First feature to look for in these measurements is the total linear absorption or reflection of the atom. To account for the efficiencies of optical coupling and the responses of the device itself, the measurements are normalized by the measurement values of the same device and frequency but without atoms, with a pulse immediately (as determined by mini-MOT lifetime) after the atom measurement concludes. We optimize our various trap loading and manipulation parameters by maximizing the




Figure 6.2: Illustration of Experiment Cycle

(A) A MOT is formed from collecting Cesium atom from the high vapor pressure in the 'Source' chamber. (B) A near resonant beam accelerates and launches the atoms from the MOT, done multiple times at rapid repetition. (C) Atom cloud ballistically flies through a thin connecting pipe, whose gas conductivity limit maintains the pressure difference between the two chambers ('Source' $\sim 10^{-8}$ Torr, 'Science' $\sim 3 \cdot 10^{-9}$ Torr). (D) A MOT is formed in the 'Science' chamber from collecting the flying atom clouds. Further cooling stages are carried out at this point. (E) The MOT cloud is compressed in volume. (F) The compressed cloud 'throw' that goes across the shadow of the edge of the chip. (G) A third 'mini-MOT' is formed by catching the thrown cloud inside the chip through-window. (H) Optical traps near the device, such as the side-illumination trap as shown here, can be loaded from the in-window MOT.

atom absorption or reflection. The resonance frequency shift of the atom spectrum can provide information regarding the trap field intensity at the position of the atom. This is particularly helpful for our experiment, due to the difficulty to directly resolve the location of the atoms by imaging. The shifting of resonance frequencies is compared with the calculated local field pattern created to infer the position of the atoms.

Time-domain decay rate measurement

The time-domain method makes use of short pulsed excitation to excite the atoms. Then, after the excitation pulse has decayed away, we record the spontaneous decay of the atoms via the waveguide channels. The importance of this method is to allow us to disentangle spectral broadening of the atom resonance from enhanced decay rate versus the inhomogeneous broadening of atoms due to light shifts from

the local environment. Attempting this measurement is more challenging than the spectroscopy, as the atom decay lifetime is on the order of tens of nano-second. The extinction of the excitation pulse, and the time resolution for the electronics of the detector must be able to operate at this time-scale. In our experiment, a cascade of AOM with an Electro-Optical Modulator (EOM) was used to produce the sharp cut-off of the excitation pulse. The EOM modulate the phase of transmitted light by applying electric field to the nonlinear material light passes through. They can typically reach >GHz bandwidth, comparing to the hundred-MHz range of AOM.

6.4 Observation of Enhanced Single-Atom Decay Rate

Our first experiment of atom-light interaction near photonic crystal (published in [60]) is done by immersing the photonic waveguide directly in a MOT atom cloud, then carry out spectroscopy of the device. This experiment uses the Cesium D1 line as probe mode, and a weak ($<1\mu\text{W}$ input power) blue-detuned beam on the D2 line was sent in as optical 'guiding' beam to help control placement of the atoms. The reflected signal is normalized in real-time by taking a reference reflectivity measurement 10ms later than the atom measurement, when the atom cloud had dispersed. We manipulate the system parameters such as the loading MOT cloud density, and the guided mode powers for trap and probe mode for our devices, to extract information regarding the atoms.

Results

With this configuration, the overall signal level is up to a 25% change in device reflection from the optical response of the atoms. We fit the signal using transfer matrix model treating atoms as frequency-dependent reflectors as in [49]. The fitting result for the mean atom number \bar{N} , assuming Poissonian distribution, was found to be 1.0 ± 0.1 , and the interaction strength Γ_{1D}/Γ' to be 0.35 ± 0.1 . From the interaction strength value, we inferred an effective mode area of $A_{eff} \simeq 0.44\mu\text{m}^2$. We then estimated the location of the atoms by comparing the effective area value to our guided mode profile, and inferred a position range of $\pm 130\text{nm}$ range from the device plane on its center mirror-plane; see Figure 6.3. The level of Γ_{1D}/Γ' allowed us to observe preliminary results of collective behavior of the atom in the form of increased saturation power with atom cloud density.




Figure 6.3: Simulated distribution of Cesium atom near device

(A) Monte-Carlo simulation of atomic trajectories gives the atom distribution near the device with blue-detuned guided mode guiding potential. We can see the depletion of atoms near the device due to the blue detuned guided mode. A stream of atom flows into the gap between the two beams of our device. (B) The atom distribution weighted by their local coupling efficiency, $\xi(\vec{r}) = \rho(\vec{r}) \cdot \Gamma_{1D}(\vec{r})/\Gamma_{1D}(\vec{0})$. Image courtesy of A. Goban [60].

Discussion

This experiment is our first demonstration that our system can observe atom-light interaction and preliminary light-mediated atom-atom interaction. Measurement from this generation of chips are more difficult to interpret since they had only one port, hence only reflection measurement is available. Additionally, they were made prior to the development and adaption of anti-reflection structure and fiber, further complicating the reflection with interference fringes between reflected signal and residual reflections on the optical path. This generation of devices was also prior to the adaption of cooling tether arrays, hence severely limiting the amount of trap power available to us. Optically, the guiding of atoms using the blue-detuned mode as designed in Chapter 5 is partially helpful in localizing atoms to the desired location, as suggested by the effective area estimation. The simulation for atom trajectory and the resulting density suggested, however, that the solid angle from which an atom would be able to arrive at the center of the device is very small, limiting our loading efficiency.

6.5 Observation of Light-Mediated Atom-Atom Interaction

In this section of experiments, we make use of a side-illumination trapping scheme with trap site above the device plane. This trap is easier to load, allowing us to reach higher number of atoms in the vicinity of our devices. The atoms are also better localized by the local field maximum, provided by a strong external beam. On the device side, the newer generation chip allows measurement in the

transmission channel for cleaner signal. We carried out time-domain decay-rate measurements that allowed us to unambiguously determine the enhanced decay rate without complex model fitting.

Result: Superradiance

In our superradiance experiment, we control the atom loading time to vary atom numbers in our trap, while maintaining all other parameters. The pulse pump-and-free-decay measurement method as described in Section 6.3 is used. The number of atoms trapped near the device can be reduced in a controllable manner by either reducing the loading time of the MOT, or by allowing the atom number to naturally decay away from its trap lifetime for some fixed time before initiating the measurement sequences. At low ($\bar{N} < 1$) trap average atom number, the probability of having multiple atoms simultaneously is low, and the decay rate asymptotes to that of a single atom. In this experiment, the single-atom decay rate was measured to be $\Gamma_{total} = \Gamma_{1D} + \Gamma' = 2.0 \pm 0.1 \Gamma_0$, where Γ_0 stands for the free-space atom decay-rate. The average atom number for each cloud density is then determined by fitting with a Poissonian distribution of atom number, using the $\Gamma_{1D, single-atom}$ and Γ' , obtained by single atom measurement and FDTD calculation, respectively. A maximum decay rate of $\Gamma_{total} = 3.0$ at a mean number of $\bar{N} = 2.6$ was achieved.

For comparison, the spectroscopy measurement was also carried out in this experiment, providing us with frequency domain information regarding the atom decay in the form of the linewidth of its resonance. At low atom cloud density, the spectroscopy yielded a single-atom linewidth of $\Gamma_{single-atom}^{Freq} = 2.1 \pm 0.1 \Gamma_0$, in reasonable agreement with the time-domain measurement of single-atom Γ_{total} . Maximum cloud density yielded a 70% transmission dip at maximum cloud density. Fitting the spectrum at this setting produced a widened linewidth of $\Gamma_{superradiant}^{Freq} = 3.4 \Gamma_0$.

Result: Band-gap mediated interaction

Useful information can be extracted from the lineshape of an atomic spectrum. In the previous sections, the discussion was focused on the atom response when its frequency is on-resonance with the device super-modes. In such configuration, the influence of the super-mode cavity compounds with the slow group velocity effect to create strongest enhancement effect. When the atom response is on resonance in such configuration, the influence from the device is purely dissipative, producing a symmetric absorption profile as in Figure 6.4 (C). Generally, at frequencies that

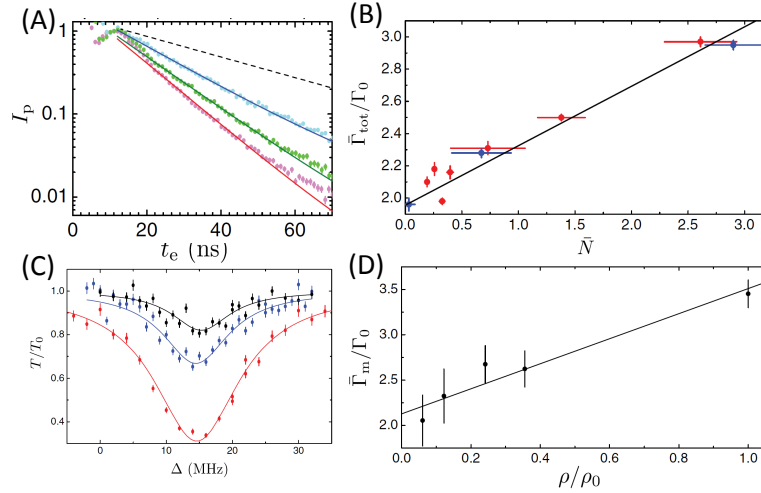


Figure 6.4: Measurements of superradiance in number-dependent decay

(A) The lifetime of the FORT allows controlled reduction of atom cloud density by delaying the measurement start time (red: 3ms, green: 13ms, and blue: 63ms). The dashed line is the unmodified free-space decay rate of the atom as reference. Empirical exponential fit to these decay curves showed the atom decay-rate to be higher for the short delay time (higher atom cloud density) than the longer delay cases. (B) Linear increase of atom total decay rate as function of fitted average atom number. (C) Increased absorption level and spectral broadening of the atom resonance as the atom cloud density increases. The traces are: red: maximum density ρ_0 ; blue: $0.24 \rho_0$; black: $0.12 \rho_0$. (D) Fitting of the resonances with Lorentzian lineshape produced consistent result with the time-domain measurement. The slightly higher single-atom decay rate may have resulted from inhomogeneous broadening. Image courtesy of A. Goban [60].

are not on the cavity resonances, the optical response of the device also contains a coherent interaction term. The lineshape of atom resonance under both coherent and dissipative effect can be described by the equation from [63]:

$$\frac{t(\delta)}{t_0(\delta)} = \frac{\delta + i\Gamma'/2}{(\delta + \sum_{m=1}^N J_{1D}^{mm}(\Delta)) + i(\Gamma' + \sum_{m=1}^N \Gamma_{1D}^{mm}(\Delta))/2}$$

where δ is the detuning between the probe frequency and the atom's natural resonance frequency without the device (but could include AC Stark shift from trapping), Δ the detuning between the atom frequency and the device band-edge, $J_{1D}^{mm}(\Delta)$, $\Gamma_{1D}^{mm}(\Delta)$ are the real and imaginary parts of the self-interaction of the Green's function of the device, calculated at the position of each atom. To explore this detuning dependence, we make use of an external heating beam directed at the Silicon substrate of the chip to thermally tune the device frequency. Spectrum data are

taken varying the device detuning Δ , from the band-edge resonances to deeper into the band-gap. The band-gap suppresses the atom decay channel into the guided mode, causing value of Γ_{1D} to vanish. The real part J_{1D} , however, remains finite for a larger frequency range into the band-gap. See Figure 6.5.

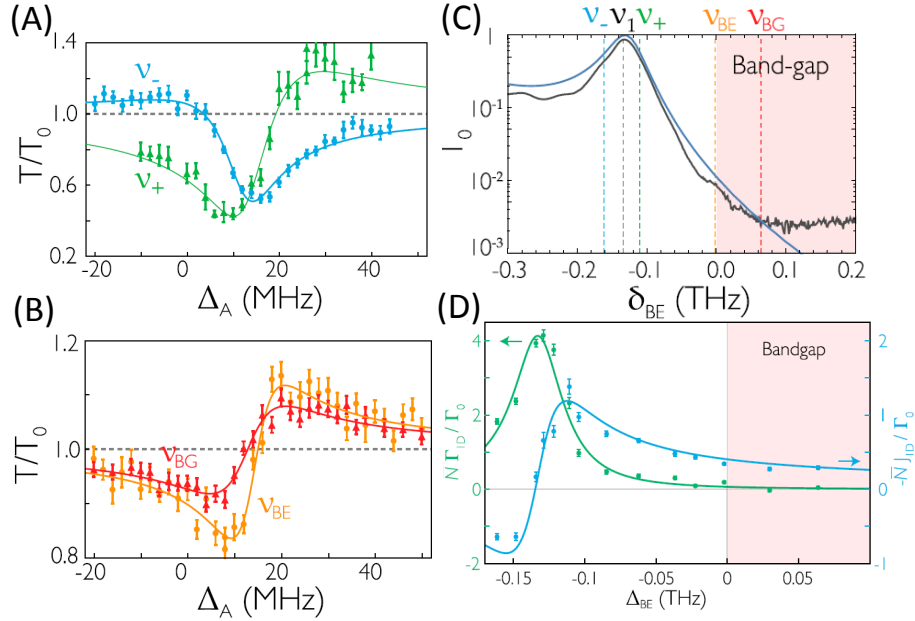


Figure 6.5: Measurements of Off-Resonance and Band-Gap Spectra

(A) As the atom frequency offsets from the device band-edge resonances, the spectral shapes become asymmetric. (B) Due to the finite length of our photonic crystal, a limited amount of light can still penetrate the crystal when the atom frequency moves into the band-gap, and a transmission spectrum can be taken. (D) Fitting the lineshapes allows us to determine the coherent (J_{1D}) and dissipative (Γ_{1D}) components of light-atom interaction. We note that the coherent component (blue) persists deeper into the band-gap frequency range. Image courtesy of J. D. Hood[63].

Discussion

We were able to load larger average atom number with the side-illumination trap, hence allowing observation of atom-atom interaction in absorptive and dispersive regimes, including inside the band-gap of the APCW. The improved band alignment to Cesium with the newer generation of devices allowed us to benefit from larger slow-group-velocity enhancement, partially mitigating the reduction of coupling strength associated with the larger effect mode area at the location of the side-illumination trap. The single-atom coupling rate to the device is still roughly equal to the residual uncollected decay rate into free-space or other channels $\Gamma_{1D} \approx \Gamma'$.

The next steps of development on the APCW devices would be to design schemes to load the side-illumination trap sites continuously into the device vacuum gap. Methods using a moving optical lattice or using gradually ramping up of a guided mode red-detuned trap are being developed. We hope that efficient loading of the guided mode traps can be achieved in the near future.

Chapter 7

ENGINEERING OPTICAL PROPERTIES WITH 2D PHOTONIC CRYSTAL SLABS

This chapter presents simulation and preliminary fabrication results of photonic structures based on a 2D photonic crystal slab. We found that our support structure developed in Chap. 2 can be adapted to support these larger devices with some minor modification. We calculated that properties such as suppression of non-collected channel decay rate Γ' can be achieved in these 2D-based structures.

For reference, the book [70] provides information on crystal and reciprocal lattices and their symmetries; also see [71]. The book [53] contain sections regarding 2D photonic crystals. Reference for projecting an 2D Brillouin zone into an 1D case can be found in [72].

7.1 Adaptation of Support Structures to Photonic Crystal Slabs

The majority of our chip layout and support structure design were developed for the 1D APCW devices. Nevertheless, the same chip layout and support structures, such as the through-window and the input coupling waveguides, are capable of supporting larger suspended structures like the 2D photonic crystal slabs. In some cases, additional mechanical support structures are required to redistribute the tensile stress in an even manner onto the larger structures to prevent damage; see Figure 7.1. The adjusted designs can typically achieve yield rate comparable to the 1D waveguides devices, despite their larger surface area and higher complexity.

Mechanical deformation of the 2D structure

Issues with mechanical stress of the Silicon Nitride material remains for 2D devices to a certain degree. The tensile force from input waveguides and various tethering structures could cause deformation to the 2D structures they support. For example, we found the mechanical gap in the 1D-embedded slot waveguide devices described in Section 7.4 to be sensitive to the way its support tethers anchor to the Silicon substrate. Keeping the anchor tether tensile was found to increase the gap size up to $\sim 100\text{nm}$ for a 350nm designed gap size. Calibration of mechanical gap must be done in a consistent manner in term of support and tethering schemes for these structures.

Coupling 2D structures to input waveguide

Coupling the input waveguides to the 2D photonic crystal slabs turned out to be relatively straight-forward. Empirically we found that a linear tapering of the typical 500nm nanobeam waveguide at a 1:10 slope is sufficient to widen the mode profile up to $10\mu\text{m}$ to directly feed into the 2D photonic crystal modes. A 2D system differs from the 1D system in that light can be injected into the crystal from a continuum of directions relative to the lattice vectors. Practically this refers to changing the relative angle between the input waveguide and the crystal lattice axes, either tilting the waveguide, or rotating the 2D lattice. We opted for the later to maintain the waveguide and mechanical support design for structural consideration. The crystal lattice vectors of the 2D crystal, and the shapes of the holes (if not circular) are rotated to test the optical response to different directions.

7.2 Optical Properties of 2D PhC Modes

We build our 2D periodic structure using the same single-layer Silicon Nitride platform. The majority of research reported in this chapter is done using a hexagonal lattice of holes (or, equivalently, a honeycomb of tethers) through a suspended slab of Silicon Nitride. The hexagonal lattice was selected over the other 2D lattices such as square lattice for the closer length of its two reciprocal lattice vectors, and hence the ease of opening a 2D band-gap. The ratio of the two high-symmetry direction reciprocal vectors of the hexagonal lattice is $2:\sqrt{3}=1.15$, closer to one than the square lattice's $\sqrt{2}:1=1.41$. The effective mode index of the larger reciprocal vector mode must overcome this ratio to open a band-gap. The geometry of the crystal is chosen to be a web of interconnected tethers forming a honey-comb; see Figure 7.2. For studies of the 2D system with square lattice, please see the thesis [73].

TE bandgap and free-space decay suppression

The first property to design for is 2D photonic band-gap. The 2D gap spans between the lower 'dielectric' mode of the K point and the 'air' mode of the M point of the Brillouin zone. With our geometry and device layer thickness, we have so far only managed to form a TE band-gap. The effective index of the TM mode for even a completely filled structure of 200nm thickness Silicon Nitride is only 1.27 (compared to TE where >1.5 can be reached), hence making it difficult to overcome the ratio between the length of the K and M reciprocal lattice vectors. Regardless, we should be able to observe strong effect on the atom decay rate with only the TE

band-gap.

A FDTD Green's function calculation demonstrates the primary features of such system; see Figure 7.3. The simulation assumed that a dipole emitter is placed in the center of one hole in the 2D lattice. We can see strong suppression to both of the dipole directions with dipole in the device plane in the photonic band-gap. The TM mode does not demonstrate a band-gap behavior like the TE modes, which makes it suitable to provide long-range interaction and probing channels.

Dual-frequency directional emission

Similar to the 1D case, it is possible to build structures that simultaneously align two band-edges to distinct atom resonance lines. For the 2D case, this would mean aligning the K point dielectric mode to the Cs D1 line, and the M point air mode to D2. We can observe in the Green's function plots that resonances associated with the finite size structure arise near the band-edges, similar to the 1D case but with more complicated 2D super-mode profiles. The more interesting feature of the 2D system is that frequency ranges can be found where the two atomic resonance lines produce different directional emission, as determined by the K and M directions [74]; see Figure 7.4.

Directional emission of a dipole emitter can be achieved by operating the crystal at frequencies with near-parallel group velocity vectors $\vec{v}_g(\omega, \vec{k}) = \vec{\nabla}_k \cdot \omega(\vec{k})$, namely, where the equal-frequency curves are locally flat. These regions generally exist on the K-direction for the dielectric band, due to the reasoning as follow: the dielectric band dispersion manifold begins at the Γ point, and increases toward its maxima at the K points; If we calculate the projection of the local curvature of the equal-frequency curves on the line $\Gamma - K$, we see it goes from negative (curving toward Γ point) to positive (curving toward K point); assuming the dispersion relation of the crystal is continuous, there would be some point on the $\Gamma - K$ line where the curvature is zero; This generates the locally flat region. A similar argument can be made for the air band, where the frequency maximum is now at the Γ point, and minima at the M points. More information can be found in the thesis [73]. This may allow for more diverse control of optical interactions of atoms trapped in the vicinity of the 2D photonic crystal, such as individually engineering nearest-neighbor and second-nearest-neighbor interactions for the D1 and D2 transitions.

Circular modes at K point air band

The air band of the K point has a 2-fold degeneracy. The hexagonal lattice has a natural 6-fold rotation symmetry. Introducing a particular \vec{k} vector on the K point breaks this symmetry down to 3-fold rotation symmetry. We see that there are two choices of K-direction vectors that this can be done, hence creating the degeneracy; see Figure 7.5. With suitable perturbation of the shape of the hole (which maintains the 3-fold rotation symmetry), this degeneracy can be lifted; see [71]. The two modes hence generated have the property that at any point where the dielectric distribution respects the 3-fold rotation symmetry, the resulting eigenmode cannot have linear polarization. A direction defined by linear polarization would violate the 3-fold rotation symmetry around this point. For the two modes on the air band of the K point, this is of particular interest, since one of the 3-fold symmetry points is the center of the hole, where the two modes have intensity maximum and opposite circular polarization. Operation of these modes would then allow coupling of atom m_f sub-levels with decay rate and emission directions. With addition of external beams or magnetic field, non-reciprocal propagation of light between atoms may be achieved. The physics of directional emission and potential chiral optical systems are being actively studied [75, 64, 65].

Atom trapping in the 2D lattice hole

We have calculated the local field patterns formed by the 2D photonic crystal structure, under illumination of an external beam incoming from the perpendicular direction to the device plane; see Figure 7.6. Interestingly, a local minimum is formed within the device layer in this condition. See Figure 7.6 panel (H). This would allow us to form a blue-detuned FORT to localize atoms inside the holes of the 2D crystal. By using a pair of counter-propagating beams with a small difference in frequency, it was possible to move the location of this local minimum continuously through the device layer.

We believe this would allow an efficient way of loading atoms, first into the minima of a blue-detuned standing wave trap, move the atoms into the device layer, then by ramping down one of the two lattice beams in a controllable manner to transfer the atom into the localized trap supported by only one beam. Loading a blue-detuned 1D optical lattice is more difficult than a typical red-detuned trap, but a blue-detuned beam is favored for trapping near dielectric structures, for its capability to prevent atoms crashing onto the dielectric surfaces. For more information on blue-detuned traps and moving lattices, please refer to the thesis [73], and also the article [56] and

its supplemental information.

7.3 Optical Properties of 2D PhC with Defects

We explore features that can be created by local modification of the 2D photonic crystal lattice. Here we present two such applications of different dimensions. We found that a local modification of one hole and its first set of nearest neighbors on the 2D lattice can further enhance the decay suppression capability of the base 2D lattice. We also develop an embedded 1D 'slot' waveguide structure that can achieve not only the optical enhancement properties of the APCW, but also suppress the uncollected scattering channels.

Point-defect super-suppressor

In a 2D photonic crystal system, the in-plane dipole emission is suppressed by the photonic crystal due to its band-gap property, but the dipole is still able to radiate out of plane into free space. If the dielectric structure around the radiating dipole can be polarized by the dipole to produce a 'image dipole' of the source, the radiation of the whole system would be approximately the coherent sum of the two. If the image dipole is designed in a way that radiates out-of-plane at equal strength but opposite phase of the source, then the out of plane radiation can potentially be suppressed.

With this idea in mind, we numerically experiment with local modification of the size of one hole, and the positioning of the surrounding nearest neighbor holes, while maintaining the rotation symmetries of the lattice. Optimizing these two parameters in FDTD simulation, we observed indeed a ~ 10 THz bandwidth dip of emission of the in-plane dipole orientations, with its center frequency adjustable by the perturbation parameters. By placing the frequency of this dip in the center of a 2D photonic bandgap, a total TE dipole emission rate suppression of 22.7 dB below free-space value can be achieved; see Figure 7.7. We hope that atoms can be placed into these special sites using technology such as the optical tweezers, and observe strongly suppressed decay rate.

1D defect slot waveguide

Local 1D modifications on a 2D photonic crystal can create embedded waveguides. This modification is typically done by removal of few number of rows of holes, but slot-like modification are also frequently explored in literature to benefit from the strong field in such slots, such as in [76]. The challenge for our experiment to adapt this technology is in whether a 1D defect can be formed to provide the set

of properties needed to interact with atoms in a similar manner to our previously developed APCW structure. We found that a defect structure, characterized by three parameters as in Figure 7.8, allows a vacuum gap to trap atoms, and is able to achieve dual-band alignment to atomic resonances. The mode profiles of the guided modes also exhibits the periodicity and contrast as in the case for APCW that we hope will eventually allow guided mode trapping of similar design as well.

7.4 Atom-Light Interaction Parameters of the Slot Waveguide

We proceed to develop the embedded 1D slot waveguide and evaluate its potential for atom-light interaction. There are then two more issues to be resolved for this structure beside band alignment. Firstly, a tapering scheme must be developed for transitioning from unpatterend nanobeams to the embedded structure. For the APCW structure, this transition can be made intuitively by reducing the modulation amplitude continuously to zero. For the defect waveguide embedded in a 2D photonic crystal, any tapering transition must be done in a way that minimize interference with the background 2D band-gap. Secondly, to achieve the goal of observing suppressed decay rate below free-space level, we must eliminate all other guided modes from the frequency in consideration, such as the anti-symmetric mode shown in Figure 7.8.

Tapering Into the Slot Waveguide

We make use of the design as described in [59], adiabatically approaching the 2D photonic crystal from outside the unpatterend beams' evanescent field. The shape of the edges of the approaching 2D crystal is made in a way that it forms the geometry of the nominal crystal once the tapering is completed; see Figure 7.9. We found that this tapering scheme can indeed adiabatically convert the unpatterend waveguide mode into the photonic crystal modes, even for the higher frequency of the two (H mode in Figure 7.8) that has a longitudinal profile with multiple nodal plains. Our FDTD calculation for this tapering indicated that a round-trip efficiency above 85% is achievable for this configuration. We found that with our stress distribution tethers and mechanical support frame design as described in 7.1, these tapering structures can be fabricated at high yield rate. For an SEM image of the tapering section of a released device, please see Figure 7.10.

Anti-symmetric mode

An anti-symmetric mode exists between the two symmetric bands, similar to the APCW; see 'N' mode in Figure 7.8. To obtain an intuitive picture for the properties of this mode, we consider the following: Consider a 2D photonic crystal slab that fills an infinite half-plane. Due to the band-gap property of the 2D crystal, a half-plane with vacuum on the other half space would guide a mode on its edge. The guided modes of our slot waveguide can be imagined as hybridization of the edge-guided modes of the left and right 2D half-planes. The frequency separation between the symmetric and anti-symmetric mode formed this way is then determined by the strength of hybridization, predominately controlled the distance between the two 2D crystal half-planes, namely, the size of the vacuum gap. We now have three bands (the two symmetric modes are named 'A' and 'H' modes, and the anti-symmetric 'N' mode) that are tuned by three geometry parameters (g , w , l). With some iterations of numerical calculation of the bands, we were able to find ranges of parameters where the higher frequency symmetric mode in question can be sufficiently removed from the anti-symmetric mode. This configuration should still allow us to observe suppressed decay rate below free space for the Cesium D2 line.

Green's function and state-dependent atom response

We compute the imaginary part of the classical Green's function for a full slot waveguide device with outgoing tapers; see Figure 7.11. While the maximum enhancement is a function of the length of the crystal, which is ultimately limited by optical loss and disorder, the band-gap decay suppression can be characterized. The dipole source output power in the photonic crystal is calculated and plotted as an enhancement ratio over the case of dipole in free space. The imaginary part of the Green's function is related to the output power by the simulated dipole source by the formula as in [58]:

$$\gamma_{tot} = \frac{2\mu_0\omega^2}{\hbar} \text{Im}(d^\dagger \cdot G(\vec{r}, \vec{r}, \omega) \cdot d)$$

The values for the components of Green's function can then be converted into atom decay rate by the formula as presented in [77]. We find that the average decay rate can be suppressed to 0.75 of its free space value, and this suppression is stronger for the stretched states, $F = 5, m_f = 5$ to $F' = 4, m_f = 4$ on the D2 line, to 0.63 of its free-space value. This is advantageous since the D2 line $F = 5, m_f = 5$ to $F' = 4, m'_f = 4$ transition is a closed transition. We hope to deploy this slot waveguide design to our atom experiment in the near future.

7.5 Characterization and Tuning of 2D-Based Devices

The 2D and embedded slot waveguide devices are connected to 1D waveguides using the same support structures. The measurement and analysis methods are similar to that of the APCW described in Chapter 5, except for the K modes of the 2D crystals. The K mode does not efficiently reflect back into its feed waveguide, likely due to mode profile mismatch, and can only be measured in transmission. These devices can also be ALD-tuned. Care must be taken during ALD tuning the slot waveguide devices, since the ALD coating simultaneously tunes the slot waveguide modes, the non-fiber-coupled anti-symmetric mode, and the background 2D crystal. The later two cannot be directly measured using transmission or reflection on the slot waveguide devices. Separated chips with 2D crystals of either crystalline orientation are fabricated and ALD coated to measure their tuning rates, and the anti-symmetric mode is estimated based on simulation.

Geometry		1D			2D	
		'A' Mode	'H' Mode	'N' Mode	M-Air	K-Dielectric
Lattice	Calc	-0.54	-0.48	-0.46	-0.46	-0.53
	Measured	-0.63	-0.57		-0.44	-0.54
Tether	Calc	-0.16	-0.29	-0.25	-0.68	-0.31
	Measured	-0.14	-0.25		-0.71	-0.27
w	Calc	-0.26	-0.19	-0.32		
	Measured	-0.24	-0.16			
l	Calc	-0.005	-0.095	0.0		
	Measured	-0.005	-0.089			
gap	Calc	0.026	0.015	-0.024		
	Measured	0.019	0.010			
ALD	Calc	-1.00	-1.00	-1.16	-1.52	-0.94
	Measured	-0.83	-0.86		-1.14	

All units in THz/nm.

Table 7.1: Geometric tuning rates for the 2D and slot waveguide devices




Figure 7.1: Mechanical support frame and tension redistribution tethering

(A) Full view of 1D slot waveguide device embedded in two slabs of 2D photonic crystal, including the tapering between crystal and waveguide, and also mechanical support frame. (B) Proper connection scheme between the waveguide and the mechanical frame are needed to prevent breaking at the junctions between the high-stress waveguide and the stress-relieved frame. (C) The design of the stress redistribution tethers junction with the primary waveguide. This design is based on adiabatic mode conversion, which prevents the guided mode from leaking into the tension redistribution tethers. The fundamental mode of the primary waveguide (i) would temporarily spread across the width of the dielectric material as the tethers split (ii), but adiabatically return to the wider center beam once the main waveguide mode and the tethers decouple (iii). Note that the color plots are the fundamental TE modes of the waveguides, with 500nm center beam and 200nm tethers. The gap size between the main beam and the tethers for (ii) and (iii) are 25nm and 100nm, respectively. (D) The current design partially redistributes the mechanical stress using side tethers. The input optical mode penetrates the frame using the tether array design.

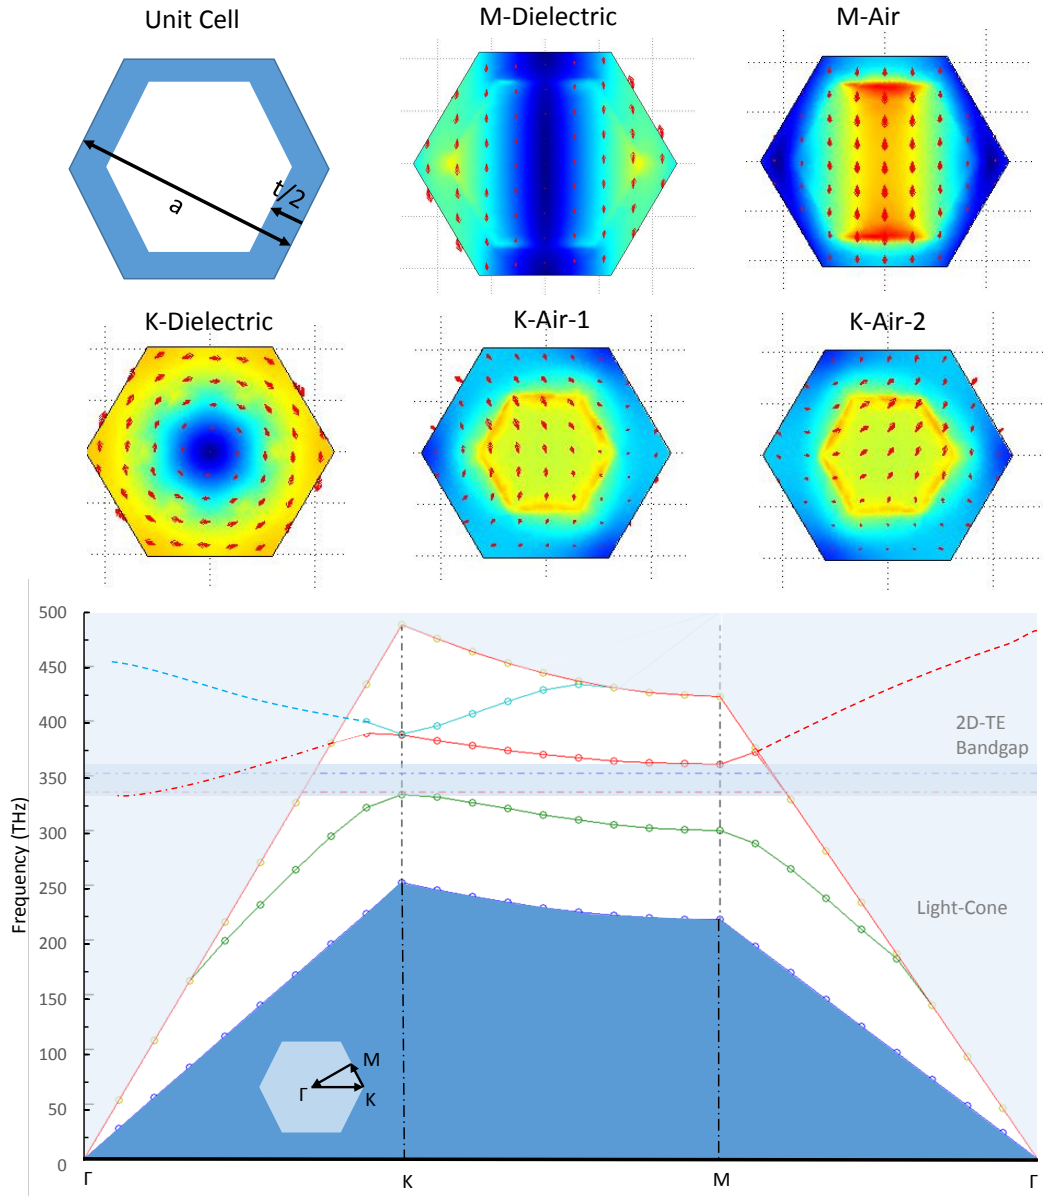


Figure 7.2: Band structure and mode profile for TE modes of 2D photonic crystal

The geometry of the 2D crystal is characterized by two parameters, the lattice constant a , and the material 'tether' thickness t , as shown on the upper left panel. The band frequencies are calculated using FEM calculation for the high symmetry directions, $\Gamma \Rightarrow K \Rightarrow M \Rightarrow \Gamma$ in the 2D Brillouin zone. The 2D photonic band-gap exists between the lower band of K point and the higher band of M point. The mode profile on the high symmetry points, including the two-fold degenerate upper K modes that are separated in a way that respects the 3-fold rotation symmetry, are also plotted. The two horizontal dashed lines on the lower panel plot indicate the Cesium resonances.




Figure 7.3: Dipole emission suppression with uniform 2D photonic crystal

This figure is the Green's function enhancement plot for a 2D photonic crystal of parameter $a = 410\text{nm}$, $t=180\text{nm}$, with the dipole emitter placed in the center of a hole. We observe up to 8dB suppression of the two TE dipole orientations (blue and green traces) in the band-gap frequencies from 320THz to 355THz. The TM mode (red) does not demonstrate a band-gap. We also note that the asymmetry of the 320THz band-edge resonances being lower than the 355THz edge is due to the mode profile of the higher band having local maximum at the dipole position.

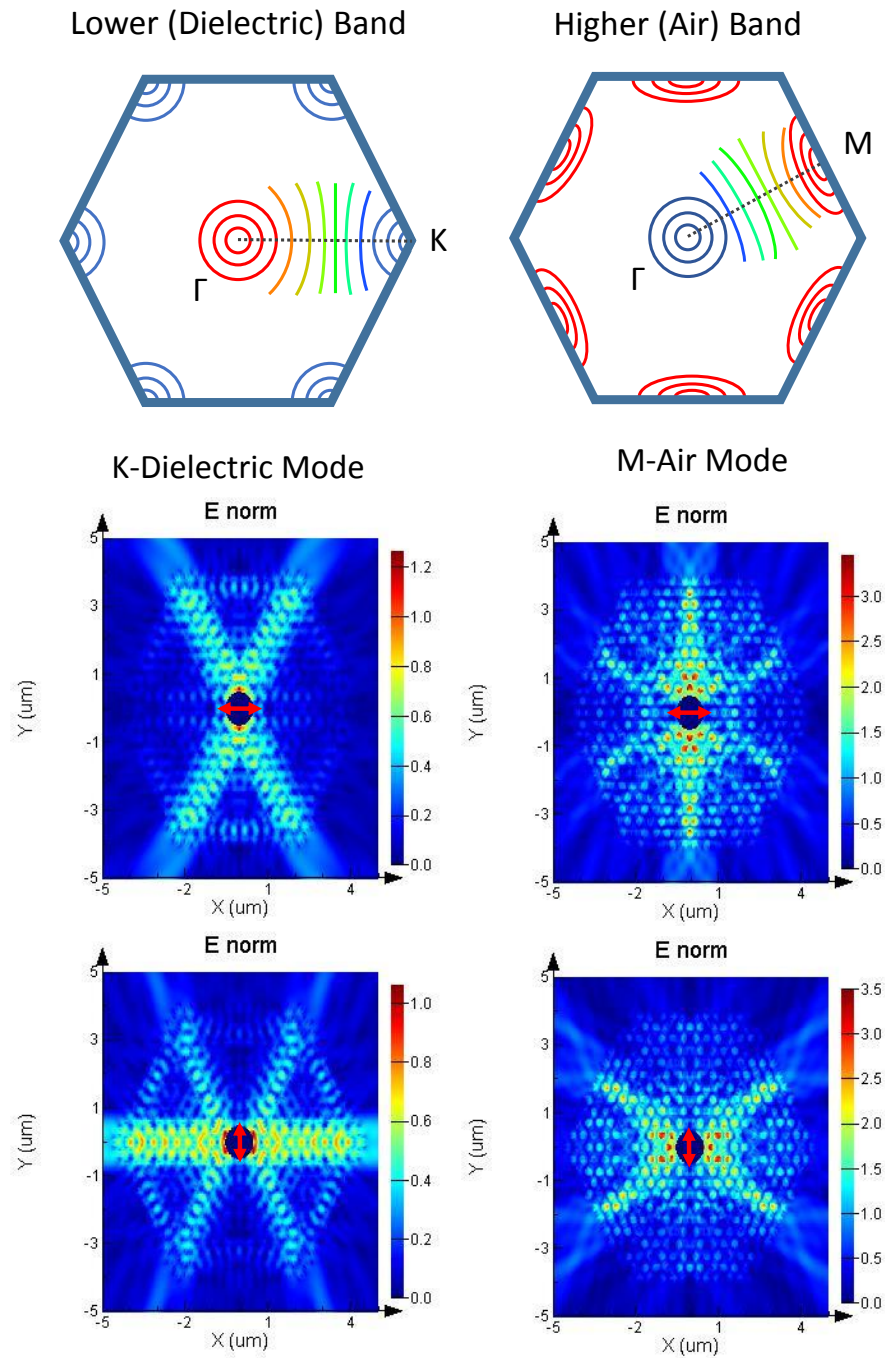


Figure 7.4: Directional emission guided by photonic crystal dispersion

Left: In the lower (dielectric) band manifold of the 2D Brillouin zone, we find some frequency range where the in-plane dipole emission propagates along the K directions of the crystal, with different efficiency toward the three K directions as determined by the dipole polarization (red arrows). Right: Similar behavior is demonstrated for range of frequencies above the 2D band-gap in the air band manifold, propagating primarily toward the M direction of the lattice.

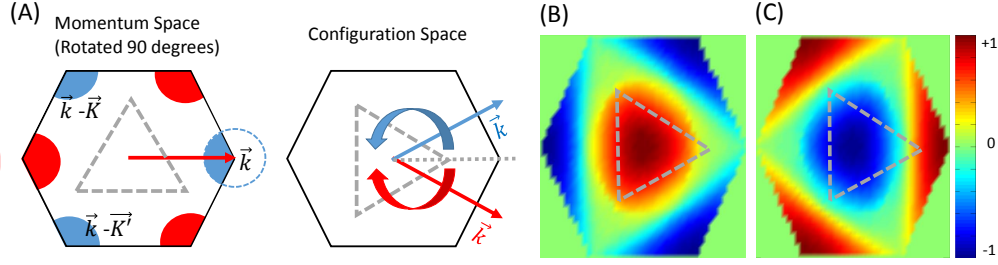


Figure 7.5: Circular polarization of the K-air modes

(A) A K-point mode is composed by equal components of all three \vec{k} values, linked by the lattice momenta \vec{K}, \vec{K}' , where $\vec{K} \cdot \vec{a}_1 = \vec{K}' \cdot \vec{a}_2 = 2\pi$. There are two sets of such modes that cannot be rotated into each other by the 3-fold rotation symmetry, as depicted by blue and red circles. Introducing a suitable perturbation that breaks the K-direction symmetry, as illustrated by the gray dashed triangle, lifts the degeneracy between the two modes. In configuration space, we observe that the circular polarization orientation can be determined by the outer product of the \vec{k} vector and the pointing of the vertices of the triangular perturbation. Note that the coordinate in (A) is 90° rotated, due to it being in the reciprocal space. (B) and (C) illustrate in the configuration space the circular polarization contrast projected in out-of-plane direction, $C_z = \frac{I_R - I_L}{I_R + I_L}$, where I_L, I_R denote the intensity of the left- and right-circular polarization intensity. We note that it has $|C_z| = 1$ for the center point of the holes.




Figure 7.6: Moving blue-detuned lattice atom loading and trapping scheme

Each of the panels in this picture is a vertical cross-section of a unit cell in the 2D photonic crystal, with the grey-out areas being occupied by Silicon Nitride. Panel (A) to (G) depict the moving lattice in the close vicinity of the 2D photonic crystal.

The lattice is formed of blue-detuned beams, and a trapped atom (white circle)

follows the intensity minimum through the hole. (H) Switching off one of the lattice beams (the downward beam in this plot) resulted in a local minimum point in the device layer.

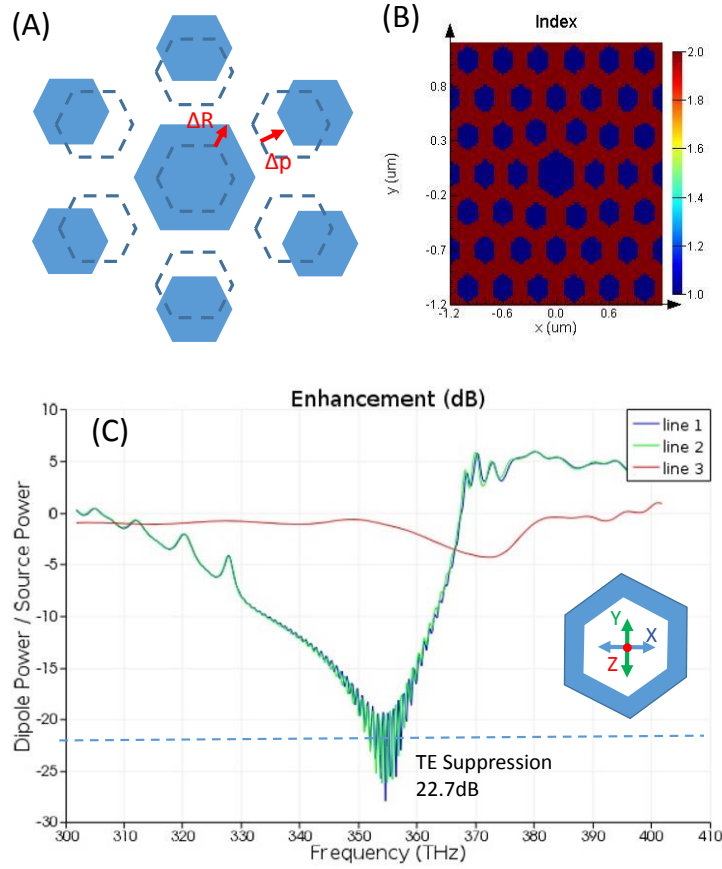


Figure 7.7: Suppression of in-plane dipole emission with point defect

(A) The geometry of the local defect is modified from the background crystal (material thickness 200nm, lattice constant = 400nm, tether width = 170nm) by $\Delta p = 45\text{nm}$, $\Delta R = 90\text{nm}$. (B) The index of refraction plot of the defect. (C) The Green's function enhancement is calculated with a dipole source polarized in X (K lattice direction), Y (M lattice direction), and Z (perpendicular to device plane). The plot shows that, along with the band-gap from 330THz to 365THz, a broad dip exists at 355THz for the two TE modes (X, blue trace; Y, green trace) that goes down to 22.7dB suppression comparing to free space. The TM mode (red) remains unaffected.

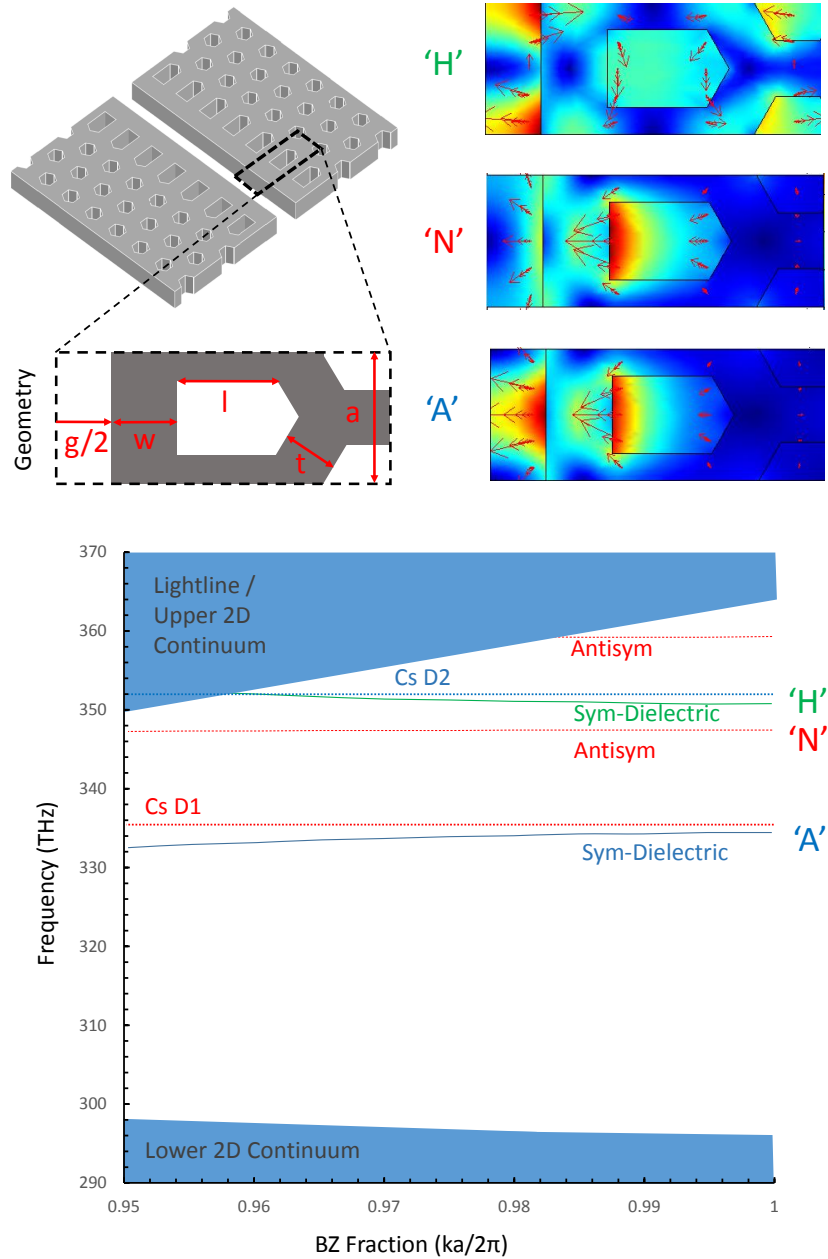


Figure 7.8: Embedded 1D defect 'slot' waveguide bands and mode profiles

A 1D defect is introduced to the 2D lattice by introduction of a vacuum gap (g), extra material on the edge of the 2D crystal (w), and lengthened inner-most hole (l), as plotted on the upper right panel. The structure is symmetric to the center-of-vacuum-gap plane. For simplicity, only the right half of a unit cell is plotted in this plot. With the parameters $(g, w, l) = (350\text{nm}, 207\text{nm}, 312\text{nm})$, along with a 2D background crystal of $(a, t) = (407.5\text{nm}, 170\text{nm})$, we can achieve the dual-band alignment to Cesium resonances as in the APCW. An 'antisymmetric' mode, denoted 'N' on this plot, exists in the band-gap range, and needs to be avoided for decay rate suppression consideration.

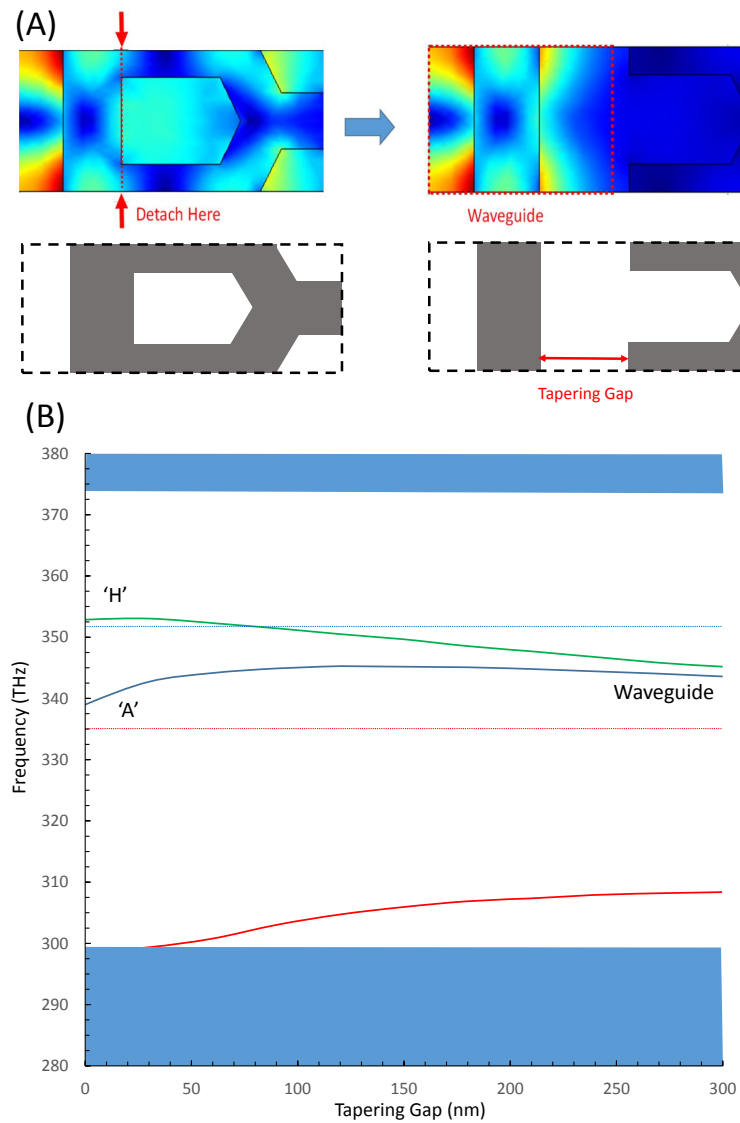


Figure 7.9: Tapering slot waveguide

(A) The 2D crystal can be 'detached' from the end of the inner-most row of holes.

As the 2D crystal recedes and forms a gap between the now-unpatterend waveguide, the guided mode becomes confined to the waveguide. (B) The band-gap closes as the 2D crystal recedes away from the evanescent field of the waveguide. The 'edge' mode propagating along the cutting edge of the 2D crystal emerges from the lower 2D continuum (solid red trace), but is sufficiently far in frequency from the modes of interest. The dotted traces are Cesium D1 (dotted red) and D2 (dotted blue) respectively.




Figure 7.10: SEM image of slot waveguide taper

This SEM image depicts the incoming unpatterned double-beam waveguide (left) tapering into the slot waveguide (right) using the tapering scheme described in Section 7.4.

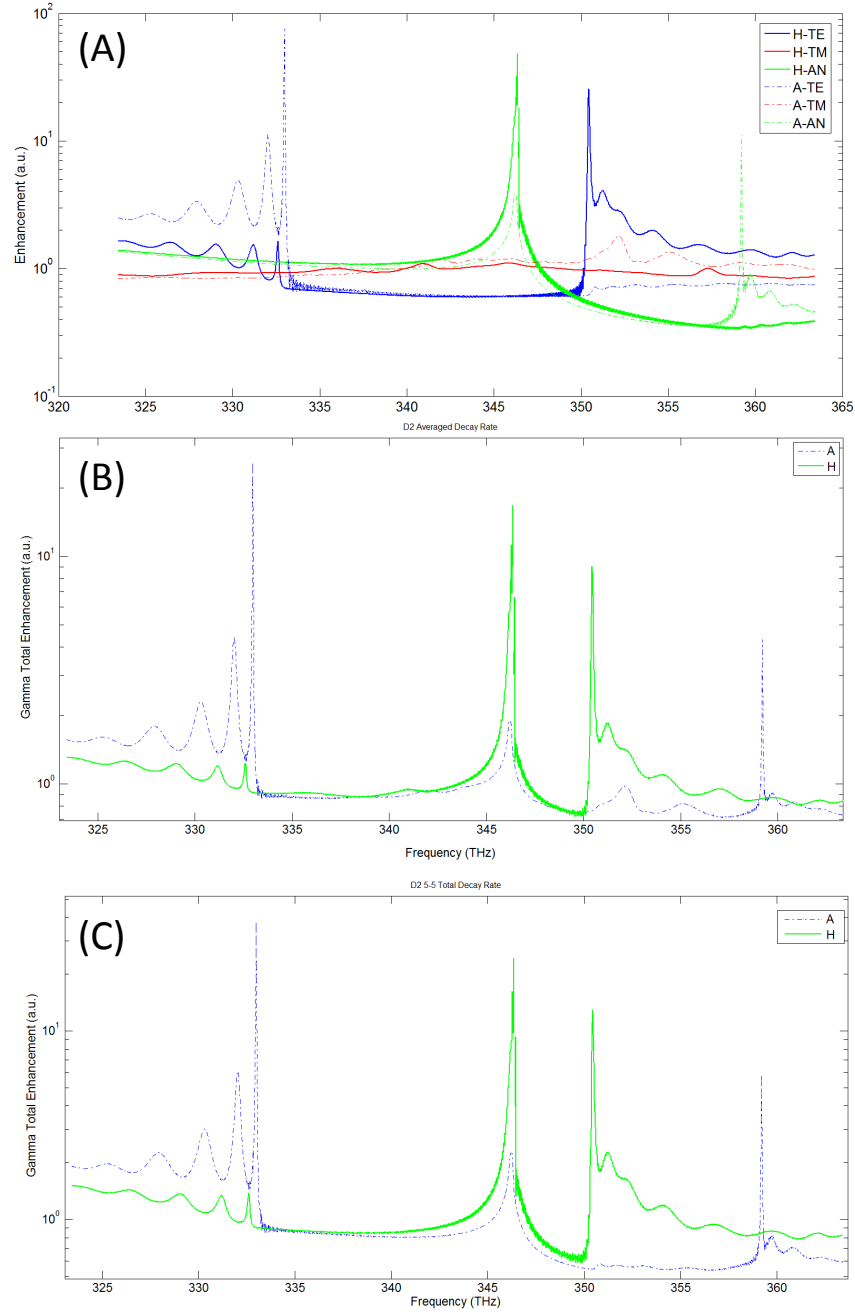


Figure 7.11: Atom decay rate calculation for the slot waveguide

(A) The enhancement for the imaginary part of Green's function over free space value calculated at the maximum of 'H' mode (solid lines) and of 'A' mode (dashed lines). The color coding are blue for symmetric TE mode, green for anti-symmetric TE mode, and red for TM mode. This information allows us to calculate averaged atom decay rate for (B) all F and m_f states averaged and (C) D2 stretched state $F=5$, $m_f=5$. The geometry parameters are as follow: $a_{lattice}=407.5\text{nm}$, $t_{2D}=170\text{nm}$, $l_{finalhole}=312\text{nm}$, $w_0=207\text{nm}$, $g_{vacuum}=350\text{nm}$; total crystal length = $30\mu\text{m}$, plus 15 sites of tapering each. The 1D defect is cladded with 6 sites of 2D crystal on each sides.

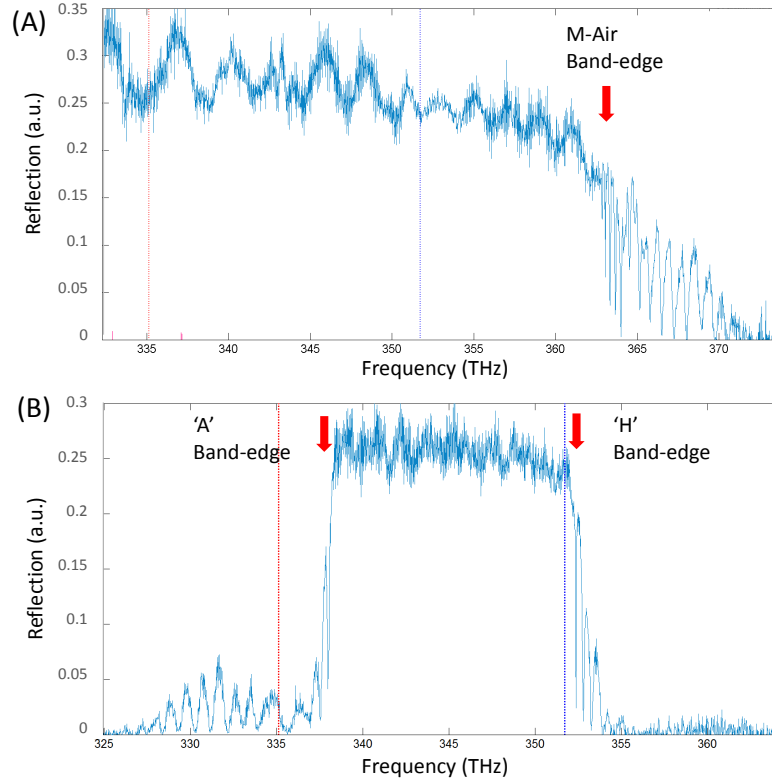


Figure 7.12: Measured optical spectra of 2D and slot waveguide devices

(A) A 2D photonic crystal device reflection spectrum from the M direction showing the M-Air band edge. We observe that the upper side of the 2D band-gap indeed covers the range of interest containing the Cesium lines. A 2D crystal with the same geometry parameters but coupled to the K orientation is needed to confirm the lower side of the band-gap. (B) Reflection spectrum measurement of a slot waveguide device. The spectrum demonstrates typical band-gap and band-edge resonances as in the APCW devices, with a slightly lower coupling efficiency (25~30% typical, where APCW devices typically show 35% round-trip reflection). The asymmetry of the lower- and upper-band-edge resulted from the asymmetric tapering behavior of the two bands as shown in fig. 7.9.

Chapter 8

OUTLOOK

Historically, the field of atomic physics has provided a reliable and clean systems that are invaluable to our understanding of quantum physics. The goal we have for our project is to also allow it the versatile capability provided by nanophotonics. In this project we have developed various optical designs with diverse properties. We hope that in the near future, these components can be combined to perform useful quantum mechanical functions. This chapter will describe the ongoing effort and future directions for our project.

8.1 Migration Toward Free-Space Coupled Devices

The next step to improve our system would be application of the free-space coupling technique. With this upgrade, we should be able to remove the limit of device footprint associated with the fiber. The fiber has a cladding diameter of $\sim 150\mu\text{m}$, but free-space coupling should only require a clearance of $10\mu\text{m}$ 2-waist (e^{-4} in intensity) diameter of the guided mode at coupling waveguide. This will allow up to a factor of 15 increase in device number per chip. Also, the smaller device size would reduce the length and hence thermal resistance of the tethers, likely increasing the thermal handling capability by a similar factor.

Ideally, a 'loadlock' design with a sample exchange chamber that is separately pumped would allow chips to be replaced without venting the main chamber. A chamber design that allows rapid turn-around of chips will utilize the scalability of lithography-fabricated devices to the fullest. We hope that the simplification of the experiment following application of free-space coupling, namely, the elimination of unconventional elements in vacuum chamber such as fiber feedthrough and gluing, would reduce the entrance barrier for future research interest in this field. We also hope that the additional flexibility afforded by the free-space optics can eventually allow high-efficiency coupling beyond Gaussian optics, so that coupling-efficiency-critical systems, such as atom-photon quantum memories, and functional quantum networks, may be realized in our system.

8.2 Deterministic Loading and Positioning with Optical Tweezers

Our current scheme for atom loading into our devices is probabilistic, in that we load from a MOT directly, and hence assumed a Poissonian distribution of atom numbers. We only have limited spatial control through our understanding of guided mode profiles and super-modes. The development for the near future would be to use the optical tweezers technology to place atoms in a controllable manner. This method should immediately allow us to directly study the band-gap-mediated finite-range interaction [56, 78], or potential directional interaction with 2D devices. The effort here would be two-fold, firstly in changing the chamber design to allow the high numerical aperture optics to be placed close to the chip surface, and secondly in designing the nano-structure reflection pattern to ensure the tweezers trap is converted to a local trap near the structure. With the advancement of the multiple-tweezers techniques as in [79], we can potentially explore multi-atom physics in a highly controlled manner, including quantum many-body physics for atoms and photons in 1D and 2D structures.

8.3 Electro-capacitive dynamic tuning of slot waveguide

We can make use of the mechanically pliable nature of the nano-string geometry to fabricate devices that can be dynamically deformed by capacitive actuation. On the small scale, a controlled deformation range of few nanometers may be used to dynamically lock our device frequencies, such as a cavity or band-edge resonance, to some reference frequency of interest. The rate of this tuning should only be limited by the mechanical resonance frequency of the double-beam section that is being tuned, which for the APCW is at few-MHz range. A faster tuning rate may be achieved if the tuning motion excites not the nano-string modes, but the longitudinal motion parallel to the beams; see Figure 8.1. Due to the faster speed of sound in this configuration ($c_{transverse} = \sqrt{\frac{T}{\lambda}} \simeq 480m/s$, compared to $c_{longitudinal} = \sqrt{\frac{K}{\rho}} \simeq 7.5km/s$). Finally, we can create large deformation of our device up to several μm . If this can be achieved in few milisecond time-scale, we might be able to dynamically manipulate the position and properties of the optical dipole trap adiabatically by capacitively actuating the device. Preliminary exploration of capacitive tuning yielded positive results; see Figure 8.2.

8.4 Potential Future Development

The strength of the nanophotonic platform developed in this project comparing to conventional atom-optics is in its ease of adaption and integration of new designs.




Figure 8.1: Vibration modes of nano-string photonic crystal

The vibration modes of our photonic crystal are simulated using FEM calculation, only one beam is plotted for simplicity. We observe two primary types of modes, first being the transverse or 'string' modes as plotted in (A), where the motion is actuated by tension on the device and is perpendicular to the waveguide. There are the 'longitudinal' modes as shown in (B), where the motion is actuated by the material's elasticity and is parallel to the waveguide. (C) The periodic modulation of the material in our photonic crystal also introduces 'phononic band-edge' modes where the motion of each consecutive lattice changes sign. (D1~D3) demonstrates the profile of the first three 'string' modes of our devices. The color in the plot indicates total displacement.

As our project develops, the processing procedure is now reliable, and the yield rate of devices sufficiently high. We believe our system can readily allow integration of more functionalities, beyond the aforementioned near-future steps.

Circular Polarization Engineering

During the development of our project, we have arrived at various designs capable of producing highly circular local polarization points. This will eventually allow the guided mode light to couple with the m_f sub-level of atoms, achieving directional emission similar to that reported in [65] and [64], but on a multiple-atom chain, or even a 2D super-lattice. The discrete nature of the atom states naturally provide nonlinearity at a single-photon level similar to that reported in [17], which can then become a platform for quantum information processing. On a more applied technology consideration, the engineering of m_f dependent scattering of atoms coupled to circular polarized waveguide modes can potentially be applied to precision magnetic field sensing.

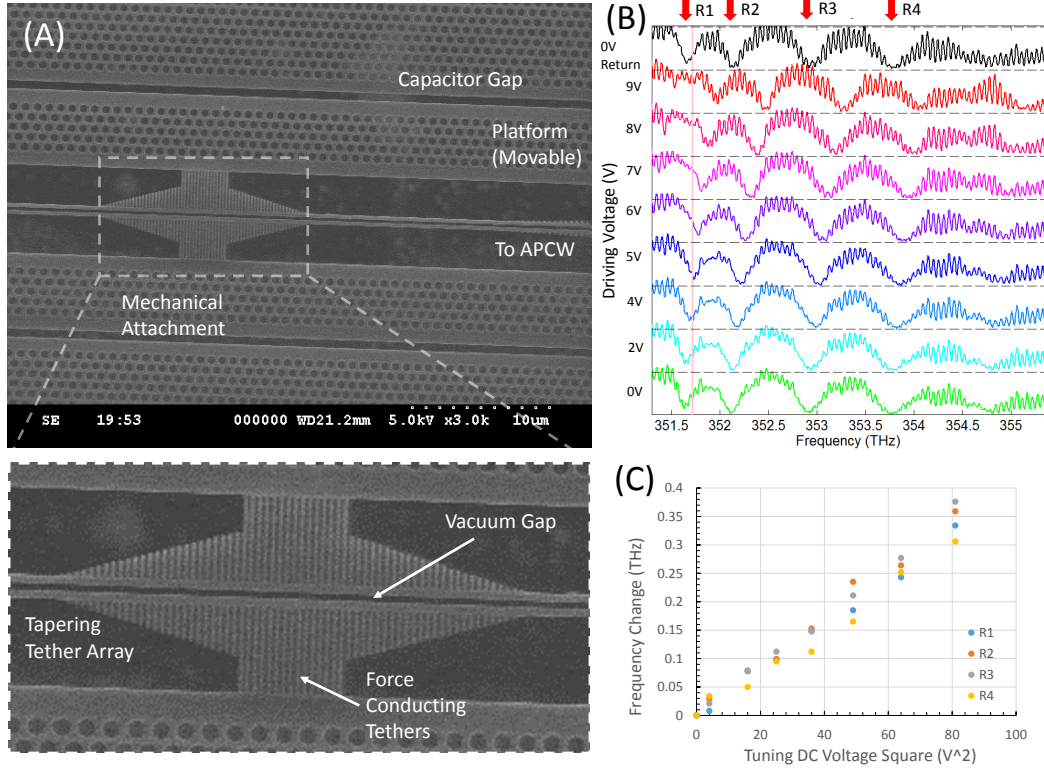


Figure 8.2: Capacitive actuation tuning of double-beam device

(A) We have built preliminary test devices with capacitor wiring placed on mechanical supporting pads. The actuation force is transduced to the double-beam structure by using a set of tethers. We note that this SEM image shows only the pads and the tethering, but not the capacitor wiring. We also note that the hexagonal structure here are only for mechanical support. (B) Electrostatic tuning is carried out on this chip. We observe the spectrum features on this device shift toward higher frequency as the capacitor voltage is increased and hence the vacuum gap of the photonic crystal is enlarged. (C) Plotting the frequency as function of driving voltage squared shows a linear tuning rate of $43\text{GHz}/V^2$.

Multi-port devices

We can imagine taking a further step toward photonic integration by building multi-port devices. Preliminary experiments with such device have been carried out; see Figure 8.3. For the full 2D photonic devices, a natural 6-fold rotation symmetry provided us with six input ports on the M direction of a hexagonal lattice. This allows us three 'beam-paths' on the same device, where we can use as local control fields, or record output from different ports to study correlation between ports. This technology can be applied to our slot waveguide devices in the following way: the 'slot' waveguide channel propagates on the K direction of the background 2D lattice, while the 2D modes are easily coupled in the M direction, perpendicular to the K

direction. We can potentially place the same tapering and launching ports as in full 2D devices perpendicular to the embedded 1D waveguide. These perpendicular ports can have a mode profile of $\sim 10\mu\text{m}$, capable of resolving the length of the 1D waveguide to several sub-sections. This could allow position resolution and selective addressing of trapped atoms in a fully on-chip manner.




Figure 8.3: Preliminary experiment of multi-port device

(A) We fabricate a six-sided hexagonal lattice slab, which is connected to our waveguide and then to conventional fibers (Port 1 to 3). The through-ports are connected directly into the Si substrate edge for this early test device. (B) Optical image of released device. Note that there are a total of 12 connections to the pad, the thinner ones are mechanical support tethers that hold the suspended slab in place. Plot (C) and (D) show injection of light through two of the ports, which transmit to their respective through-ports.

Potential for 'random-access' quantum system

In the foreseeable future, we believe our photonic platform has potential to develop further into a platform for quantum information science. Here we present a potential direction toward random-accessible quantum memory using our system. Our current

set of photonic structures is capable of building 2D devices that can have two sets of input/output ports. We can imagine that we make use of them in a 'bit-line, word-line' configuration (see Figure 8.4), where atoms can be placed in this 2D array. Each atom can then be individually addressed by a coordinate system defined by the indices of the two orthogonal sets of ports of the device. Then, the array can potentially provide quantum information processing capability, such as selected information storage and retrieval in atoms, or selectively interact a subset of atoms on the photonic platform to engineer entanglement between atoms. We hope that the union of neutral atoms and nanophotonics can bring together the advantages of both worlds to create a clean and controllable hybrid system with engineer-able quantum functions.




Figure 8.4: Potential development of devices with random access capability

We hope that by combining the various structure designs we have developed in this project, future devices can have random-access site addressing capability. By addressing localized atoms from two orthogonal sets of ports, selecting of single site, or selecting pairs of atoms to create interaction, should be achievable on our devices.

BIBLIOGRAPHY

- [1] Nicolò Spagnolo et al. “Experimental validation of photonic boson sampling”. In: *Nature Photonics* 8 (2014), pp. 615–620. doi: 10.1038/nphoton.2014.135.
- [2] Lian-Wee Luo et al. “WDM-compatible mode-division multiplexing on a silicon chip”. In: *Nature Communications* 5 (2014), p. 3069. doi: 10.1038/ncomms4069.
- [3] Vladimir M. Shalaev. “Optical negative-index metamaterials”. In: *Nature Photonics* 1 (2007), pp. 41–48. doi: 10.1038/nphoton.2006.49.
- [4] Mindy Lee and Philippe M. Fauchet. “Two-dimensional silicon photonic crystal based biosensing platform for protein detection”. In: *Optics Express* 15 (8 2007), pp. 4530–4535. doi: 10.1364/OE.15.004530.
- [5] Z. Hu and H. J. Kimble. “Observation of a single atom in a magneto-optical trap”. In: *Optics Letters* 19 (22 1994), pp. 1888–1890. doi: 10.1364/OL.19.001888.
- [6] Peter Lodahl, Sahand Mahmoodian, and Søren Stobbe. “Interfacing single photons and single quantum dots with photonic nanostructures”. In: *Reviews of Modern Physics* 87 (2015), p. 347. doi: 10.1103/RevModPhys.87.347.
- [7] Igor Aharonovich, Andrew D. Greentree, and Steven Prawer. “Diamond photonics”. In: *Nature Photonics* 5 (2011), pp. 397–405. doi: 10.1038/nphoton.2011.54.
- [8] H. J. Kimble. “Strong interactions of single atoms and photons in cavity QED”. In: *Physica Scripta* 1998.T76 (1998), pp. 127–137. URL: <http://stacks.iop.org/1402-4896/1998/i=T76/a=019>.
- [9] Darrick E. Chang, Vladan Vuletic, and Mikhail D. Lukin. “Quantum nonlinear optics - photon by photon”. In: *Nature Photonics* 8 (2014), pp. 685–694. doi: 10.1038/nphoton.2014.192.
- [10] Ranojoy Bose et al. “All-optical coherent control of vacuum Rabi oscillations”. In: *Nature Photonics* 8 (2014), pp. 858–864. doi: 10.1038/nphoton.2014.224.
- [11] Monika H. Schleier-Smith, Ian D. Leroux, and Vladan Vuletic. “Squeezing the collective spin of a dilute atomic ensemble by cavity feedback”. In: *Physical Review A* 81 (2010), p. 021804. doi: 10.1103/PhysRevA.81.021804.
- [12] D. Reitz et al. “Backscattering properties of a waveguide-coupled array of atoms in the strongly nonparaxial regime”. In: *Physical Review A* 89 (2014), p. 031804. doi: 10.1103/PhysRevA.89.031804.

- [13] C. Sayrin et al. “Storage of fiber-guided light in a nanofiber-trapped ensemble of cold atoms”. In: *Optica* 2 (4 2015), pp. 353–356. doi: 10.1364/OPTICA.2.000353.
- [14] D. E. Chang, J. I. Cirac, and H. J. Kimble. “Self-Organization of Atoms along a Nanophotonic Waveguide”. In: *Physical Review Letters* 110 (2013), p. 113606. doi: 10.1103/PhysRevLett.110.113606.
- [15] Xu Yi et al. “Single-mode dispersive waves and soliton microcomb dynamics”. In: *Nature Communications* 8.14869 (2017). doi: 10.1038/ncomms14869.
- [16] K. M. Birnbaum et al. “Photon blockade in an optical cavity with one trapped atom”. In: *Nature Letters* 436 (2005), pp. 87–90. doi: 10.1038/nature03804.
- [17] Wenlan Chen et al. “All-Optical Switch and Transistor Gated by One Stored Photon”. In: *Science* 341 (6147 2013), pp. 768–770. doi: 10.1126/science.1238169.
- [18] G. D. Fuchs et al. “A quantum memory intrinsic to single nitrogen–vacancy centres in diamond”. In: *Nature Physics* 7 (2011), pp. 789–793. doi: 10.1038/nphys2026.
- [19] R. W. Andrews et al. “Bidirectional and efficient conversion between microwave and optical light”. In: *Nature Physics* 10 (2014), pp. 321–326. doi: 10.1038/nphys2911.
- [20] Kejie Fang et al. “Optical transduction and routing of microwave phonons in cavity-optomechanical circuits”. In: *Nature Photonics* 10 (2016), pp. 489–496. doi: 10.1038/nphoton.2016.107.
- [21] V. S. C. Manga Rao and S. Hughes. “Single Quantum Dot Spontaneous Emission in a Finite-Size Photonic CrystalWaveguide: Proposal for an Efficient “On Chip” Single Photon Gun”. In: *Physical Review Letters* 99 (2007), p. 193901. doi: 10.1103/PhysRevLett.99.193901.
- [22] A. Furusawa et al. “Unconditional Quantum Teleportation”. In: *Science* 23 (5389 1998), pp. 706–709. doi: 10.1126/science.282.5389.706.
- [23] Nicolas Gisin and Rob Thew. “Quantum communication”. In: *Nature Photonics* 1 (2007), pp. 165–171. doi: 10.1038/nphoton.2007.22.
- [24] T. E. Northup and R. Blatt. “Quantum information transfer using photons”. In: *Nature Photonics* 8 (2014), pp. 356–363. doi: 10.1038/nphoton.2014.53.
- [25] H. J. Kimble. “The quantum internet”. In: *Nature* 453 (2008), pp. 1023–1030. doi: 10.1038/nature07127.
- [26] John D. Jackson. *Classical Electrodynamics*. 3rd ed. Wiley, 1998.
- [27] A. W. Snyder and J. Love. *Optical Waveguide Theory*. 1st ed. Springer US, 1998.

- [28] Sean M. Meenehan. “Cavity Optomechanics at Millikelvin Temperatures”. PhD thesis. Caltech, 2015.
- [29] Inc. Lumerical Solutions. *Knowledge Base*. URL: <https://kb.lumerical.com/>.
- [30] C. Wuttke and A. Rauschenbeutel. “Thermalization via Heat Radiation of an Individual Object Thinner than the Thermal Wavelength”. In: *Physical Review Letters* 111 (2 2013), p. 024301. doi: 10.1103/PhysRevLett.111.024301.
- [31] Paul E. Barclaya et al. “Integration of fiber-coupled high-Q SiNx microdisks with atom chips”. In: *Applied Physics Letters* 89 (2006), p. 131108. doi: 10.1063/1.2356892.
- [32] R. L. Edwards, G. Coles, and W. N. Sharpe Jr. “Comparison of tensile and bulge tests for thin-film silicon nitride”. In: *Experimental Mechanics* 44 (1 2002), pp. 49–54. doi: 10.1007/BF02427976.
- [33] Scott S. Verbridge and Jeevak M. Parpia. “High quality factor resonance at room temperature with nanostrings under high tensile stress”. In: *Journal of Applied Physics* 99 (2006), p. 124304. doi: 10.1063/1.2204829.
- [34] Justin D. Cohen, Seán M. Meenehan, and Oskar Painter. “Optical coupling to nanoscale optomechanical cavities for near quantum-limited motion transduction”. In: *Optics Express* 21 (9 2013), pp. 11227–11236. doi: 10.1364/OE.21.011227.
- [35] S. J. McNab, N. Moll, and Y. A. Vlasov. “Ultra-low loss photonic integrated circuit with membrane-type photonic crystal waveguides”. In: *Optics Express* 11 (22 2003), pp. 2927–2939. doi: 10.1364/OE.11.002927.
- [36] D. Taillaert et al. “An out-of-plane grating coupler for efficient butt-coupling between compact planar waveguides and single-mode fibers”. In: *IEEE Journal of Quantum Electronics* 38 (7 2002), pp. 949–955. doi: 10.1109/JQE.2002.1017613.
- [37] Eric R.I. Abraham and Eric A. Cornell. “Teflon feedthrough for coupling optical fibers into ultrahigh vacuum systems”. In: *Applied Optics* 37 (10 1998), pp. 1762–1763. doi: 10.1364/AO.37.001762.
- [38] L. M. Bennie et al. “A versatile high resolution objective for imaging quantum gases”. In: *Optics Express* 21 (7 2013), pp. 9011–9016. doi: 10.1364/OE.21.009011.
- [39] Michael D. Henry. “ICP Etching of Silicon for Micro and Nanoscale Devices”. PhD thesis. Caltech, 2010.
- [40] Oxford Instruments. *Etching, Deposition, and Growth Processes*. URL: <https://www.oxford-instruments.com/products/etching-deposition-and-growth/processes>.

- [41] Y.-C. Cheng. *Macroscopic and Statistical Thermodynamics*. World Scientific Publishing Co. Pte. Ltd., 2006.
- [42] Ultratech/CNT. *Atomic Layer Deposition Tutorials Knowledge Center*. URL: <http://cambridgenanotechald.com/ald-tutorial.shtml>.
- [43] K. R. Williams, K. Gupta, and M. Wasilik. “Etch rates for micromachining processing - Part II”. In: *Journal of Microelectromechanical Systems* 12.6 (2003), pp. 761–778. doi: 10.1109/JMEMS.2003.820936.
- [44] Steven Chu et al. “Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure”. In: *Physical Review Letters* 55.48 (1985). doi: 10.1103/PhysRevLett.55.48.
- [45] J. D. Miller, R. A. Cline, and D. J. Heinzen. “Far-off-resonance optical trapping of atoms”. In: *Physical Review A* 47.6 (1993), pp. 4567–4570. doi: 10.1103/PhysRevA.47.R4567.
- [46] H. Schroeder, E. Obermeier, and A. Steckenborn. “Micropyramidal hillocks on KOH etched 100 silicon surfaces: formation, prevention and removal”. In: *Journal of Micromechanics and Microengineering* 9.2 (1998), pp. 139–145. doi: 10.1088/0960-1317/9/2/309.
- [47] Paul E. Barclay. “Fiber-coupled nanophotonic devices for nonlinear optics and cavity QED”. PhD thesis. Caltech, 2007.
- [48] Richard A. Norte. “Nanofabrication for On-Chip Optical Levitation, Atom-Trapping, and Superconducting Quantum Circuits”. PhD thesis. Caltech, 2015.
- [49] Jonathan D. Hood. “Atom-Light Interactions in a Photonic Crystal Waveguide”. PhD thesis. Caltech, 2016.
- [50] Andrew McClung. “Photonic Crystal Waveguides for Integration into an Atomic Physics Experiment”. PhD thesis. Caltech, 2017.
- [51] B. Brandstaetter et al. “Integrated fiber-mirror ion trap for strong ion-cavity coupling”. In: *Review of Scientific Instruments* 84 (2013), p. 123104. doi: 10.1063/1.4838696.
- [52] Gunter Krautheim et al. “Mechanical stress in ALD-Al₂O₃ films”. In: *Applied Surface Science* 252 (2005), pp. 200–204. doi: 10.1016/j.apsusc.2005.01.118.
- [53] John D. Joannopoulos et al. *Photonic Crystals: Molding the Flow of Light*. 2nd ed. Princeton University Press, 2008.
- [54] R. Shankar. *Principles of Quantum Mechanics*. 2nd ed. Plenum Press, 1994.
- [55] Stefan Yoshi Buhmann et al. “Casimir-Polder forces: A nonperturbative approach”. In: *Physical Review A* 70 (2004), p. 052117. doi: 10.1103/PhysRevA.70.052117.

- [56] A. Gonzalez-Tudela et al. “Subwavelength vacuum lattices and atom–atom interactions in two-dimensional photonic crystals”. In: *Nature Photonics* 9 (2015), pp. 320–325. doi: 10.1038/nphoton.2015.54.
- [57] S.-P. Yu et al. “Nanowire Photonic Crystal Waveguides for Single-Atom Trapping and Strong Light-Matter Interactions”. In: *Applied Physics Letters* 104 (2014), p. 111103. doi: 10.1063/1.4868975.
- [58] C.-L. Hung et al. “Trapped atoms in one-dimensional photonic crystals”. In: *New Journal of Physics* 15 (2013), p. 083026. doi: 10.1088/1367-2630/15/8/08302.
- [59] Steven G. Johnson et al. “Adiabatic theorem and continuous coupled-mode theory for efficient taper transitions in photonic crystals”. In: *Physical Review E* 66 (2004), p. 066608. doi: 10.1103/PhysRevE.66.066608.
- [60] A. Goban et al. “Atom–Light Interactions in Photonic Crystals”. In: *Nature Communications* 5 (2014), p. 3808. doi: 10.1038/ncomms4808.
- [61] Thompson JD et al. “Coupling a single trapped atom to a nanoscale optical cavity”. In: *Science* 340 (6137 2013), pp. 1202–5. doi: 10.1126/science.1237125.
- [62] A. Goban et al. “Superradiance for Atoms Trapped along a Photonic Crystal Waveguide”. In: *Physical Review Letters* 115 (2015), p. 063601. doi: 10.1103/PhysRevLett.115.063601.
- [63] J. D. Hood et al. “Atom–Atom Interactions Around the Band Edge of a Photonic Crystal Waveguide”. In: *Proceedings to National Academy of Science* 113.38 (2016), pp. 10507–10512. doi: 10.1073/pnas.1603788113.
- [64] Immo Söllner et al. “Deterministic photon–emitter coupling in chiral photonic circuits”. In: *Nature Nanotechnology* 10 (2015), pp. 775–778. doi: 10.1038/nnano.2015.159.
- [65] Fam Le Kien and A. Rauschenbeutel. “Anisotropy in scattering of light from an atom into the guided modes of a nanofiber”. In: *Physical Review A* 90 (2014), p. 023805. doi: 10.1103/PhysRevA.90.023805.
- [66] Jr. Joseph R. Buck. “Cavity QED in Microsphere and Fabry-Perot Cavities”. PhD thesis. Caltech, 2003.
- [67] J. Dalibard and C. Cohen-Tannoudji. “Laser cooling below the Doppler limit by polarization gradients: simple theoretical models”. In: *Journal of the Optical Society of America B* 6 (11 1989), pp. 2023–2045. doi: 10.1364/JOSAB.6.002023.
- [68] W. Wohlleben, F. ChevyK. Madison, and J. Dalibard. “An atom faucet”. In: *The European Physical Journal D* 15 (2 2001), pp. 237–244. doi: 10.1007/s100530170171.

- [69] Agilent Technologies. *High and Ultra-High Vacuum for Science Research*. URL: http://www.agilent.com/cs/library/training/Public/UHV_Seminar_Handbook.pdf.
- [70] Neil W. Ashcroft and N. David Mermin. *Solid State Physics*. Thomson Learning Inc., 1976.
- [71] Feng Wen et al. “Two-dimensional photonic crystals with large complete photonic band gaps in both TE and TM polarizations”. In: *Optics Express* (16 2008), pp. 12278–12289. DOI: 10.1364/OE.16.012278.
- [72] Steven G. Johnson et al. “Linear waveguides in photonic-crystal slabs”. In: *Physical Review B* 62.12 (2000), pp. 8212–8222. DOI: 10.1103/PhysRevB.62.8212.
- [73] Juan A. Muniz. “Nanoscopic Atomic Lattices with Light-Mediated Interactions”. PhD thesis. Caltech, 2017.
- [74] J. Witzens, M. Loncar, and A. Scherer. “Self-collimation in planar photonic crystals”. In: *IEEE Journal of Selected Topics in Quantum Electronics* 8.6 (2002), pp. 1246–1257. DOI: 10.1109/JSTQE.2002.806693.
- [75] Peter Lodahl et al. “Chiral quantum optics”. In: *Nature* 541 (2017), pp. 473–480. DOI: 10.1038/nature21037.
- [76] Marcelo Davanço et al. “Slot-mode-coupled optomechanical crystals”. In: *Optics Express* 20 (22 2012), pp. 24394–24410. DOI: 10.1364/OE.20.024394.
- [77] D. A. Steck. *Cesium D Line Data*. URL: <http://steck.us/alkalidata/>.
- [78] J. S. Douglas et al. “Quantum many-body models with cold atoms coupled to photonic crystals”. In: *Nature Photonics* 9 (2015), pp. 326–331. DOI: 10.1038/nphoton.2015.57.
- [79] A. M. Kaufman et al. “Two-particle quantum interference in tunnel-coupled optical tweezers”. In: *Science* 345 (6194 2014), pp. 306–309. DOI: 10.1126/science.1250057.
- [80] Su-Peng Yu. *Software Package for Nano-Photonic Platform for Atom-Light Interaction Project*. DOI: 10.22002/D1.241. URL: <https://caltechdata.tind.io/records/241>.

*Appendix A***FEM SIMULATION METHODS**

This chapter describes our application of the Finite-Element Method (FEM) to our project. The calculations are carried out using *COMSOL Multiphysics* version 3.5a with a *Matlab* interface. Some examples from our COMSOL code package can be accessed in the CaltechDATA archive; see reference [80]. For our calculation, the RF module and the MEMS structural mechanics module will be used. The steps for operation of the COMSOL program are as follow:

- An application mode is selected, and determines a set of governing differential equations to be solved.
- The geometry to be simulated is created from basic geometry shapes or imported from CAD drawings, with suitable simulation volume defined to contain it.
- Physical properties such as material index of refraction or Young's modulus are specified for the structures.
- Boundary conditions are specified for the simulation volume to determine symmetry.
- The program runs its 'meshing' function on the geometry, assigning 'elements' of the finite-element method of the suitable type in its simulation volume.
- The program converts the governing differential equations into numerical relations between the elements, forming a matrix of parameters
- A suitable numerical solver program is called to solve the problem, typically in the form of inverting of the parameter matrix or calculating its eigenvalues and eigenvectors.
- Calculation and plots can then be produced from the solutions.

With version 3.5a, these steps can be done either via the program visual interface, or by using a Matlab script file. Two typical calculations we do for our project will

be described in the following sections. More information can also be found in the COMSOL program manuals and example model libraries. We note that COMSOL calculation is demanding on the computer memory space, and a typical simulation of 500,000 degrees of freedom can easily use up 50Gb of memory during the solution process.

A.1 Optical Band Structures

We calculate the band structures of a photonic crystal using the FEM method. The 'RF Wave' application mode is selected, and solves the Maxwell's equations. The band simulation is typically done in 3D.

Geometry and boundary conditions

We typically generate the geometry in Matlab by drawing poly-lines, then extrude this into COMSOL geometry object. The geometry is specified with its index of refraction for the material, where we have used the bulk refractive index for Silicon Nitride for our purposes. We then specify the Bloch periodicity condition for the boundaries on the propagation direction. In COMSOL this is under the name of Floquet periodic condition, where a k -vector is used to specify the phase difference between two parallel, identically-shaped surfaces. This can also be applied for 2D structures, using multiple pairs of Bloch condition surfaces that contain the 2D volume. In addition to the Bloch surfaces, a symmetry plane is placed parallel to the device plane at half the thickness of the material. Selecting a 'perfect electric conductor' (PEC) condition for this plane results in TM modes, and 'perfect magnetic conductor' (PMC) produces TE modes. The remaining boundaries should be placed sufficiently far so the solution mode profile is negligible at their locations. Either PEC or PMC condition would work for these boundaries.

Meshing and solving

We found it to be important that the Bloch boundaries to have identical mesh points, lest the occurrence of failure to converge increases. This can be done by meshing one of the Bloch surfaces and then copying it to the other surface. This should be done for every pairs of Bloch surfaces. Once that is done, the remaining volume can be meshed using the automatic mesh function from COMSOL. We note that narrow gaps, sharp corners, and use of elliptical shapes (such as for the simulation volume) should be avoided, as the meshing algorithm tends to place extremely dense mesh in these points and their vicinity.

We call the eigenfrequency solver to solve our system. The output information using this calculation method yields a set of frequencies for each k-vector. Running a loop over all k values in the Brillouin zone of interest produces the dispersion curves. We note that there is a 'shift' parameter in the eigenvalue solver, when specified, allows the solver to search for eigen-solutions with frequency near this shift parameter. We note that for 2D-embedded waveguide structures, the shift parameter should be specified, to avoid the upper- and lower-2D continuum modes.

Information output

The information of interest in these simulations are typically the dispersion curve and the mode profile. Since the program looks for a set number of eigen-modes near a specified shift frequency, then orders it by its distance to the shift frequency, the n-th output frequencies from one solution run may not be the same mode as the n-th output from a different run. Additionally, when looking for modes above the light-line, un-guided modes start to appear manifesting themselves as modes bouncing around the outer boundaries of the simulation volume. Sorting these modes which has non-vanishing intensity at the far boundaries needs to be done to remove them from the actual spectra of interest. The program can be requested to output electric and magnetic fields interpolated at any given set of coordinates in the simulation volume. Further information such as intensity or polarization can be calculated from the field patterns.

A.2 Mechanical stress and resonance frequency

We use the finite element method to calculate the mechanical deformation of our devices as they relieve their intrinsic stress upon being released from the Silicon substrate. The MEMS module for structural mechanics is used. A static calculation can calculate the relaxation of the stress and deformation, and following that, we can calculate the resonance modes of a structure with consideration of the residual stress.

Geometry and boundary conditions

The geometry typically simulated in this scenario is a section of the device, such as the double-beams for the photonic crystal, or the coupling waveguide and its tethers. The mechanical simulation does not require an 'air-box' around the device geometry like the optical simulation. Boundary conditions are typically set to 'fixed' for parts of the device and tethers that are anchored to the Silicon substrate, and the rest set to

'free' as default. COMSOL supplies a material library of MEMS materials, where we make use of the data for Silicon Nitride such as density and Young's modulus. The program allows input of initial stress, where we typically assume an isotropic tensile stress of 800MPa.

Meshing and solving

Many of the structures simulated here has very large aspect ratio, due to our devices being very thin and long. We found that the default 'triangular (advancing front)' meshing method to sometimes fail, setting the meshing method to 'triangular' seems to partially resolve this issue. We also found that drawing and meshing our structure in 2D and then, if needed, extruding the mesh into 3D, seems to help with the meshing process. Firstly, the static solver is called to calculate the steady state equilibrium configuration of the structures due to intrinsic stress. This step is typically fast and reasonably reliable, but since our system can potentially produce very large deformation, the 'large deformation' toggle on the 'Application Mode' setting must be turned on. When the 'large deformation' is turned on, the program treats the system as nonlinear, and carries out several iterations of the solver linearizing the system around the previous run to achieve convergence. With the static stress calculated, we can then switch to eigenfrequency solver linearized around the static solution to obtain the system's eigen-modes and eigenfrequencies.

Information output

The solution produces the deformation field, which is the displacement vector of each point of the material from their initial position. The resulting stress pattern can be calculated from the field, as part of COMSOL's default set of variables it can extract. The eigen-mode analysis produces mode profiles and frequencies of vibration.

Appendix B

FDTD SIMULATION METHODS

The time-domain simulation of our devices are quite intuitive. We draw the structure in the program, illuminate the structure with a broad-band optical pulse, the program propagates the electromagnetic field in the simulation volume in some small time steps, and the simulation is terminated when all the excitation has left the simulation volume, absorbed by the Perfectly Matched Layer (PML). The majority of our FDTD simulation was done using the commercial *Lumerical FDTD* package. Some examples from our Lumerical code package can be accessed in the CaltechDATA archive; see reference [80]. The steps for such a simulation are as follow:

- Draw the structure geometry to be simulated, and specify its index of refraction.
- Place excitation sources and specify their frequency range.
- Specify 'monitors' that instruct the program to record information during the simulation run.
- Specify the simulation volume that properly contains all of the elements above.
- Run the simulation for sufficient run-time, typically the time it takes light to propagate through the simulation volume, but can be drastically lengthened if simulation contains slow-light devices or resonators
- Analyze the results from the monitors to extra information such as transmission, reflection, or field patterns.
- Additionally, for analyzing transmission into a waveguide guided mode, the projection of the transmitted in the mode in question is calculated.

The program Lumerical typically consumes relatively little computer memory, but is very demanding on CPU speed. The program has built-in parallel computing function and can distribute the simulation onto many separate threads for each core of the CPU. We note that the simulation time scales rapidly with the length of the simulation volume in the propagation direction, due to the increased load associated

with the larger volume compounding with the longer time required for the excitation to propagate through the volume.

B.1 Transmission and Reflection

The simulation of transmission and reflection using the FDTD method is relatively straight-forward, with perhaps a couple of details to keep track of, such as the mode projection.

Source placement

Lumerical FDTD provides a handy set of pre-made sources. For our purposes, we are interested in particular in the 'mode' source, where the program assumes the index refraction distribution at the source plane to be an infinite transverse-invariant waveguide of the same index distribution, and launches its eigen-mode. The directionality of the sources from Lumerical are typically very reliable, and non-source fields can propagate through the source plane without interference with the source mode. It is non-trivial, however, to launch the eigen-modes of non-transverse-invariant structures, such as our photonic crystal waveguide. The simulation of these non-invariant structures are typically done by launching light onto an unpatterned waveguide, then taper into the crystal section like the actual devices. We typically place the source $1\mu\text{m}$ away from the inner boundary of the PML, allowing some clearance and also spare space to place the reflection monitor.

Monitor placement

We typically place three types of monitors. First are the power monitors, whose function is to both record the field profile on them, and also to integrate the power transmitted through them for each frequency analyzed. These are placed $0.5\mu\text{m}$ from the injection side PML boundary for reflection, and same distance from the end of the simulation volume for transmission. Second are the field profile monitors, which we typically place on the symmetry planes of the simulation volume. Information such as local intensity and photonic crystal super-modes can be extracted from these monitors. Finally, we place an index of refraction monitor in the device layer, as a means to double-check for geometry errors.

Field projection into guided mode

Initially, we use the eigen-mode solving function of the excitation source to compute the guided mode profile inside the Lumerical program, which we record and save the field profile. We then let the simulation run until conclusion. The mode profiles of

the simulation run at the power monitors are also recorded, which generally would not be the same as the guided mode profile at those position. The measured power at the power monitors are then projected into the guided mode of interest by using the inner product of the the measured mode profile with the saved guided mode profile. This value is calculated for each frequency of interest. For newer versions of Lumerical, a built-in mode-projection monitor has been included to perform this projection. More information regarding mode projection monitors can be found on the Lumerical Knowledge-Base website.

Data analysis and application

The transmission and reflection data can readily be applied to rapid structure designing, allowing optimization of device geometries to maximize transmission and minimize reflection and loss. It is also helpful to calculate the sensitivity of these values to the various geometry parameters and potential misalignment between sections of the optical path. For photonic crystal devices, the frequency spectra demonstrates the band-edge resonances, and the recorded mode profiles from the profile monitors show the super-mode corresponding to the frequencies.

B.2 Green's Function Calculation

The primary purpose of calculating Green's function for our project is to calculate decay rate of a quantum emitter. The decay rate depends on the imaginary parts of the classical Green's tensor of the electromagnetic field. For high-symmetry points, which are typically our points of interest, such tensor is diagonal, and one simulation run is carried out for each principle axis.

Source placement

We use the dipole source for the injection of excitation. For each geometry, we place the dipole on each of the principle axes and run three separate simulations. The symmetry setting of the simulation volume needs to be correctly adjusted as the dipole orientation is changed. A full 3-symmetry-plane system cuts the simulation volume by a factor of eight, which greatly increases the simulation speed. Since the three orientations of the dipole are independent simulations, it is advantageous to run them in parallel on multiple simulation resources.

Monitor placement

The imaginary part of the Green's function can be calculated directly by monitoring the field at the dipole source, where the real part is divergent but the imaginary part

is finite, or can be calculated from the total power radiated by the dipole per unit frequency. The Lumerical FDTD package provides in-built functions that track the dipole power, and also a calibration function that provides the dipole power that would have been at the frequency with uniform index background. It is also helpful to keep track of the field patterns generated by the source as function of frequency and dipole orientation, since the propagation of a photonic crystal structure can be highly anisotropic.

Data analysis and application

The 'enhancement' of the Green's function imaginary component is calculated from the ratio of the dipole power measured to reference dipole power function. For rapid designing of a device, this ratio is typically maximized (or minimized for suppression devices) as function of geometry. The enhancement ratio can then be converted to the classical Green's tensor values, and then to the (F, m_f) state-dependent decay rates for atoms. Analysis of the emission pattern can potentially yield information of the Brilliuon zone, such as the actual profile of the 2D band-gap or the quasi-guided modes above the light-line, in compliment to the FEM results.

Appendix C

DIMENSIONS FOR CHEMISTRY PROCESSING HOLDER

The following figure contains the detailed dimensions, in inches, for our Teflon wet-chemistry holder.




Figure C.1: Dimensions for Chemistry Processing Holder

Appendix D

TYPICAL DESIGN AND FABRICATION PROCESS FLOW

Here we present a step-by-step typical process example, from optical design to installation into atom-light experiment. The goal of this chapter would be to provide a recipe for future readers, should they decide to fabricate devices using the platform developed in this project. A flow chart is included; see Figure. D.1. For the simulation steps, the designing of the embedded slot waveguide, as described in Chapter 7, will be quoted as an example. Our pattern generation code package can be accessed in the CaltechDATA archive; see reference [80]. The following fabrication and assembling processes are general to all of our devices.

1. Establish the functional goal of a design

- We hope to improve the atom-light interaction performance beyond APCW, by suppressing free-space decay rate component of Γ' .
- Our previous study indicated the 2D photonic band-gap of a photonic crystal slab can suppress the local density of state, and hence the Γ' value for the TE polarization.
- The goal is then to create an 1D waveguide embedded in 2D photonic crystal slab, with sufficient degrees of freedom to achieve band alignment to Cesium resonances similar to the APCW.
- Determine that the 1D defect is to be parallel to the K direction of the 2D lattice, due to Brillouin zone projection properties of the 2D lattice.
- Determine that a vacuum gap design similar to the APCW will be used to create sufficiently strong electric field to interact with atoms

2. Experiment with photonic design geometries to verify desired properties.

- Numerically experiment with types of defects to add. We calculated the dispersion and mode profile using FEM calculation, to experiment with cutting the 2D lattice at different positions, adding extra material in the 2D lattice cuts, and deformation of the inner-most hole. We confirmed that combination of these can indeed create two 1D guided modes confined to this defect, with roughly correct sized band-gap.

- We verify using FDTD that a guided mode can be coupled into this slot waveguide. Several tapering schemes were experimented, until the method used in Section 7.4 was developed and found to be sufficiently efficient.
- We carry out FDTD Green's function simulations to check if adding the defect to the background 2D lattice can still benefit from the Γ' suppression properties of the 2D crystal, and also the slow group velocity Γ_{1D} enhancement of a properly aligned 1D waveguide.
- We observed an anti-symmetric mode that creates undesired decay channel in the band-gap of the symmetric modes.

3. Numeric optimization of optical design.

- We find the geometry parameters that would place the upper and lower bands of the waveguide in the vicinity of Cesium D1 and D2 lines. We also calculate the tuning rates of these bands versus the geometry parameters.
- By using the calculated tuning rates, we managed to shift the undesired anti-symmetric mode away from the upper band of the symmetric mode, allowing a true band-free range in its vicinity.
- We calculate the Green's function to verify that Γ' suppression is achieved near the aforementioned band-free range that covers the upper symmetric band-edge. We then search around this set of parameters to find the strongest suppression to Γ' .
- We determine the remaining support structures, such as the lengths of tapering, and the number of lattice sites for the cladding 2D photonic crystal that would be sufficient to prevent light leakage.
- We use simulation to estimate the level of influence from fabrication error. Here we checked the vacuum gap deformation potentially associated with the 2D crystal slabs being strained by their support structure.
- We calculate a set of parameters that are pre-corrected for post-release tuning using 10nm thickness of ALD-grown Al_2O_3 .

4. Establishment of sets of 'Nominal' parameters

- We gather the information from simulation to determine a set of target parameters as an initial guess to fabricate the actual device. Separate

sets of parameters are created for non-ALD chips and ALD pre-adjusted chips.

- We determine the proper parameter sweep range for each of the geometry parameters. For initial test chips, large parameter sweeps are needed.
- The nominal parameters are updated using feedback from optical measurements of previously fabricated devices. After a few iterations, the nominal parameters will become finalized. We then do finer parameter sweep for actual chips designed to be used in atom-light experiments.

5. Keep track of difference between design dimension and actual dimensions ('Magic Factors')

- There are many sources that can create discrepancies between the dimensions in a design file and the actual resulting design. With sufficiently consistent processing, we typically see a roughly $\sim O(50nm)$ shrinkage of the device dimensions compared to the design dimensions, with a $\sim 5nm$ fluctuation within a beamwrite batch, and some level of long-term drift.
- We keep track of these shrinkage 'Magic' factors for critical dimensions of our device and pre-compensate for them during the design phase.
- Sometimes abrupt change of these 'Magic Factors' occur. In our experience, this was often caused by lithography resist contamination or aging, or sometimes inappropriate chamber conditioning of the plasma etcher.

6. Creating lithography write-files

- We calculate the 'write dimensions' that are our nominal parameters pre-compensated by the 'Magic' factors, as the nominal dimension for a design. Typically we make linear sweep of parameters around the nominal design parameters across the devices on a chip.
- A Matlab script takes these parameters, then create the shapes of the structures on a chip. This is done by printing AutoCAD script commands with the Matlab file, where the files contain global information such as creation of Layers and local information such as the poly-line and arcs that numerically draw the shape of the structures.

- We run the script in AutoCAD to generate a CAD drawing (.dxf) file that contains the design. These files are also convenient for later use of keeping track of the device types and dimensions of each chip.
- We use the software (GenISys BEAMER) provided by our cleanroom facility to convert the CAD files into readable format for our electron beam lithography machine (.gpf). Such software allows us to write 'flow files', namely a flow chart of several subroutines included in the software, that process the CAD files. We carry out the separation of coarse and fine exposure layers, as described in Chapter 3, using a customized flow file program. This program also allows specification of beamwriting precision (which with beam current determines writing time).
- This BEAMER program converts the CAD file into a set of trapezoids that the lithography machine actually writes, and also the placement of main-fields. We found that the settings used for this 'fracturing' of CAD pattern into trapezoids, such as the directions of the trapezoids and the way the code deals with circles and curves, can sometimes influence the outcome of written shapes. Some trial-and-error may be required.
- With the .gpf files ready, we can proceed to do the electron beam lithography.

7. Electron beam lithography

- We do our lithography using Caltech KNI cleanroom's Vistec-Raith EBPG 5000+ lithography machine. This machine has 2.5nm minimum step size, 50MHz maximum operation frequency, and a current range of 100pA to 200nA. In KNI, the machine is kept at 100kV acceleration voltage.
- We prepare our chips as follow:
 - a) Pre-clean the substrates with acetone, isopropyl alcohol, then deionized water.
 - b) Clean the chips further with Piranha ($H_2SO_4 : H_2O_2=2:1$) solution for 10 minutes.
 - c) Thoroughly rinse the chip with deionized water, then blow dry with dry Nitrogen.

- d) Spin-coat electron beam resist (Zeon Chemicals ZEP520A) on the chips. This is carried out at 4000rpm spin speed with 2500rpm ramp up. We keep the chip spinning for 60 seconds.
- e) The spun chips are baked on a hotplate at 180°C for 2 minutes.
- With an on-site microscope that has a stage identical to that inside the lithography machine chambers, we pre-align our chips to a lithography machine chip holder. We carefully align the chip edges to the coordinate system of the stage. For some sets of wafers we used, their dicing precision was insufficient, and we have to 'pre-rotate' the chips at this stage to compensate for the error in dicing direction.
- We load the chips into the lithography machine, and begin writing. The beam-writing typically takes approximately one hour per chip.
- The exposed chips after the lithography are developed using the resist developer supplied by the resist company (Zeon Chemicals ZED-N50), carried out at room temperature for 2 minutes 30 seconds, then rinsed using micro-filtered Methyl-Isobutyl Ketone (MIBK) for 30 seconds, and then blow dried with dry Nitrogen.

8. Plasma etching pattern transfer

- We pre-condition the chamber of the plasma etcher (Oxford Instruments Plasmalab 100) for our recipe (C_4F_8 : SF_6 = 25sccm : 12sccm at 15mTorr. 1000W ICP power and adjust RF to get a DC bias of 82V, etch time 3 minutes) for a total of 6 minutes, to stabilize the chemistry coating of the etch chamber.
- We hold our chips on Silicon carrier wafers, with a small amount of thermal conduction paste to provide adhesion and thermal conduction. We experimented with up to having three chips on the same wafer in the same etch run to no ill effect.
- The etched chips are cleaned as follow:
 - a) Pre-clean the substrates with acetone, isopropyl alcohol, and then deionized water.
 - b) Clean the chips further with Piranha (H_2SO_4 : H_2O_2 =2:1) solution for 10 minutes.
 - c) Thoroughly rinse the chip with deionized water, then blow dry with dry Nitrogen.

9. Pattern check and 'Magic' factor measurement

- The cleaned chip after the plasma pattern transfer would have the device dimensions transferred into the Silicon Nitride layer. With resist removed, we can inspect the chips in an SEM (Hitachi S4300).
- We first check to make sure the pattern is correctly defined. Namely, look for user errors during file preparation, contamination of chip substrate (which will be imprinted into the Nitride layer by the etching), or other unexpected issues.
- We measure the 'Magic' factors for the critical dimension, such as the coupling waveguide size, and generally the geometry parameters of the photonic crystals, if any.
- The 'Magic' factors are compared with their previously measured values to keep track of any long-term drifts. Care must be taken particularly when switching batches of resists, or carry out aggressive cleaning of plasma etching chamber.

10. Wet chemistry processing

- a) Load chips into the Teflon holder. We typically load them with the patterned side facing inward toward the holder, in hopes of avoiding accidentally scratching and damaging the devices during handling.
- b) Etch the Silicon substrate until through-windows on all chips are completely clear. The nominal recipe is a 73°C 30% KOH in water solution for 2 hours 30 minutes. A magnetic stir bar is placed and spun at 200rpm. Note that the Teflon ring of the holder should be removed during the KOH etch to allow liquid flow, and also visual access to see the chips.
- c) Prepare a clean bath of deionized water heated to above the KOH temperature. Once the KOH etch concludes, unload the holder and rinse it with the heated water for 30 seconds, by opening the Teflon ring on the holder submerged in water. Then, rinse again with another bath of room-temperature water.
- d) Transfer the holder to NanoStrip for cleaning. We typically use NanoStrip at 85°C and process for 20 minutes. We have observed no ill effect for running this cleaning step longer.

- e) Rinse with room-temperature water using 3 clean baths of water for 30 seconds each.
- f) Transfer the holder into a beaker of water suitable for Hydrofluoric acid processing. For our case we used Nalgene beakers. We put the holder in with additional 70mL of deionized water. Then, 20mL of 1:10 Buffered HF is added to the beaker. This cleaning step then lasts for 2 minutes.
- g) Rinse with room-temperature water using 3 clean baths of water for 30 seconds each.
- h) Rinse with high-purity isopropyl alcohol using 3 clean baths for 30 seconds each. These steps are for removing water, so the holder can be transferred into the critical point dryer.
- i) Transfer the holder, now submerged in isopropyl alcohol, into the critical point dryer. We found that one should use as little isopropyl alcohol as possible in the critical point drier, as long as it is sufficient to cover the chips. The CPD machine exchanges finite CO_2 volume, and excess alcohol will result in residue accumulating on the devices, causing them to stick together. We use a Leica CPD300 machine, with 30 cycles of CO_2 exchange and maximum heating temperature setting ($40^\circ C$). The CPD run takes approximately 3 hours.
- j) Once the CPD run concludes, the chips would be in air, and can be carefully unloaded from the holder. We found that for our holder design, this is most easily done by removing the Teflon pins first, then use a suitable tweezers to remove the now-loose chip.

11. Post-process dry cleaning

- We typically check the released chips optically at this stage regarding device yield, and the cleanness condition. With a sufficiently good microscope and some practice, it is also possible to confirm whether the various small gaps have stayed open.
- The chips are then cleaned using Oxygen plasma using a Plasma-Therm RIE system, with a recipe of 20mTorr processing pressure fed by a 20sccm Oxygen flow. This machine is not equipped with ICP function, and we operate only with an RF power of 80W and the resulting DC bias of 140V. Our typical process time is 4 minutes.

- The Oxygen plasma cleaning process is carried out typically twice on each chip. Once after the wet chemistry process, and once after the ALD coating, using the same recipe.
- After all dry cleaning is concluded, we package the chips in plastic chip trays and sealing the trays with Parafilm for transferring outside cleanrooms and intermediate storage.

12. Optical characterization

- Since the optical characterization would leave a chip exposed to air for an hour or more, it is desirable to do this in a suitably clean environment. Our setup was a special optics table with protective curtains and HEPA filters.
- For optical characterization, we place the chip on the Macor holder, and then temporarily secure it down with a copper clip. Note that excess pressure at this stage due to said clip could cause the chip to fracture.
- If a chip is designed for free-space coupling, they need to be cleaved at this point.
- With the chip on the characterization setup, we typically align the chip rotation by placing the input cleaved optical fiber in the chip's groove, then inspect if the fiber and the grooves are aligned to each other. This is most easily done by comparing the distance from the fiber surface to the edge of the V-groove on either sides of the fiber. Rotate the chip using its rotation stage until they are aligned.
- Once the chip is aligned in rotation, the back-side fiber can also be lowered into the V-groove, if needed for transmission measurements. Here we assumed the front- and back-fibers have been pre-aligned to be parallel.
- Use the characterization setup to measure the reflection and transmission of the device, as described in Chapter 5. Move to the subsequent device after the measurements finish. Repeat until every device is measured.
- The optical measurement results can now be used to determine the quality of the chip, and the number of cycles required for the ALD process.

13. ALD tuning

- We carry out our ALD process using the Ultratech Fiji F200 G2 ALD system. We make use of the default process supplied by the manufacturer for thermal (instead of plasma-enhanced) process using trimethyl Aluminum and water as reaction precursors at 250°C substrate temperature. This process gives us approximately 0.11nm per cycle.
- The number of cycles needed is simply calculated as desired frequency shift divided by tune rate.
- Once the ALD process concludes, the chip should be Oxygen-cleaned and optically characterized again to confirm the frequency placement.

14. Chip assembling and fiber gluing

- Here we assume the thermal-cure glue is used. The curing process would require the heating of the chip holder and stage. A throughout pre-cleaning using the typical vacuum cleaning process is carried out on all holder parts, using a sequence of solvent cleaning in a ultrasonic agitator with acetone, methanol, and isopropyl alcohol, each for 15 minutes at room temperature.
- With the cleaned and dried parts ready, we assemble the holder platform in the characterization setup. Using a ceramic (Ohmic) heater and a DC power supply, we heat up the holder stage, while monitoring its temperature with a thermal-couple.
- The temperature of curing was typically operated at 100 to 125°C, where a few-minute cure time is achieved. It is also possible to pre-cure the thermal glue by heating it to partial curing before cooling it back down. Care must be taken here, as the glue's viscosity changes drastically with temperature. Overly pre-cured glue can become difficult to apply.
- We first glue the chip to the holder. It is preferred to use the Macor holder instead of aluminum for the thermal glue due to thermal expansion coefficient differences. The glue is applied by a hand-tool where we have taped a cleaned and cleaved-end fiber section to a suitable handle. Wait for the glue to completely cure, at which point it turns a dark red-brown color.
- Prepare a reservoir of glue near the gluing robot arm, but away from the chip. For our system, we move the motorized stage to drive the robot arm to the position of the glue, then record its coordinates. It will also

be convenient to record the coordinate at which the robot arm tip comes into the microscope view, but at a higher altitude than the chip.

- We then begin to install the fibers into the V-grooves as follow:
 - a) Load a clean anti-reflection coated fiber onto the translational stage in the characterization setup. The fiber is held down using two magnetic clips.
 - b) Place the fiber into the groove, and align the chip's rotation to the fiber, similarly to the usual optical characterization.
 - c) Optimize the reflected signal in said fiber, creating a pre-gluing optical spectrum during the process.
 - d) Instruct the robot arm to pick up a droplet of glue from the reservoir. Visually inspect the size of the droplet, wipe away excess glue with a clean hand-tool. Our experience was that the glue droplet size should be below 0.2mm at its widest point. Err on the low side for glue droplet size. A device can be repeatedly glued if the glue volume is insufficient.
 - e) Move the robot arm to the pre-recorded coordinate that is above the chip. It should now be possible to move the microscope focus upward to look at the glue drop. Use the microscope images at the two different focal planes as proxy to align the glue drop to the fiber in the V-groove. We found it is most efficient to place the drop approximately 1mm away along the groove from the starting point of the device, with the center of the droplet aimed at the gap between the V-groove and the fiber in the groove.
 - f) Lower the glue drop fiber onto the chip surface. The liquid surface tension should pull the droplet into the space between the groove walls and the fiber. With the correct amount of glue, the liquid surface should approach the end of the cleaved fiber, but not go around it or touch the device.
 - g) Wait 10 minutes for the glue to cure (adjust depending on temperature / pre-cure status). Then, use another length of fiber to lightly tap the back side of the glued fiber to confirm it has cured. The tip of the glued fiber should not move relative to the substrate.
 - h) Use the hand-tool described earlier to apply more glue to the back side of the glued fiber, where it touches down on the Macor holder.

This reinforces the fiber-chip system by offloading the tension force from the chip-fiber bound to the fiber-holder bound. More glue can be used at this point, and it is favorable to get the glue droplet to touch the Macor surface first, then smear it toward the fiber, to avoid the glue flowing along the fiber toward the chip.

- i) Wait sufficient time for the glue to cure. Since it is on the Macor holder now, the color change should be visible. Gently apply tension to the glued fiber again to confirm the glue has completely cured.
 - j) Unload the back-end of the now attached fiber from the translational stage. This side of this device is coupled successfully.
 - k) Take another optical measurement at this point for reference. Compare this trace with the pre-glue case to look for difference or damage.
 - l) Move on to the next device.
- Once all desired devices are coupled on one side, rotate the chip around using the rotation platform. This step should be done slowly and carefully so as not to accidentally pull on the glued fibers. Repeat the same gluing process for the opposite side of the chip.
 - Our typical holder has an aluminum arm connecting to the Macor holder, with which we attach the Macor piece to the chamber. We fold the fibers, now coming out on both sides of the Macor holder, to a pair of one-inch loops, as shown in Figure 2.14. Beware that the 780HP fiber can generate bending loss or physically break if the bending radius is too small.
 - The chip is now fiber-coupled and ready for experiment.

15. Installation into atom-light experiment chambers

- To install the chip assembly to the experiment chamber, we need to vent the 'Science' chamber, thread the fibers through their feed-through, and then place the chip assembly into the chamber.
- For vacuum chamber cleanness, keep the chamber at positive pressure related to the outside with dry Nitrogen. Control the Nitrogen flow carefully, and do not over-pressure the chamber if it were closed at any point.
- Before starting to work on the fibers, make sure that all vacuum flanges have their gaskets pre-attached. Once the fibers are threaded, they may

interfere with placement of gaskets on certain flanges and could make the installation impossible.

- Thread the fibers through the Teflon feed-throughs. If multiple-hole feed-throughs are used, make sure to fill in any unused holes with a length of fiber. The feed-throughs are attached to the vacuum system using a special welded flange with a Swagelok tube fitting. Do not tighten the Swageloks at this point. Remember that tightening of these Swageloks onto the Teflon pieces is irreversible.
- Finally, insert the chip assembly into the chamber. Monitor this from multiple viewports with care, to avoid the chip assembly from colliding with the chamber. The fibers may need to be pulled through the feed-throughs in the process to reduce the amount of loose fiber inside the chamber.
- Tighten the flanges normally, and tighten the Swageloks. For our case, we found one and three quarters turns of the Swagelok suffice to make a seal, but this may vary with the style of the Teflon pieces.
- The chip installation is complete. The atom-light interaction experiment may commence once the vacuum reaches sufficient level.

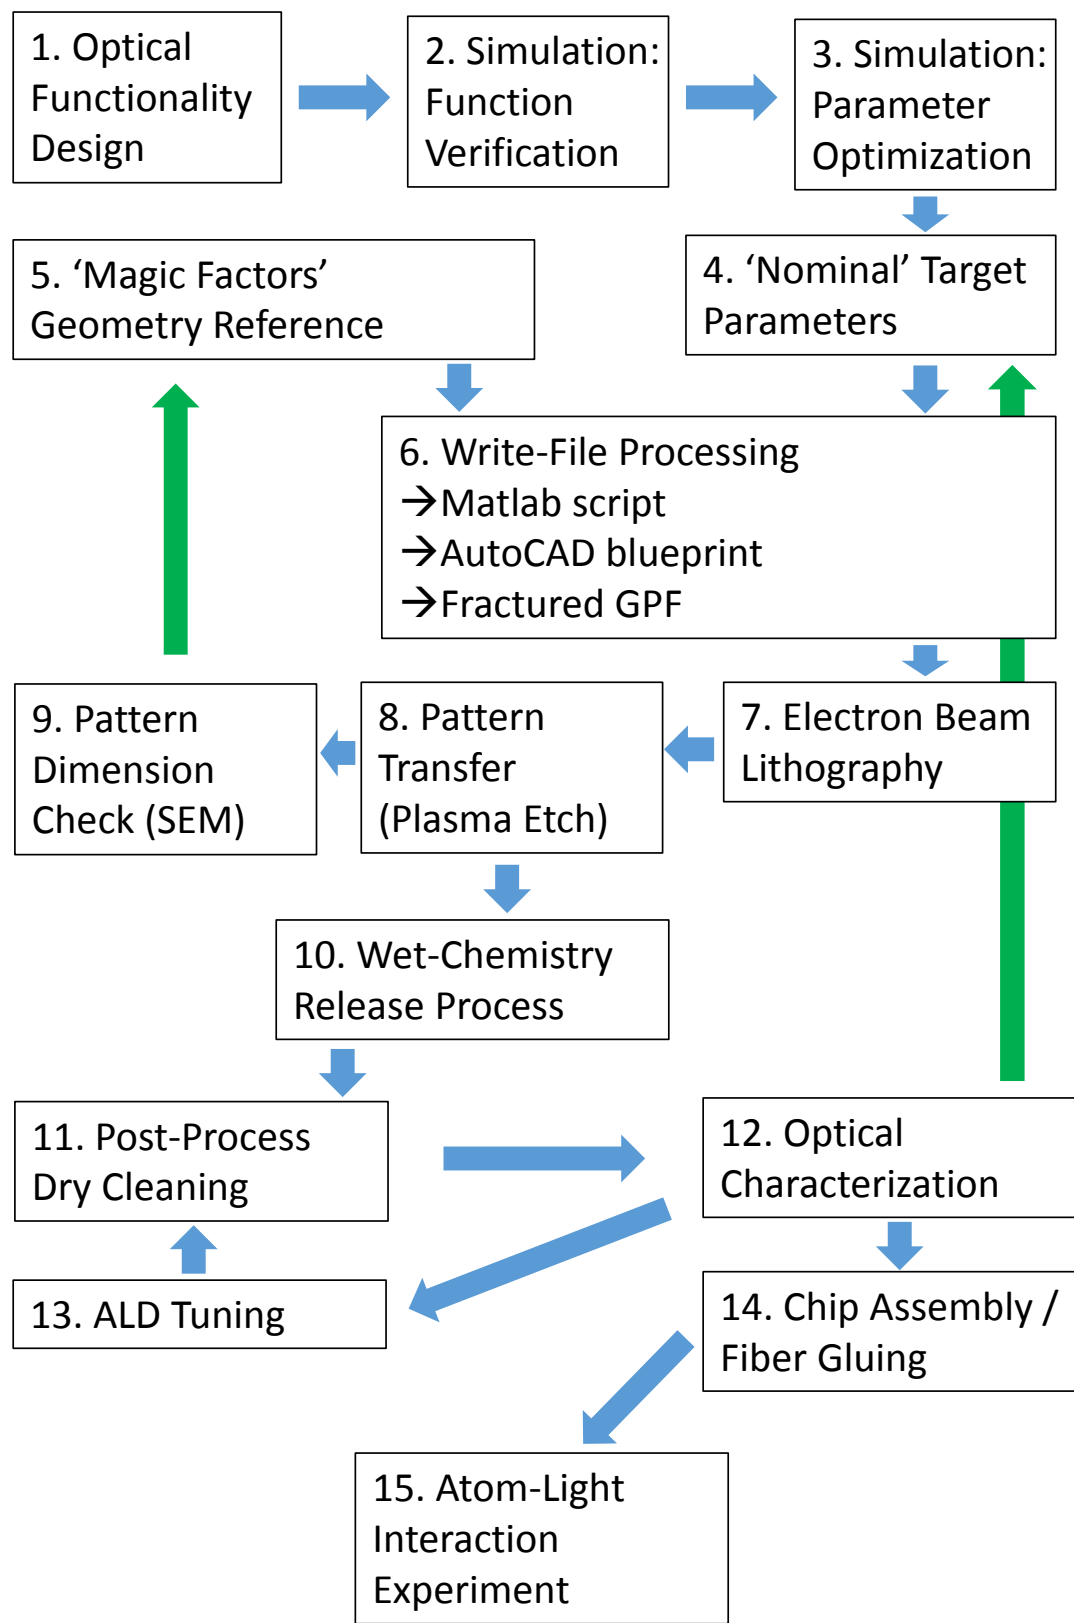


Figure D.1: Flow chart of a typical design-to-device process

INDEX

- acusto-optic modulator, 89
- adiabatic tapering, 6, 25, 57, 63, 73, 78, 98, 102
- annealing, 50
- anti-crossing, 53
- atom faucet technique, 88
- atomic layer deposition, 24, 49, 76
- avalanche photo-diode, 1

- band-edge, 52, 58, 61, 73, 75, 76
- band-gap, 52, 53, 58, 70, 74, 95, 98, 101, 103
- bit-line word-line configuration, 121
- Bloch boundary condition, 54
- Brillouin zone, 53, 63
 - edge time-inversion invariance, 54

- Casimir-Polder force, 52
- cavity, 2, 47, 58, 73, 77
 - finesse, 76
- Cesium, 25, 56, 60, 63, 68, 103
 - deposition on device, 85
 - source, ampule, 87
- chemical vapor deposition, 27, 48, 62, 78
- cleaning, 76
 - chemical, 43
 - Oxygen plasma, 45, 46, 77, 78
- cleave, 22, 33, 39
- COMSOL Multiphysics, 12
- cooperativity, 2
- coupling, 52, 68
 - alignment into free-space device, 71
 - enhancement, 74, 76
 - to external optics, 12, 17, 20, 38, 116
- critical point drying, 24, 44, 46, 78

- decay rate, 57, 90, 103

- single atom, 93
- superradiance, 93
- suppression, 99, 101
- device tuning
 - atomic layer deposition, 56, 60, 77, 104
 - capacitive, 117
 - geometry, 60, 103
 - thermal, 56, 94
- dipole, 57, 99
 - cancellation, 15, 17, 101
- directional emission, 63, 100, 118
- dose to clear, 29
- electro-optical modulator, 91
- electron beam evaporation, 49
- etching
 - plasma, 32
 - Potassium Hydroxide, 33, 34
 - surface quality, 78
 - test pattern, 33
 - timing, 36, 37, 42
- external cavity diode laser, 68, 89
- extinction ratio, 74
- fabrication, 24, 56, 70, 77
 - imperfection, 25, 76
 - requirement, 24
- fabrication imperfection, 56, 73, 74
- far-off-resonance trap, 25, 52, 100
- feedthrough, 20, 86
- figures, 4, 8, 10–16, 18, 19, 21–23, 26, 27, 30, 31, 34–38, 40, 41, 43, 45–48, 51, 54, 55, 59, 64–67, 69–72, 75, 80–84, 87, 90, 92, 94, 95, 105–115, 118–121, 136, 149
- finite element method, 12
- finite-difference time-domain method, 13, 56, 93
- Gaussian beam, 12, 13, 17
- glue, 18, 47

- cure, 19
- Green's function, 52, 57, 94, 99, 103
- group velocity, 52–54, 57, 63, 73, 93, 95, 99
- guided mode, 117
- holder, 20, 41
- hybrid system, 3, 24, 116, 121
 - design target, 6
- interaction, 53
 - coherent, 94, 117
 - enhancement, 103
 - light-mediated, 3, 92
 - with atom, 91
- lithography, 24, 28, 39, 48, 60
 - main field, 31
 - proximity effect, 29
- local density of state, 2, 55, 57
- Lumerical, 6, 13
- magneto-optical trap, 25, 87, 93
 - throwing, 89
- metal deposition, 48
- mode profile, 52–54, 60, 61, 104
 - air mode, 53, 98
 - anti-symmetric, 103
 - dielectric mode, 53, 98
 - intensity enhancement, 57
 - projection, 14
 - symmetry, 62, 100
- nanophotonics, 1, 2, 52, 116, 117
- Nanostrip, 44
- numerical aperture, 12, 21, 22, 117
- optical characterization, 68, 104
- optical fiber, 11, 12, 17, 21, 62, 69, 74, 116
 - placement and alignment, 70

- polarization maintaining, 69
 - splicing, 71
- optical lattice
 - moving, 96, 100
- optical measurement, 56, 73
 - broad-band spectrum, 68
 - real-time calibration, 89
 - scattering, 76
 - spectroscopy, 93
 - time domain, 90
- optical spectrum analyzer, 68
- optical transistor, 3
- optical tweezers, 101, 117
- photonic crystal, 47, 52, 53, 56–58, 60, 68, 73, 77, 95
 - 2D slab, 97, 99
 - atom-light interaction, 91
 - hexagonal lattice, 98
- piranha, 43
- piranha solution, 28
- plasma etching, 24
- Poissonian distribution, 91, 93, 117
- polarization, 48, 52, 62
 - alignment to device axes, 69
 - circular, 63, 100, 118
- polarization gradient cooling, 88
- Potassium Hydroxide, 42
- propagation loss, 73, 78, 85
- quantum emitter, 1, 55, 116
 - strong coupling, 2
- quantum many-body physics, 117
- quantum network, 3, 116
- reflection, 68, 73, 76, 91, 102
 - anti-, 16, 71
 - atom spectroscopy, 89
 - reduction due to annealing on Al_2O_3 coated devices, 50

- residual, 15, 58, 72
- resist, 28, 38
- saturation spectroscopy, 68
- scanning electron microscope, 32, 43, 45, 73
- side illumination, 100
- side rail, 11, 35
- side-illumination, 61, 92
- Silicon Nitride, 8
- slice-grid, 30, 35
- spatial light modulator, 23
- stitching error, 25
- stress, 9, 35, 50, 97
 - balancing, 9
 - concentration, 9
 - deformation, 97
 - deformation prevention, 9
 - intrinsic tensile, 8
- substrate, 27, 40
- super-mode, 58, 59, 61, 93, 99, 117
- superluminescent light-emitting diode, 68
- superradiance, 93
- support structure, 6, 97
- surface tension, 18, 26, 41, 42, 44, 45
- tables, 17, 20, 23, 56, 104
- tether array, 9, 50
- thermal dissipation, 7, 11, 14, 49, 92, 116
 - limitation on guided mode power, 76
- through-put, 25, 29, 42
 - chip replacement, 86
- transmission, 13, 57, 68, 73, 74, 93
 - atom spectroscopy, 89
 - through cavity mirror, 48
 - wide unpatterned cross-beam, 37
- tuning
 - atomic layer deposition, 49

- capacitive, 49
- v-groove, 17, 34, 50, 70, 73
 - precision cleaving, 38
- vacuum, 7, 22, 71, 85, 116
 - dual-chamber, 86
 - for power testing, 76
 - requirement, 85
- Van Hove singularity, 73, 75
- waveguide, 2, 6, 31, 52, 53, 62, 74, 97, 98
 - coupling section, 15, 50, 71, 73, 77
 - mode matching, 12
 - slot, 102
- yield rate, 15, 24, 50, 118