Integrated Diamond Nonlinear Optics

Integrated Diamond Nonlinear Optics

a dissertation presented
by
Pawel M. Latawiec
to
The Department of Applied Physics

in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy
in the subject of
Applied Physics

Harvard University
Cambridge, Massachusetts
December 2017

Thesis advisor: Professor Marko Lončar Pawel M. Latawiec

Integrated Diamond Nonlinear Optics

Abstract

The past decade has seen an explosion of development in new devices, modalities, and architectures for advanced sensing, computation, and communication. This growth has been driven by the massive investments of large consumer-tech companies in data-hungry domains of customer interest. As these companies develop their products and seek a competitive edge, technologies which can provide advantage across application areas by making computation cheaper, or sensors more knowledgeable, or communication faster, are attracting more interest in and outside academia. Photonics, or the study of how to manipulate light, is a broad-based platform that inherently delivers on these promises. Although photonics has been a part of daily life since the invention of the laser, modern needs and aspirational designs which meet those needs far exceed the capabilities of the platforms which have traditionally constituted the photonics toolkit.

There are two specific areas where standard photonic materials are lacking - support for visible wavelengths and functionality for quantum optics. Currently, there are many implementations of systems which use visible light or quantum optics on the tabletop, but the drive to scale these technologies and deliver them to the consumer encounters a roadblock in their size and cost. Visible light sensors, for instance, could greatly benefit from conversion into an integrated photonics platform, where the optical components are printed onto a small substrate. However, the standard materials used for such systems (silicon and indium phosphide) absorb visible light, and other materials (like silicon nitride) do not have the attributes required for advanced functionality. More compellingly, there exist very few material platforms which can host the quantum defects required for next-generation technologies like quantum computation or secure communication.

The past decade has seen extensive research into an emerging material platform which can deliver on the ability to bring integrated visible photonics and quantum optics to the same chip - diamond. With a broad transparency region (from UV to the far infrared) and high refractive index, diamond makes an excellent material for run-of-the-mill photonic devices. More importantly, it harbors atom-like defects in its crystal lattice - the nitrogen-vacancy, silicon-vacancy, and germanium-vacancy centers, among many more - which provide a direct interface between the photons circulating within the diamond device and the quantum world. Taken together, these simple properties of diamond constitute a powerful approach towards implementation of advanced quantum technologies, such as repeaters for secure quantum communication or even, one day, computers.

Before that future arrives, diamond's capability must first be developed. Diamond has only been available in significant quantities for research use for about a decade, and as such, the processing technologies which are taken for granted in silicon or silicon nitride-based work must be redeveloped. Within the context of photonic devices, we discuss different fabrication strategies and their implications for device design. For visible light devices especially, we investigate methods to reduce sidewall roughness for etched structures.

Interfacing with diamond photonic structures also brings its own unique set of challenges due to its high refractive index and very small material footprint. To this end, we develop a method termed "loaded tapered-fiber coupling" to optically access large, free-standing diamond devices. For size-limited diamond resonators and waveguides, we also develop auxiliary optical waveguides to act as spot-converters for visible wavelengths, an improvement on previous techniques for telecom-range wavelengths.

Finally, by synthesizing these improvements in fabrication, we embark on a set of experiments in diamond microresonators. Nonlinear optics is an excellent testbed for such improvements, both for diamond's superlative properties in that regard and their high sensitivity to improvements in fabrication. First, we look at a process known as Raman lasing by constructing long path-length resonators.

This technique involves the shifting of an input photon by 40 THz via the interaction with the diamond's crystal lattice. The first demonstration in a diamond integrated device was shown with a pump wavelength of 1600 nm, with an output that was tuned over 100 nm and a low threshold. Leveraging the developments in visible photonics for diamond, a second Raman laser was demonstrated, this time at near-visible (720 nm) wavelengths. This too had a Stokes output which was tunable to over 100 nm and a very low threshold. Raman-specific effects such as polarization conversion were also investigated.

A cousin to the Raman effect, the Kerr nonlinearity plays a key role in so-called microresonator (or Kerr) frequency combs. We survey the possibilities and difficulties in implementation for a diamond platform, highlighting the competition between the Kerr and Raman nonlinearities. Although the outlook for a technologically relevant Kerr comb in diamond is ultimately pessimistic, avenues for improvement are highlighted. Namely, we investigate via simulation a related class of effects, known as supercontinuum generation. Taking advantage of subtleties of the Raman process in diamond, we can find a preferred geometry which enables generation of a broadband light spectrum natively in diamond.

Finally, we look at alternative applications of diamond, particularly in high-power systems. Due to its excellent thermal conductivity and power-handling capacity, optics fashioned purely out of diamond can show significant advantages compared to other materials. We end with a discussion of the design principles and results stemming from this work.

Contents

Author List

Portions of Chapter 2 appeared in the following:

- Pawel Latawiec, Michael J. Burek, Vivek Venkataraman, and Marko Lončar, "Waveguide-loaded silica fibers for coupling to high-index micro-resonators," Appl. Phys. Lett. 108, 031103 (2016)

Portions of Chapter 3 appeared in the following:

- Pawel Latawiec, Michael J. Burek, Young-Ik Sohn, and Marko Lončar, "Faraday cage angled-etching of nanostructures in bulk dielectrics," J. Vac. Sci. Tech. B. 34, 041801 (2016)
- Haig A Atikian, Pawel Latawiec, Michael J Burek, Young-Ik Sohn, Srujan Meesala, Normand Gravel, Ammar B Kouki, Marko Lončar, "Freestanding nanostructures via reactive ion beam angled etching," APL Photonics 2, 051301 (2017)
- Pawel Latawiec, Vivek Venkataraman, Amirhassan Shams-Ansari, Matthew Markham, and Marko Lončar, "Integrated diamond Raman laser pumped in the near-visible," Opt. Lett. 43, 318-321 (2018)

Portions of Chapter 4 appeared in the following:

- Pawel Latawiec, Vivek Venkataraman, Michael J. Burek, Birgit J. M. Hausmann, Irfan Bulu, and Marko Lončar, "On-chip diamond Raman laser," Optica 2, 924-928 (2015)
- Pawel Latawiec, Vivek Venkataraman, Amirhassan Shams-Ansari, Matthew Markham, and Marko Lončar, "Integrated diamond Raman laser pumped in the near-visible," Opt. Lett. 43, 318-321 (2018)

Portions of Chapter 5 appeared in the following:

- Yoshitomo Okawachi, Mengjie Yu, Vivek Venkataraman, Pawel M. Latawiec, Austin G. Griffith, Michal Lipson, Marko Lončar, and Alexander L. Gaeta, "Competition between Raman and Kerr effects in microresonator comb generation," Opt. Lett. 42, 2786-2789 (2017)

The following authors contributed work to Chapter 5: Yoshitomo Okawachi, Vivek Venkataraman,, Amir-Hassan Shams-Ansari, Mengjie Yu, and Austin G. Griffith
The following authors contributed to Chapter 6: Alex Muhr and Xiao Xiong

Listing of figures

For my wife Melisa Rojas

Acknowledgments

Neither my acknowledgements nor my Phd could start without the most important person of all, Marko Lončar. Marko has been the most amazing PI that any graduate student could ask for, whether through his unwavering support or his continuously churning dynamo of ideas and exuberance. I owe a lot to Marko, not only in helping me through my PhD, but also in lending his insight and commitment to helping me develop as a professional. As PI, he has the opportunity to shape the outlooks of generations of students, a task I know he takes seriously and thoughtfully, with more purpose than any other professor I have come across. Even more importantly, he sets a stunning human example for his students to follow. A large part of why I went to Harvard was this man, and the group that he leads.

I'd like to thank the rest of my committee: Evelyn Hu, Amir Yacoby, and Frederico Capasso. I have met with Evelyn and worked with her students many times - they always speak warmly about her as a mentor. Her perspective on fabrication and clear understanding of the issues we students face is refreshing. Likewise, I've heard similar great things about Amir. Although we did not interact as often, from meeting and collaborating with his students I get the same impression. Finally, Frederico has been a near-constant force during my PhD, from one of the first classes I took at Harvard. I

later served as a teaching fellow for him and have been in frequent contact with either him or his students. Frederico's perspective is always challenging and refreshing, and, most importantly, he forces you to be scientifically honest and confront your gaps in knowledge. This can be uncomfortable for many of us but I've learned through him it is always worth the effort.

Outside of my formal boss, there have been a number of mentors I've had the privilege of knowing throughout my career at Harvard. Vivek Venkataraman has been instrumental in helping me get to speed on a large portion of this research, particularly in nonlinear optics, and he has also been a great friend. In particular, I wanted to thank Vivek for his help building optics experiments and invaluable aid during the paper-writing process. He always had an insight into how to communicate something more clearly, a skill which I desire to reflect in my own writing.

Throughout my time I have also had the opportunity to serve a senior role for other students in the group, and I could hardly think of a more deserving student than Amirhassan Shams-Ansari. Besides his stellar work ethic, every conversation with him is a treat, and he may be one of the most genuinely kind individuals I know. I hope I learned from him in that regard.

There have been a number of members of my group who have impacted my research through collaboration or discussion. First of all, the senior members Mike Burek, Birgit Hausmann, and Raji Shankar were a great source of fabrication tips and insight. I've always admired Mike's organizational methods and keep them in the back of my head when designing some system. Birgit was also the person I learned the basis of diamond fabrication from, and many of the techniques I use today are ones originally taught by her. I've had a few collaborators from lab, like Boris Desiatov and Stefan Kalchmair, who provided motivation and excellent insight into what was going on inside our microresonators, as well as provided companionship during treks to conferences. Boris made an excellent hiking buddy in that regard. Closer to my cohort, I'd like to thank Haig Atikian for great conversations about research and the stuff surrounding research. A few sections from this thesis came from our direct collaboration. Anna Schneidman has also been great to have around in the

lab and run into around Harvard when not in the lab. Srujan Meesala has been a fantastic friend and positive force in the lab, I have definitely learned a lot of the more physics-side of diamond experimentation from him. I had a lot of fun with him and Young-Ik Sohn, both in and outside the McKay. Young-Ik has been an interesting person to talk to and I enjoy hearing his perspective. Cheng Wang and Mian Zhang have been great partners at conferences. Our research has some overlap so it's been great to see the different ways we approach similar issues. Both Mian and Cheng are very thoughtful and come to decisions very analytically, while still keeping in mind the bigger picture. One of the people I've learned the most from was Zin Lin, perhaps because our spheres of knowledge rarely intersected. There's no doubt that Zin thinks very deeply about things and brings an unbridled enthusiasm to any project, and, most importantly, is willing to share that enthusiasm to someone like myself. I'd like to thank I-Chun Huang for always excellent discussion regarding fabrication puzzles and his genuine presence in the lab. There's also a number of ambitious undergraduate and graduate interns who I have interacted with over the years, especially Jake Rochman and Xiao Xiong, who I wish success in their future endeavors. Jake worked on very early, exploratory investigations on a number of topics that were refined in this thesis, from chip-edge coupling for angle-etched devices to long path-length angle-etched resonators for Raman lasing. I'd like to highlight Xiao too, who really helped with a lot of the lens measurement and metasurface work.

There are also a number of collaborators outside our lab who have contributed meaningfully to the research presented in this thesis defense, especially Nathalie Picque's group at Max Planck and Alex Gaeta's team at Columbia. The direction of this research was certainly influenced by their insight. I also thank Carmen Palacios-Berraquero of Mete Atuture's group for a very fruitful collaboration. For someone who fabricated as many devices as I have, it is always a pleasant surprise when some of them just work on the first try. Stefan Bogdanovich from Ron Hansen's group deserves special mention as an excellent host and guinea pig for diamond fabrication knowledge transfer. In then end, one wants research to be useful to someone else, so I thank Stefan for volunteering to be one

of the first, as well as showing me around Delft. I would also like to thank our collaborators at Element Six, including Dan Twitchen, Matthew Markham, and especially Alex Muhr. Not only for an invaluable supply of research material (i.e. diamond) but for early funding and a close collaboration on a real-world project with tangible impact. It was always a pleasure to discuss research with Alex and think of new ways to use industrial-scale diamonds. In the end, I should also thank the team at the DARPA SCOUT program for funding a majority of the research contained inside this thesis, and, of course, for choosing to fund us in the first place.

I'd also like to acknowledge the Harvard Photonics club and its leadership, as a leader myself it's helped me to broaden my horizons outside of the lab bench, from engaging in educational outreach to children to organizing an entire photonics conference on Harvard's campus.

My time at Harvard has been enriched by a number of friends who I have met through the program or through classes, especially Rob Devlin and Daryl Vulis. It's been a pleasure to cross paths once again with many of my classmates from Caltech as well, including Will Steinhardt and Brian Penserini.

I should also acknowledge a number of organizations and people across Harvard as an institution which were instrumental in helping this thesis come to conclusion. Both Kathleen Masse and Stacia Zatsiorsky have been phenomenal secretaries for the group, and they have truly helped the lab and our schedules run much more smoothly.

I owe a debt of gratitude to the staff at CNS, who helped to keep the tools running so that us grad students could (ab)use them. They are truly one of Harvard's greatest assets. I also wanted to thank the custodial and cafe staff at the LISE and McKay buildings for their conversations and, of course, excellent work helping to keep our operations running smoothly.

Lastly, I thank my mom and dad, whose unending support I've leaned on heavily over these past few years. I know they have put a lot of energy into helping me achieve my goals, all the way back from when we first moved to the USA some 25 years ago. My siblings Matt, Caroline, and Tommy

have been a drive as well, and it's been fun watching them grow and learning from their experiences. Finally, I'd like to thank my wife Melisa, who constantly challenges me to improve and develop as a person. I can imagine no one I'd rather spend my time with. To the many adventures ahead.

1

Introduction

The new devices, architectures, and designs required of the next generation of consumer, medical, and industrial technology requires new inputs and fresh perspectives. Although photonics, the study and engineering of devices in which light interacts with matter, has been relevant ever since the invention of the laser, new frontiers and boundaries are being explored as the potentials of light as sensor, detector, imager, communicator, rangefinder, and computer open up. In this thesis, we will explore one material - diamond - which has tremendous potential for next generation optical

Property	Symbol	Value
Density	ρ	3.515e3 kg/m^3
Lattice type	-	Cubic (Fd3m)
Lattice constant	a_0	0.357 nm
Bandgap	E_g	5.5 eV
Toughness	n	2.4 -
Thermal conductivity	k	2200 W/m/K
Thermal expansion	α	1.05 ppm/K
Speed of sound[138]	V_s	12,000 m/s

Table 1.1: Table of standard material properties of diamond. Data from Colbran [42] unless otherwise noted. Where the property is dependent on temperature, the value reported is given at 300 K.

and photonic systems, and attempt to blaze a path for future researchers to study. This investigation will not be limited to developing methods and techniques for handling and processing this fascinating material, but will also extend to physics and devices never-before demonstrated in the diamond, or any, integrated optical platform.

1.1 Diamond

In the public consciousness, diamond has enjoyed an overwhelmingly popular role as the commercial gemstone of choice, likely owing to material properties which make it truly unparalleled. However, diamond has played and will continue to play an ever-increasing role in industrial processes, particularly those in which a tough cutting tool or abrasive is needed. Furthermore, its excellent thermal conductivity, again the highest of any material, has encouraged its use in new applications such as high-power electronics as a heat spreader.

The material properties of diamond are summarized in Table 1.1. Again, as an engineering material diamond has few rivals, and its limited use to-date stems mostly from its difficult manufacture. There are two different methods with which to grow diamond - high pressure/high temperature (HPHT) processes and chemical vapor deposition (CVD). The characteristics of synthesized dia-

mond from these two techniques vary strongly. HPHT is mainly used to grow diamond plates. In this technique, the growth chamber is pressurized and heated to simulate the earth's core in a process very much similar to the one that naturally forms diamond[42]. This has the advantages of being relatively cheap and scalable, capable of producing diamond at a quick rate. However, the diamonds from this process are often tinged yellow due to impurities in the growth chamber. CVD synthesis, then, is the preferred method to create research-grade diamond for photonics. Diamonds from this process are available in a number of grades, depending on their impurity concentration and types of impurities present. To create a diamond via CVD, first a seed crystal of diamond is placed in the chamber. Then, a plasma of methane and hydrogen mixed in precise ratios is struck. The methane acts as a carbon source, growing the diamond atom-by-atom, while the hydrogen acts as a pruning mechanism - removing any poorly-bonded carbon or impurities from the growth surface[163]. This process is much slower but can yield diamond with incredibly low impurity concentrations.

In diamond, impurities are not necessarily a bad thing and in fact make the basis for quantum optics in diamond. These impurities, or defects, take the form of a literal catalog of possibilities. The most well-studied is the nitrogen-vacancy center, or the NV. This defect has attracted significant attention for magnetic sensing[2] and quantum communication protocols[68]. The silicon-vacancy center (SiV) shows several distinct advantages over the NV, especially for quantum optics, and recent demonstrations have shown its tunability[160] and suitability in various protocols[144,158]. A newcomer is the germanium-vacancy, or GeV, which promises to even further improve on select characteristics of the SiV center[16]. Although a fascinating subject matter, for this thesis we will be focusing our efforts on the opposite end of the spectrum, or defect-free diamond, and push our technology on that front. It should be emphasized that advances in the general processing of diamond have direct applicability to the future devices these quantum emitters will inhabit.

1.2 Photonics in Diamond

The focus of this thesis will be the creation of integrated optical devices in diamond. Typically, these devices are made in more familiar material platforms such as silicon nitride, silicon, glass, or indium phosphide. Diamond, as we will see, has several compelling advantages over this traditional pantheon of optical materials, but techniques to process it are considerably less mature. What techniques exist typically do not serve to create the technologies that we are interested, at least not at the requisite quality. This goes doubly so for the area where diamond shows tremendous promise - visible wavelength integrated photonics. Of the material platforms mentioned previously, only glass and to some extent silicon nitride have the ability to perform adequately for visible wavelengths due to the other's material absorption. Diamond, absorbing essentially no light in this wavelength region, has a unique advantage here. However, the issues of integrated photonics - loss, coupling, measurement - are only exacerbated in at these challenging wavelengths.

For integrated optical devices, light is contained in waveguides of very small cross-sectional area, typically on the order of a few square wavelengths of light. The central conceit required to maintain this confinement and keep light moving happily on a chip is the presence of refractive index contrast. In short, the waveguiding material (in our case diamond) must be surrounded in all sides by other materials with lower index. For traditional integrated devices, such as silicon-on-insulator (SOI), the optical material (silicon) rests directly on top of a spacer material (insulator, or silicon dioxide). For a number of reasons, the situation in diamond is considerably less straightforward, and new geometries or techniques have to be proposed and studied. In this thesis we study two geometries in particular - one similar to the SOI platform where a diamond film rests on oxide, and another where the diamond is three dimensionally sculpted and nearly suspended in air, surrounded in this case by free space.

Waveguides alone often aren't enough to solely comprise a meaningful device architecture. Of-

tentimes, some sort of feedback or resonant enhancement is required to get devices to interact strongly with their constituent matter. A cutting-edge example of such a system is the photonic crystal, where a periodic arrangement of holes or nanostructure can effectively act as a mirror, causing light to bounce back and forth and build up inside a cavity. Another, more straightforward example is the microresonator, where a waveguide is simply wrapped around to meet itself. Microresonators will form the core of the experiments shown in this thesis, and we will discuss their physics at length.

1.3 Nonlinear Optics

To create functional devices, light has to interact with something - be it a detector or a laser gain material or a crystal. Nonlinear optics is concerned with this process in a general sense, specifically with the results of light-matter interaction. These nonlinearities may occur at the single-photon level, or they may only begin to become apparent at higher power levels. Using techniques from integrated photonics, we can build cavities which increase the interaction of the light inside with the material system. Coupled with feedback, we can hit regimes where a process goes above threshold and new properties of light are seen.

Diamond's single photon emitters, or quantum defects, can be seen in some way as a nonlinear component of next-generation, specifically quantum, optical systems. When light interacts with such an emitter, it is absorbed and re-emitted in a fundamentally nonlinear way. This interaction can be used for a wide variety of applications, including sensors, repeaters, or communication. However, because of the atom-like structure and single-photon interaction, the use case of such systems can be inherently limited to low-power and low-bandwidth applications. There are, however, other material properties inherent to diamond and nonlinear phenomena which may be used to provide interesting functionality. Although diamond has been little-explored in these regards, we will show throughout the course of this thesis its tremendous promise.

Although there exists a wide range of schemes and useful systems taking advantage of nonlinear optics, we will focus on a few which can be classified into Raman- and Kerr-mediated. Even though Kerr-mediated and Raman-mediated effects can never be truly separated, there exist regimes where one forms the dominant explanatory mechanism or classification. Raman-based devices rely on the scattering of a photon off of an optical phonon in the crystal medium to convert the wavelength of the scattered photon, known as Stokes scattering. If the process occurs inside a resonator, it has the opportunity to feed back onto itself - resulting in Raman lasing (or, alternatively, stimulated Raman scattering) where more scattered phonons are being generated and encouraging the scattering of more light. This can produce very high powers of Stokes-shifted light, to the point where these types of lasers are technologically relevant in converting light for medical and industrial applications, particularly at visible wavelengths.

Kerr-mediated devices rely on the electronic nonlinearities of the crystal the light propagates through. This interaction causes a coupling term to occur between different frequencies of light and is directly dependent on the intensity of light. Seen at its most basic, it is the modification of the refractive index of light n to a constant component n_0 and an intensity-dependent component $n_2 I$:

$$n = n_0 + n_2 I \tag{1.1}$$

Though conceptually simple, the implications of this concept lead to a wide range of useful applications like four-wave mixing and supercontinuum generation, which we will explore in later chapters.

Outside of these named nonlinear mechanisms, there of course exist others (such as the second-order nonlinear effect and the electro-optic effect, among others). We do not discuss them because they are not prevalent in diamond, mainly due to diamond's crystal symmetry. Because the crystal itself is isotropic, it cannot possess such second-order nonlinear effects. However, diamond's crystal

Property	Symbol	Value
Refractive index	n_0	2.4
Temperature sensitivity	$\mathrm{d}n/\,\mathrm{d}T$	9.6e-6 /K
Nonlinear refractive index[104]	n_2	1.3e-19 m^2/W
Nonlinear susceptibility[104]	$\chi^{(3)}_{1111}$	5e-20 $\mathrm{m}^2/\mathrm{V}^2$
Nonlinear anisotropy[104]	$\rho \equiv 3\chi^{(3)}_{1122}/\chi^{(3)}_{1111}$	1.2
Raman frequency shift[123]	Ω_R	40 THz
Raman gain[123]	g_R	15 cm/GW
Raman linewidth[123]	Γ_R	60 GHz
Electrostrictive constant[138]	γ_e	15 cm/GW
Brillouin linewidth[129]	Γ_b	20 MHz
Brillouin gain (est.)	g_b	1.5 cm/GW

Table 1.2: Table of optical material properties of diamond. Data is reported at room temperature and at 1064 nm probe wavelength. All data taken from Colbran [42], unless noted otherwise.

lattice is precisely the reason why it has such excellent properties otherwise. Because it is made of a single species of atom, there are no dipole moments in its unit cell, meaning that low-frequency light cannot be absorbed by the crystal. This makes diamond one of the best materials for infrared optics, which we will touch on later in the thesis. However, because the unit cell nonetheless consists of two atoms, the lattice supports phonon modes which can couple to and scatter light (the so-called optical phonons).

Table 1.2 lists some of diamond's linear and nonlinear optical properties, which are used and referenced throughout. Included are values for the nonlinear electronic susceptibility, the properties related to stimulated Raman scattering, and to motivate future interest, the values relevant for stimulated Brillouin scattering. In these parameters there are a number of points of interest which set diamond apart from the materials commonly used in photonics today. For instance, diamond has superlative properties for Raman interactions. Not only is its Raman frequency shift the largest of all known materials, it also has the highest Raman gain. This is taken as a direct result from its crystal lattice, where the constituent carbon atoms are tightly packed and very light, resulting in

high-frequency optical phonons. Taken together, this places diamond as the material of choice for Raman laser applications. From a passive-optics point, diamond possesses the rare combination of a high refractive index and a large bandgap. This means that diamond maintains its transparency over a large wavelength range, but, more importantly, it also maintains a naturally low material dispersion over wavelengths of technological interest. With the high refractive index contrast, this means that diamond waveguides can be easily engineered to provide the preferred dispersion where needed, without resorting to higher-order modes or mode-mode interactions. This is a useful property for Kerr frequency combs and other applications, especially as the search for a compact, visible broadband source continues. More mundanely, these properties are necessary conditions for four-wave mixing and supercontinuum generation to occur, which we report on in this thesis. From a broader perspective, the threshold-like behavior and high sensitivity of these processes to device quality make these excellent testbeds for fabrication improvements.

1.4 Outline

After introducing diamond, we will first develop the concepts behind the most fundamental micro-optic devices - resonators and waveguides. In particular, we will take the time to highlight the issues and concerns with implementing these types of devices in diamond to ground the discussion in future chapters. Finally, we will develop different platforms for interfacing with diamond devices, including a discussion for why these platforms need to be developed. For diamond-on-insulator devices, we discuss how mode-converters are used to design around the difficulties inherent in diamond's small size. For angle-etched devices, especially those where the optical mode is well-confined in the diamond itself, we develop a technique known as loaded tapered-fiber coupling which aids in the probing of such microresonators. Both of these coupling techniques are broadly applicable to other material platforms.

In the next chapter, we discuss approaches to the fabrication of diamond. The sections will detail the various ways diamond optical devices can be fabricated, staring with a detailed investigation of Faraday cage angled etching. Diamond devices were first made with Faraday cage angled etching and then later transitioned to the reactive-ion beam undercut etching technique. This technique proved more robust and possessed much finer control. Then, fabrication techniques for the diamond-on-insulator (DOI) platform are presented, with an emphasis on a step-by-step accounting of the fabrication process. This section details the intricacies in diamond fabrication and discusses the methods used to attain record-setting quality factors in diamond photonic devices. Importantly, the difficulties in processing will be discussed to provide heuristics which will aid further investigation. Finally, we include a discussion of mode-converter fabrication for high-power applications.

Having laid out our basic architecture and fabrication process, we move on to experiments done in diamond towards Raman lasing. After providing a historical overview and contextualizing these developments, we present a telecommunication-band Raman laser implemented in diamond microresonators. Next, we present an evolution of this technology towards a Raman laser functioning at visible wavelengths. This demonstration synthesizes most of the fabrication and theoretical efforts developed in this thesis and shows excellent performance. The diamond resonators presented show the highest quality factors across all diamond devices at visible wavelengths, and are the shortest-wavelength CW-pumped integrated Raman lasers ever demonstrated (to the knowledge of the author). Owing to the clean experimental setup, we also investigate polarization conversion effects due to Raman in the diamond device.

Next, we investigate Kerr nonlinearities in diamond photonics. We begin our discussion with supercontinuum generation in diamond, showing the potential for broadband generation well into the UV. We highlight potential architectures for such supercontinuum generating waveguides. We then discuss the potential for a frequency comb in diamond at visible wavelengths. Motivation and simulations to this end are presented, followed by an in-depth discussion of the subtleties and diffi-

culties involved in crystalline resonators. Although the prospects for frequency comb generation in diamond are ultimately pessimistic, a full accounting of the system model reveals interesting avenues for further exploration.

Before concluding, we take a look at other devices developed during the course of this PhD for high-power optics. Specifically, we develop the theory behind meta-optic devices and show implementations in diamond for different wavelengths. We focus initially on monolithic anti-reflective coatings, followed by a quarter wave plate using form birefringence, and lastly on a metasurface-based focusing lens, including characterization results.

Lastly, we conclude the thesis and include some outlook on the future of diamond photonics. Throughout, we have emphasized points of advantage of diamond in nonlinear and high-power systems. To finalize, we will discuss how these developments extend beyond this corner of photonics and apply to emerging systems such as integrated quantum optics in diamond.

2

Diamond integrated optics

Benchtop optical systems, some even including diamond as a component, have served a variety of uses. Mostly, because of their easy extensibility and customization, they form a core component of optical research in the lab. The drive for miniaturization and the packaging of many functional devices onto as small a form factor as possible, however, precludes their use for serious commercial-scale deployment. To this end, integrated optics has promised to bring the silicon chip revolution to the world of optics. By squeezing components down to scales comparable to the wavelength of

light, not only is light-matter interaction enhanced, but other gains such as lower power consumption, smaller size, and cheaper production realized. Additionally, because of the fabrication techniques used in making these devices, there is no post-construction fine-tuning of systems as occurs in benchtop systems, and the performance of these devices remains relatively stable without constant maintenance.

In this chapter, we will give an overview of the physics of integrated optics, particularly in relation to the device architectures which will be employed in the rest of the thesis. As the main figure of merit throughout is the quality factor of the constructed resonators, we will briefly analyze and discuss sources of loss and the theory underpinning them. Finally, we will conclude with an overview of the different ways to optically interface with the constructed devices. A new technique with which to do so will be presented, which addresses the challenges of coupling to certain diamond optical device architectures.

2.1 Whispering gallery mode resonators

The central object of our study is the whispering gallery mode resonator. In this section, we will develop, derive, and enhance our intuition of the physics of these systems.

2.1.1 Sources of loss

Sources of optical loss form perhaps the biggest and hardest-to-tackle challenge in any micro-optical system. But, because thresholds for many nonlinear processes scale as the inverse of $\sim Q^2/V$, we immediately see that bringing the quality factor Q as high has possibly reduces our threshold by a factor of Q^2. To wit, we can write the total loss of any linear optical mode as the sum of its components[17]:

$$\alpha_T = \alpha_{\text{s.s.}} + \alpha_{\text{s.a.}} + \alpha_{\text{rad.}} + \alpha_{\text{b.a.}} + \alpha_{\text{coupling}} \tag{2.1}$$

Where $\alpha_{\text{s.s.}}$ is loss due to surface scattering, $\alpha_{\text{s.a.}}$ is loss due to surface absorption, $\alpha_{\text{rad.}}$ is radiative loss, $\alpha_{\text{b.a.}}$ is loss due to bulk absorption, and α_{coupling} is loss to any other coupled optical modes, be they within the same resonator or from a separate waveguide which interfaces with it. This last term is not usually included in discussion of loss but it is phenomenologically consistent and should be included nevertheless. To convert from loss (with units 1/m) to quality factor (a dimensionless parameter), one can use the following formula[132].

$$Q_T = 2\pi n_0/(\lambda \alpha_T) \tag{2.2}$$

Applied to the individual loss components, one may speak of the "coupling Q" or equivalent "surface scattering Q".

Let's start to pick apart this formula and examine where the predominant sources of loss in our diamond system.

Surface scattering is the most visible target to improve - you can see a rough surface on an SEM - therefore most effort goes into trying to reduce this roughness. Although new techniques have been developed which give an actual quantification of the roughness of an etched sidewall, by and large visual inspection and experimental testing of quality factors with a fabrication parameter sweep remain the main methods for developing processes which reduce surface scattering loss.

To gain an estimate for the effect of surface scattering loss given a sidewall roughness, we can consider it akin to Rayleigh scattering, where the loss is proportional to the scattering cross-section, or[12]

$$\alpha_{\text{s.s.}} \sim \left(\varepsilon_{\text{wg}} - \varepsilon_{\text{cladding}}\right) \frac{\sigma_{\text{rms}}^2 d_r}{\lambda^4} \tag{2.3}$$

where ε_{wg} and $\varepsilon_{\text{cladding}}$ are the dielectric constants for the waveguide and cladding respectively, σ_{rms} is the root-mean-square value of the surface roughness, and d_r is the correlation length of the

roughness. Looking at this heuristic, we immediately see that higher index contrast means more loss (a problem for diamond) and that there is a strong dependence on loss with respect to wavelength (a problem for going to shorter wavelengths). More formally, we can write the scattered power per length in the following way[12]

$$\frac{P}{4L} = \oint \left(\vec{S}_{\text{element}} \cdot \hat{r}\right) \tilde{R}\left(\beta - nk_0 \cos\theta\right) \mathrm{d}A \tag{2.4}$$

where the integral is essentially a surface integral describing the Poynting vector generated by scattering $\vec{S}_{\text{element}}$ multiplied by the spectral density of the roughness $\tilde{R}$, itself dependent on the propagation vector of the mode with respect to the scattered angle θ. By approximating a scattering element (the roughness) as an infinitesimal line at the surface, we simply need to account for the Poynting vector at the surface to construct a reasonable model of scattering loss.

The next loss term, $\alpha_{\text{s.a.}}$, stems from absorption due to impurities and adsorbates on or near the surface. This can be caused by ion implantation due to etching or simply gas-phase molecules present in the laboratory sticking on the surface. Indeed, it is known from the diamond literature that chlorine-based etch chemistries can leave residual chlorine in the diamond lattice, up to a distance of a few nm, based on x-ray photoemission spectroscopy (XPS) and studies with nitrogen vacancy centers, which are particularly sensitive to such dipolar molecules[40]. For this reason, chlorine-based etches are never used to finish the etching of diamond optical structures, though it is used in other circumstances as we will see. Within the context of diamond surface chemistry, the literature provides us with strong suggestions for proper treatment. In some systems such as silicon, a piranha acid-clean greatly helps in cleaning up the exposed surface, creating a thin oxide layer, and reducing losses[17]. For diamond, because the exposed surface is simply carbon atoms with possible etch damage, a way to ensure a clean surface is paramount. To this end, there are two main ways to ensure such a surface. The first is to perform a tri-acid clean, which removes any contaminants, slightly

etches the diamond surface, and leaves a clean surface termination. Because this process is quite violent, for especially sensitive structures such as those appearing in this thesis it is not recommended. The other process is an oxygen-based anneal[130]. Although adopted for NV-center studies[40], it has the added benefit of etching away any residual carbon on the diamond surface (which may absorb) and terminating the surface with oxygen molecules, ensuring a pristine surface with no midgap states which can absorb optically[166]. A fluorine-based chemistry can also guarantee such a surface state, whereas hydrogen-terminated diamond surfaces actually are conductive and potentially absorb light[181,166]. Although in principal this loss mechanism is easily modeled, it is typically only investigated experimentally. All the work done in this thesis relied on a final oxygen anneal step to treat the surface and ensure good theoretical surface quality, though more experimentation could surely be done in this regard.

The next loss mechanism, $\alpha_{\text{rad.}}$, is loss due to bending and coupling to free-space radiative fields. For the size of resonators which we study, bending loss is negligible, particularly because our diamond waveguides are completely etched. The bending loss can be thought of intuitively as a coordinate transformation which changes the cylindrically symmetric problem of a bending microresonator into one where the microresonator is straightened out[37]. This coordinate transformation has the side-effect of changing the refractive index, so that it is now an increasing function of the radial coordinate. If we think of the refractive index as a potential well, where the modes of the optical waveguide sit, then we see that because the transformed refractive index increases outside of the waveguide, it will allow the mode to effectively tunnel into free-space. A detailed analysis[37] reveals that the bending loss increases exponentially as the inverse of the bending radius of curvature.

For certain device architectures (such as tapered-support angle-etched devices or even in the diamond-on-insulator platform), where the device cross section can slowly vary over the waveguide, we can perform a more general treatment and derive the lost power over the geometry change via[174]

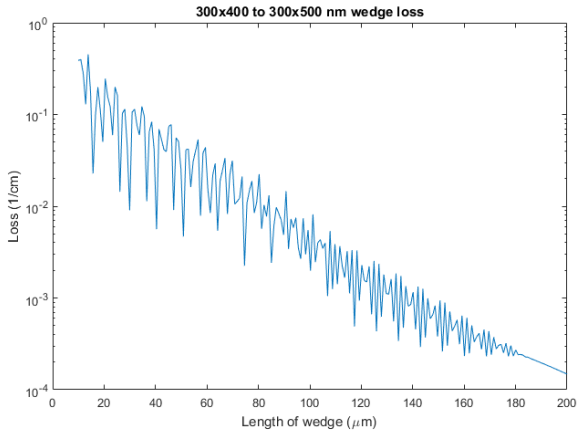

Figure 2.1: Simulated radiative loss of diamond due to wedging or tapering. The diamond waveguide cross-section varies from 300x400 nm to 300x500 nm over the length of the wedge. The loss is given in units of inverse centimeters. The equivalent Q factor of the implied loss remains above 10^6, even when the tapering occurs for a distance as short as 20 μm. The wavelength used in the simulation was 700 nm.

$$P_{\mathrm{L}} = \int_0^{k_0 n_z} |q_T(\kappa_s, L)|^2 \, \mathrm{d}\kappa_s \tag{2.5}$$

where

$$q_T(\kappa_s, L) = \int_0^L \eta_{gr}(\kappa_s, z) \exp\left\{ i \int_0^z \left[\beta(\kappa_s) - \beta_0(z')\right] \mathrm{d}z' \right\} \mathrm{d}z \tag{2.6}$$

and η_{gr} is the overlap integral of the guided mode and the radiative mode given by κ_s. Fig. 2.1 shows simulations of the loss due to a wedge in the diamond film. For typically wedge values (100 nm of height increase across 100 μm of diamond film) the loss is negligible compared to measured results.

The final loss mechanism before we discuss coupling is $\alpha_{\mathrm{b.a.}}$, the loss due to bulk absorption. As mentioned previously, diamond plays host to a number of different color centers which can absorb light at specific wavelengths. Depending on the diamond growth process, the impurity concentration can be a dominant source of loss, as can be easily understood simply by looking at a yellowish piece of HPHT grown diamond. For research use, we exclusively use electronic-grade diamond with

ultra-low impurity concentration. At visible wavelengths (532 nm, see Fig. 2.2) we see that at nitrogen concentrations typical of standard type IIa diamond ($\sim$ 100 ppb) bulk loss is non-negligible, at roughly 10^{-1} cm^{-1}. Electronic grade diamond, with concentrations of $\sim$ 20 ppb, shows much lower loss of 10^{-3} cm^{-1} [42].

2.1.2 Coupling

The question of coupling is related to, yet distinct, from the other loss considerations. Although phenomenologically equivalent from the point of view of the resonator mode, the main distinction between coupling and other sources of loss is the fact that loss due to coupling ends up in a mode which is "useful". This may mean that the resonator mode is injected through this channel, or it is but one output port among many which control it. An accessible introduction and analysis of the coupling of a microresonator to a waveguide is provided in Yariv [190].

From a practical standpoint, system design is usually undertaken with a finite-element method solver (COMSOL). This fast simulation technique is able to easily grab coupling parameters for a given geometry, and from there the system parameters can be chosen to develop an optimized coupling scheme. An advantage of this technique over perturbative approaches is the accurate coupling parameter retrieval, particularly when the waveguides are close. The two waveguides are placed a distance apart d, and the difference between the symmetric and antisymmetric mode indices reveals the coupling parameter κ [77]. Other techniques such as FDTD can be used for full 3D or 2.5D simulations.

2.1.3 Dispersion

Especially for nonlinear optical devices, dispersion is a vital system parameter. At the most basic level, we can define a the propagation constant for an optical mode traveling along a waveguide:

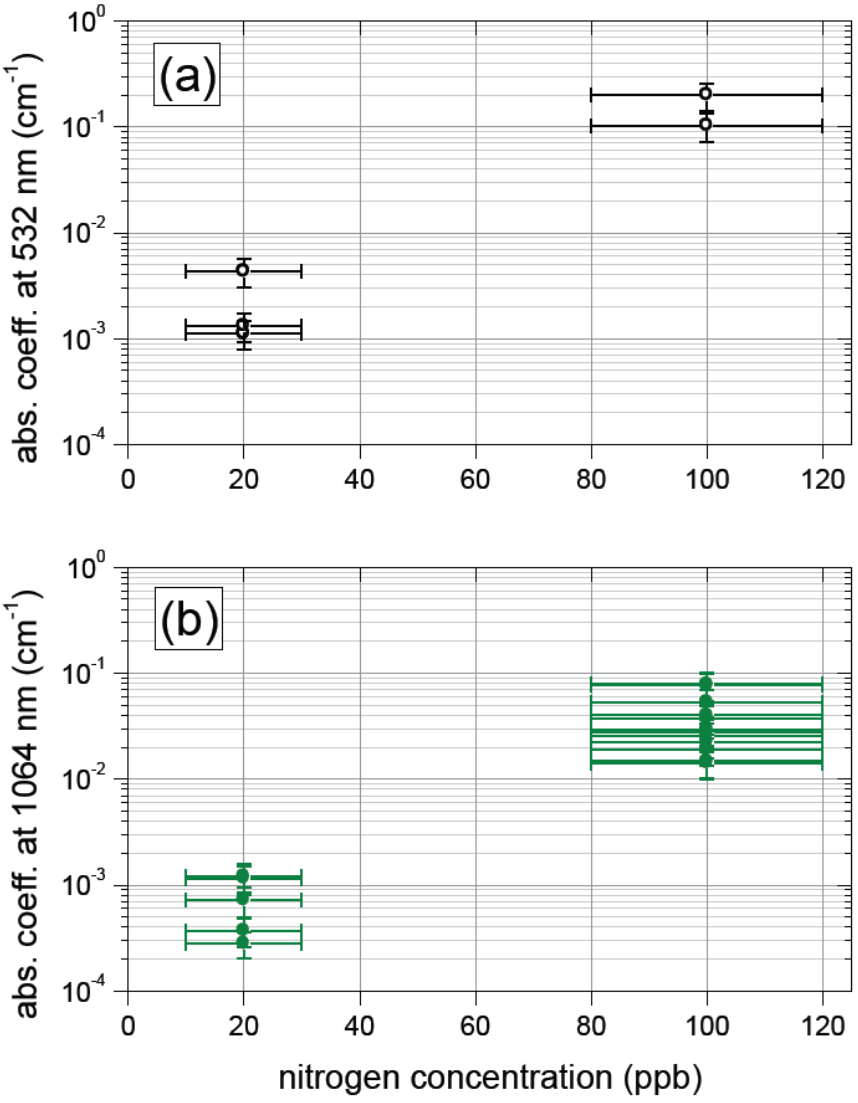

Figure 2.2: Measured losses for single crystal diamond material of varying impurity concentration at (a) 532 nm and (b) 1064 nm probe wavelength. 100 ppb concentration is typical of standard CVD (type IIa) diamond, whereas 20 ppb is typical of electronic-grade diamond[42]

$$\beta(\omega) = \omega n_{\text{eff}}(\omega)/c \tag{2.7}$$

Where $n_{\text{eff}}(\omega)$ is the frequency-dependent effective index of a mode (either in a waveguide or free-space), dependent on the material parameters and precise cross-sectional geometry. The group velocity v_g, or the speed at which a pulse moves along the waveguide, is simply given by the inverse of the derivative

$$v_g = \left(\frac{\partial \beta}{\partial \omega}\right)^{-1} \tag{2.8}$$

The group velocity dispersion (GVD) is likewise the second derivative of the propagation constant:

$$\text{GVD} = \left(\frac{\partial^2 \beta}{\partial k^2}\right) \tag{2.9}$$

This term describes how a pulse spreads out (or narrows) as it propagates. Typically, the GVD is normal, meaning that higher frequency light propagates slower than lower frequencies. The inverse relationship is termed anomalous dispersion, and is important for the zoo of nonlinear effects we will study and demonstrate in diamond later.

By taking consecutive derivatives of the propagation constant, we can reach higher order terms of the dispersion of the waveguide. Although higher order terms are usually much smaller in magnitude, their effects can still have a visible result on experiment.

Lastly, the final relevant parameter in microresonator design is the free spectral range (FSR) of the resonator. This parameter is just the spacing in frequency-space between subsequent resonances of the same mode family in a resonator. It is related to the group velocity v_g and the total resonator round-trip path length L by $\text{FSR} = v_g/L$.

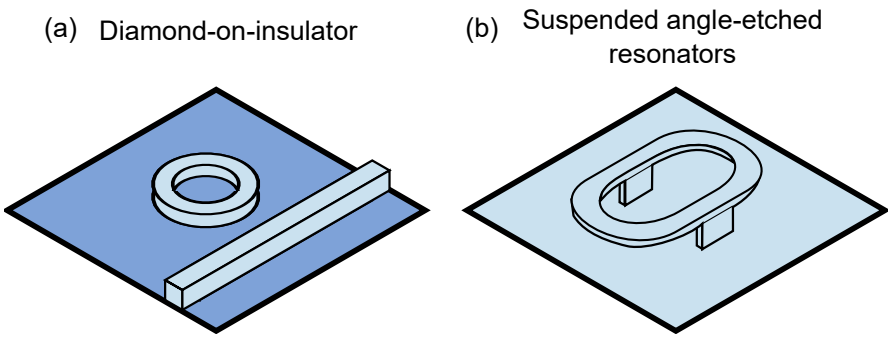

Figure 2.3: Diamond device geometry comparison. (a) Thin film, diamond-on-insulator example of resonator on dielectric with nearby bus waveguide. (b) Angle-etched resonator with supporting sections. In some designs, the support can run throughout the waveguide, instead of tapering in and out.

2.2 Diamond device geometries

Diamond's difficult synthesis yields tough constraints on the available device geometries and architectures. Fig. 2.3 shows the two architectures studied in this thesis for photonic structures. The first (Fig. 2.3(a)) is the diamond-on-insulator, or DOI approach. As the name suggests, it is analogous to silicon-on-insulator devices, which are prevalent in industry. It comprises of a thin diamond film resting on top of a lower-index dielectric, typically glass.

The other architecture (Fig. 2.3(b)) shows the angle-etched photonic geometry, where the waveguide is suspended above the substrate. This architecture has some advantages and disadvantages compared to DOI, which we will discuss in detail here.

2.2.1 Diamond-on-insulator architecture

The diamond-on-insulator (DOI) architecture consists of a thinned-down diamond plate on a dielectric. This design was first developed for quantum optics[63], and has historically been the surest choice for a high-Q resonator[62]. Due to the fragility of the diamond film, this process has low yield and takes several hours simply to prepare a substrate for further processing.

Typically, the diamond used for DOI devices is a thin plate ($\sim$ 30μm-thick) fashioned by polishing. The diamond is further thinned in an etching tool, then transferred to the final dielectric substrate. A number of suitable substrates were investigated and judged on the metrics of refractive index, chemical compatibility, and ability to fabricate on. Sapphire, although an inert substrate which bonds well to the diamond, was ultimately unsuitable due to its relatively high refractive index, which makes constructing auxiliary waveguides such as mode-converters impossible. Calcium fluoride (and other fluorites) were investigated as well, but their brittle nature caused difficulties in downstream processing, with the diamond film easily breaking just be putting the prepared diamond/substrate on a hot plate. Ultimately, a fused silica substrate was chosen. Although fused silica is not particularly compatible with all of the chemistry one would like to do on the device, it met the requisite low refractive index and broad transparency. In particular, fused silica was chosen over thermal oxide grown on silicon because the latter suffered from absorption of visible light, to the extent that the substrate would melt, given powerful-enough light.

2.2.2 Angle-etched diamond architecture

The angle-etched diamond architecture was investigated over the course of this thesis, though it saw limited use in nonlinear optics until recently. Now commonly used in quantum optical devices[28,64,158,16], this architecture has a number of benefits compared to DOI, though it still does not possess the high quality factors and coupling can be difficult.

In an angle-etched device, the waveguide is fashioned on top of a bulk, single-crystalline diamond plate and typically triangular in cross-section , suspended above the diamond substrate. Supports can either be periodically defined along the device, or the device may be supported throughout, resting on a narrow rib. Both microresonators and photonic crystals are readily made in this platform, and they can be interface either through tapered fiber coupling[25,99,27] or by mating tapered diamond with a single-side tapered optical fiber[158,29]

Although both architectures can be used to create devices with great optical properties, the main reason the angle-etched architecture is prevalent in quantum optics work is due to the ability to anneal the structure. This anneal is vital to ensure that the quantum defects are properly situated and prepared in the diamond lattice, and because the structure is monolithically defined, it is naturally robust to the extreme temperatures required of annealing. The DOI platform, on the other hand, should be limited to temperatures below 450 C, for two reasons. For one, the thermal expansion mismatch between diamond and the substrate can cause the film to become stressed and break at these extreme temperatures. For two, the high temperatures typically cause a small degree of graphitization, even in an ultra-high vacuum furnace, which must be cleaned off[130]. Although an oxygen anneal may in principle be used to perform such a cleaning procedure, in practice a tri-acid clean is more widely used. This tri-acid process is too damaging to the DOI film and can easily cause device failure.

2.3 Mode-converter end-facet coupling

Especially in the DOI platform, interfacing the small diamond portion of the device with the outside world can be a difficult task. For these devices, a high coupling efficiency is required, particularly when studying nonlinear optics. To bring the light to the devices, we use a mode-converter platform. On the facet-side, the mode converter waveguide has a relatively large cross-section but low refractive index. Its dimensions are optimized numerically to maximize coupling into the mode converter waveguide from a free-space beam of light, with a theoretical efficiency of 70%. The light is then funneled to the diamond device, where the mode converter waveguide tapers down and the diamond waveguide, which it overlaps, tapers up. This ensures an efficient mode transfer from one to the other.

Mode converter waveguides are usually built on top of the diamond waveguide structures. Be-

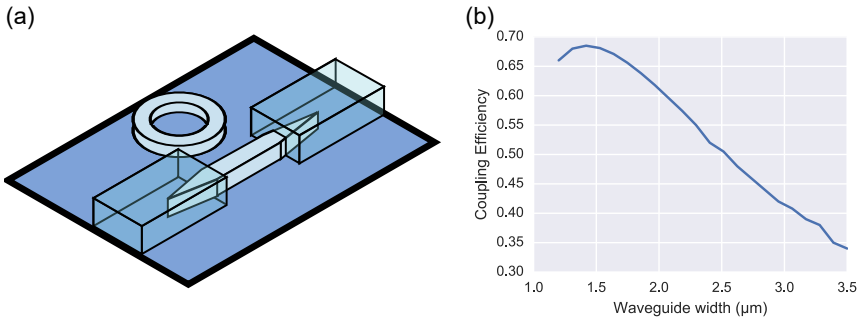

Figure 2.4: Mode converter for DOI platform. (a) Schematic of mode converter collecting light and transferring it to diamond waveguide (b) Plot of mode converter coupling efficiency to input light beam. In this case, $\lambda = 800\text{nm}$ and the beam waist radius is $2\mu\text{m}$, set to TE polarization. Efficiency is maximized for a cross-section of $1.4\mu\text{m}$ by $1.5\mu\text{m}$ with a refractive index of $n = 1.5$ on top of an oxide substrate. The peak efficiency is 70%.

cause of this, a facile fabrication method is important to ensure minimal disturbance to the underlying diamond. Early iterations of this technique[67] used electron-sensitive polymer (SU-8), but this proved to be insufficient for visible wavelengths and the high powers required in nonlinear optics. A new material platform which contains the same advantages of the polymer design was developed and is discussed in a later chapter.

2.4 Loaded tapered-fiber coupling

Efficient coupling of light to and from an integrated chip is essential for many applications in nonlinear optics[50,33,115,62,184,39,97], optomechanics[126,89], and quantum optics[52,67]. For these reasons, a wide range of techniques with which to deliver light on-chip have been developed, such as end-fire coupling[22], planar gratings[169], or optical fiber coupling to dispersion-engineered photonic crystal waveguides[10,11]. However, many of these applications are based on materials where full integration with couplers is difficult or rapid testing is desired. In this case, silica fiber tapers, where an optical fiber is pulled to have a thin coupling region with thickness on the order of the wavelength, have formed the foundation for many optical experiments[95,32] in diverse environments, including cryogenic temperatures[168,143]. In particular, when the fiber and resonator are made of the same material,

this platform has shown ideal coupling to silica resonators[164], with the ability to tune from under- to over-coupled based on taper positioning.

When the fiber and resonator are made from different materials, for instance when probing high-index micro-resonators, the refractive index contrast results in a large propagation constant mismatch between the resonator and the tapered silica fiber, limiting the coupling efficiency into the resonator[167]. Although the interaction length may be small, on the order of a few wavelengths, the coupling into such systems is still limited. This problem is exacerbated in larger, multimodal resonators where the fundamental mode is confined mostly in the material. Existing schemes, such as prism coupling[79], or separate on-chip waveguides brought near to the resonator of interest[23], are bulky and do not allow for rapid testing of multiple components. We demonstrate a method to improve upon the phase matching by introducing an ancillary, support waveguide attached directly to the optical fiber. This technique enables stronger coupling to a variety of high-index resonators, both small and large mode volume devices. Because it extends the familiar tapered-fiber approach, it is expected to be applicable to numerous materials and environments. Using this platform, we show efficient coupling to the fundamental modes of a multimodal silicon micro-resonator. In particular, we fabricate a tapered, free-standing, angle-etched silicon waveguide which we then detach and affix to the tapered section of an optical fiber. This achieves adiabatic mode conversion between the optical fiber and the silicon waveguide, effectively changing the propagation constant of the input light to that of silicon, matching it to the resonator.

2.4.1 Coupled-mode theory

Energy transfer between waveguides can be modeled via coupled mode theory[189]. For two waveguides in close proximity with propagation constants β_1 and β_2, field amplitudes A_1 and A_2, and mutual coupling κ, the coupled mode equations may be written as

$$\begin{cases} \frac{dA_1}{dz} = & i\kappa A_2 e^{i(\beta_2-\beta_1)z} \\ \frac{dA_2}{dz} = & i\kappa^* A_1 e^{-i(\beta_2-\beta_1)z} \end{cases} \tag{2.10}$$

where z is the length dimension along the direction of propagation. These equations can be solved analytically. In particular, if light is only injected into one waveguide, then the power in the second waveguide as a function of distance is given by

$$P(z) = \frac{1}{1+(\delta/\kappa)^2} \sin^2(\sqrt{\kappa^2+\delta^2}z) \tag{2.11}$$

where we define $\delta = (\beta_2 - \beta_1)/2$. Under this formulation, the maximum power coupled into the second waveguide cannot exceed $P_{\max} = 1/(1+(\delta/\kappa)^2)$. The corresponding length such that $P(L_c) = P_{\max}$ is given as

$$L_c = \frac{\pi}{2\sqrt{\kappa^2+\delta^2}} \tag{2.12}$$

which indirectly depends on wavelength through δ and κ. In particular, for a fixed coupling constant κ, the power coupled in ($P_{\max}$) decreases rapidly as a function of increasing effective index mismatch (δ). This underlies the difficulty of coupling to high-index micro-resonators from optical fibers.

2.4.2 Device geometry

In contrast to uniform coupling waveguides, adiabatic coupling requires the breaking of translational symmetry in order to transfer optical energy across waveguides. This technique has been used to couple light out of optical fibers and into target waveguides made of silicon[59], silicon nitride[172,173] and diamond[134]. Here, we adiabatically transfer the mode from an optical fiber into such a waveguide (referred hereafter as the "loader waveguide"), using the converted mode to efficiently couple to

an index-matched micro-resonator.

The loader waveguide and resonator were fabricated via angled-etching within a Faraday cage[28,148]. After a resist mask is defined, the pattern is first etched vertically. In a second step, a Faraday cage is placed around the sample, directing the incident etching ions to the substrate at an angle defined primarily by the cage geometry[28]. In order to leave the structure suspended, the etch is timed to end before the structure is completely undercut[25]. The final etch profile shows a three-dimensional taper since the etch depth is defined via the width from the etch angle. This can yield adiabatic coupling over shorter intervals compared to thin film platforms[134,23]. Because the operating principle of angled-etching is agnostic to the etch chemistry, it has been used successfully in a number of materials, including diamond[28,30,25,15], quartz[98], and silicon[148,98].

Once sample etching is complete, the resist is removed and the loader waveguide is manually detached from the substrate using an eyelash probe. This situates the waveguide on the surface so that it is easier to pick up. Thereafter, it is transferred to a tapered optical fiber by bringing the fiber into contact as it lies on the substrate. The bi-conical, adiabatically tapered silica optical fiber is formed by pulling a single-mode fiber as it is heated[95], ending with a taper diameter of $\sim$ 1μm over a length of several hundreds of microns. As the waveguide is completely detached from the substrate, the fiber's attractive forces are sufficient to pull the device off. Afterwards, the loader waveguide is manipulated against features on the sample until it lies parallel to the fiber. During the loading process no visible damage is done to the fiber, although the fiber transmission reduces to $\sim$ 0.92%. In a separate experiment, we investigated the impact the loading process has on the overall fiber transmission. When the flat fiber is brought close to contact with the waveguide, it tends to make contact with the substrate as well. Each time this happens, a transmission drop of 10-30% is observed from dirt on the substrate sticking to the fiber. Upon transferring the waveguide from the substrate to the fiber, a further drop of $\sim$80% is observed. Overall, the loading process reduced the transmission of the fiber to 8.8% of that before loading. Alternately, cleaner methods of loading[134,173] can increase

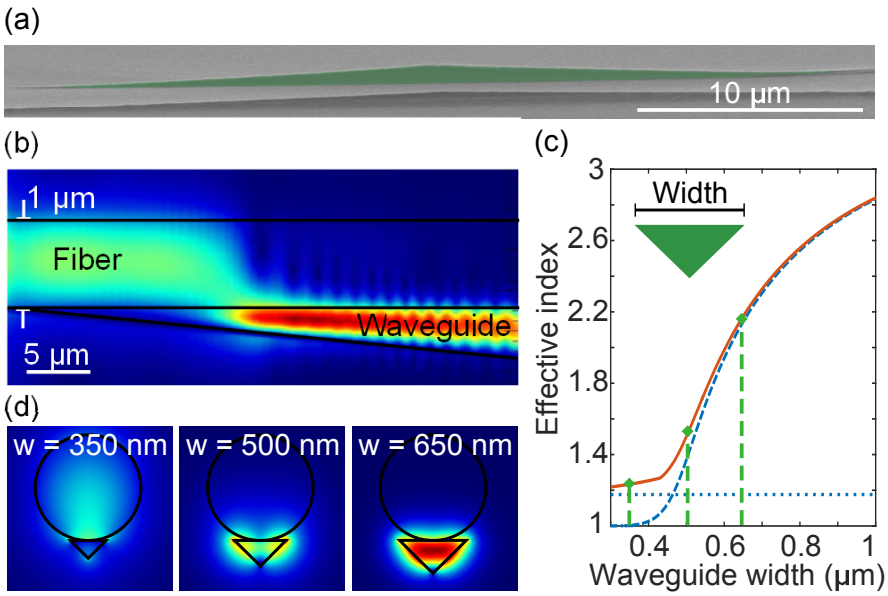

Figure 2.5: Silicon adiabatic coupler loaded to tapered fiber. (a) Scanning electron micrograph (SEM) of a silicon loader waveguide fabricated via Faraday cage angled-etching. The device (green) tapers down from a nominal width of $1\mu\mathrm{m}$ in the center to points at the ends over a length of $20\mu\mathrm{m}$. The supporting fin is etched through, collapsing the device onto the substrate (b) Finite-difference time-domain (FDTD) simulation of a silicon adiabatic coupler loaded to a tapered silica fiber with diameter $1\mu\mathrm{m}$. The mode originally stays in the silica fiber before being drawn into the waveguide. The waveguide is shown coupling over $40\mu\mathrm{m}$. The multimodal nature of the thicker end of the silicon waveguide visibly manifests itself as interference fringes in the normalized electric field. (c) COMSOL simulations of the effective index of the tapered fiber (blue, dotted), silicon coupler device (blue, dashed), and supermode (red, solid) as the width of the silicon coupler is increased from 300nm to $1\mu\mathrm{m}$. The dashed green lines correspond to the cross-sections shown in (d). The waveguide's cross-section is an isosceles triangle with equal angles of 30°. (d) Supermodes obtained at different cross-sectional waveguide widths.

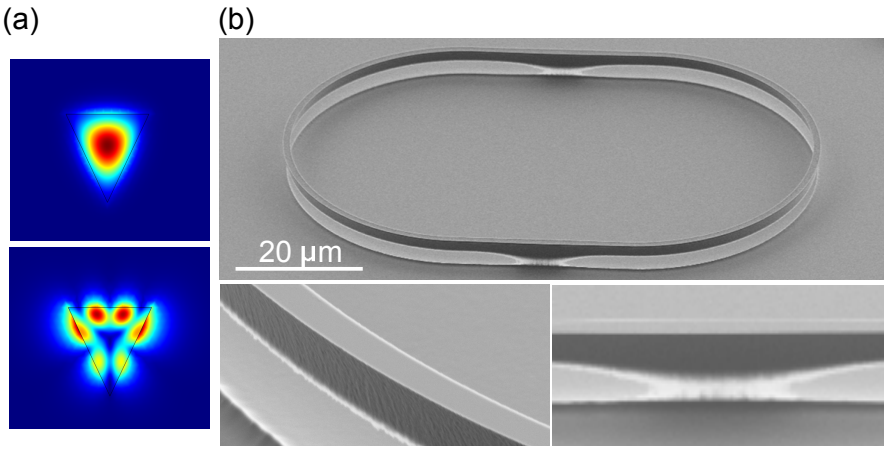

Figure 2.6: Studied silicon micro-resonator with nominal width 1μm. (a) Simulated modal profiles at $\lambda_0 = 1.52\mu m$ for the fundamental (top, $n_{\text{eff}} = 3.19,\ n_g = 3.85$) and higher-order (bottom, $n_{\text{eff}} = 1.52$) modes using an extracted etch angle of 65°. (b) SEM image of silicon micro-resonator (top) fabricated with angled-etching. The bending radius is 25μm while the straight region is 25μm. A hydrogen silsesquioxane (HSQ) mask is defined via electron-beam lithography and the structure is etched in a SF_6/C_4F_8 chemistry. Slight waveguide roughness (bottom left) is visible on the SEM. The structure is supported above the substrate by silicon fins (bottom right). This is done by timing the etch so that material remains under the widened support sections.

the overall transmission of this design.

2.4.3 Broadband coupling to high-index resonators

In this study, we investigate the optical properties of a suspended silicon resonator etched via angled-etching (Fig. 2.6(b)). As reported elsewhere [25], by selectively widening the patterned area and precisely timing the etch, we can create a fully suspended resonator with supporting sections ensuring sufficient distance ($\sim$ 2μm) from the substrate. The nominal width of the resonator is 1μm, while the support region is 1.1μm wide. The bending radius of the resonator is 25μm and the etching angle is seen from SEM to be $\sim$ 65°. The cross-sectional area is sufficiently large to support several transverse modes in the structure. Fig. 2.6(a) shows the mode profile of two such modes simulated in COMSOL. The top mode is the fundamental mode for the structure, with a calculated effective index of 3.19. The higher-order mode (bottom) shows a calculated effective index of 1.52, signif-

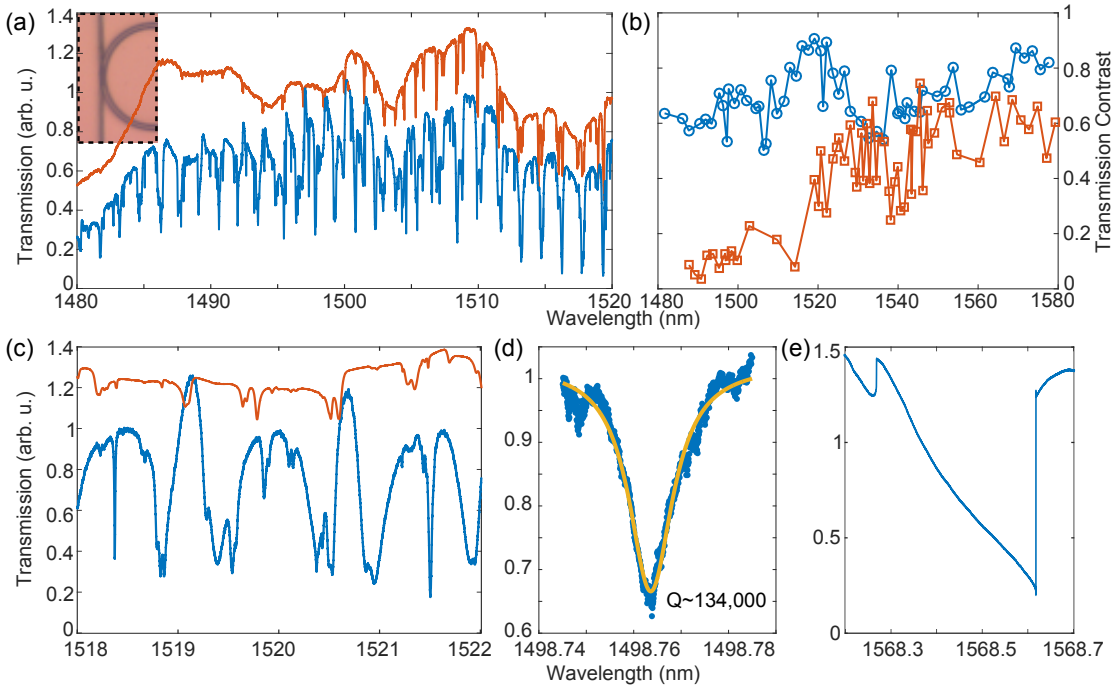

Figure 2.7: Silicon micro-resonator spectra with bare fiber (red) and loaded fiber (blue) (a) Transmission measurements of silicon micro-resonator (inset, optical image) from 1480nm to 1520nm. The bare fiber result is shifted from the loaded fiber result for clarity. In the bare fiber case, the transmission dips are not consistent across the entire spectrum, petering off at shorter wavelengths. In contrast, the loaded fiber shows a consistent coupling for all modes in the same family, with transmission dips of $\sim 40 - 60\%$. (b) Extracted transmission contrast for highly-coupled resonances as a function of wavelength under bare (red) and waveguide-loaded (blue) coupling. (c) Optimized coupling to high-quality factor (Q) resonances at $\sim 1518.4\text{nm}$ and $\sim 1521.5\text{nm}$. The measured Qs were $\sim 130,000$ and $\sim 40,000$, with transmission dips of 60% and 80%, respectively. These modes are not visible when coupling with the tapered fiber alone. (d) Lorentzian fit to a high-Q mode at short wavelengths, showing a Q of $\sim 134,000$. (e) "Shark-fin" shaped transmission dip characteristic of optical nonlinearities when pumping at $\sim 138\mu\text{W}$ of measured power.

icantly lower than that of the fundamental. Seen through the lens of coupled mode theory, this implies that the higher-order mode should couple more easily to the bare optical fiber (effective index $\sim$ 1.4). Furthermore, as seen in the mode profiles, the higher-order mode has more electric field concentrated on the sides of the waveguide. This implies that it interacts more with any surface roughness and adsorbed molecules, limiting the Q-factor of this mode family. In contrast, the fundamental mode is relatively isolated from the surface, resulting in a higher predicted Q-factor. In addition to these considerations, the multimode nature of the silicon resonator may increase losses during the transition region by the supports, causing coupling between mode families and limiting the observed Q-factors. For these reasons we do not necessarily expect our Q-factors to be limited by material absorption. Crucially, however, we anticipate that the fundamental mode has a much larger Q-factor than any other mode.

After loading the fiber with the silicon waveguide, the fiber was brought near the micro-resonator. The transmission spectrum was monitored continuously over a narrow bandwidth as the fiber was moved closer. The position of the fiber was controlled in 50nm increments via stepper motor. Because the effective index of the loaded waveguide changes along its direction of propagation, all three spatial dimensions were used to tune the coupling to the resonator. Within the context of coupled mode theory, the coupling constant (κ) was tuned by changing the height of the fiber off the substrate and its lateral distance to the resonator, while the effective index contrast (δ) was modified by translating the fiber along its length.

The data shown in Fig. 2.7 were taken at optimized coupling locations either at a bare fiber section (red) or at the section containing the loaded waveguide (blue). A tunable telecom laser (Santec TSL-510) scanned the resonator as the transmission collected by a photoreceiver was monitored. Under a large scan range, the difference between the loaded and unloaded section is apparent when looking at shorter wavelengths (Fig 2.7(a)). The increased coupling bandwidth is a result of the better index-matching.

In the studied geometry, the length over which there is significant mode overlap between the resonator and the fiber is fixed due to the curvature of the resonator. Furthermore, the effective indices of the silicon device and resonator both vary strongly as a function of wavelength, whereas the index of the bare fiber varies weakly so. Because the loader waveguide and resonator have similar cross-sections, we can expect that their indices have a similar dependence on wavelength. In the context of coupled-mode theory, this implies that δ for the waveguide-loaded fiber-resonator system depends weakly on wavelength whereas δ for the bare fiber-resonator system depends on it strongly. This, in turn, imparts a strong wavelength-dependence on the coupling into high-index resonators from bare fibers (Fig. 2.7(b)).

In addition to increased coupling bandwidth, transmission measurements of the resonator with the waveguide-loaded fiber showed the ability to access the high-effective index, fundamental modes of the device. Fig 2.7(c) shows a scan taken at an optimized coupling position for wavelengths around 1520nm. High-Q resonances can be seen at $\sim$ 1518.4nm and $\sim$ 1521.5nm. The group index can be calculated from the free spectral range (FSR, $\Delta\nu$) as $\Delta\nu = c/(n_g L)$ where L is the resonator path length. These two modes give $n_g \approx 3.6$, which is close to the calculated value of $n_g = 3.85$. The discrepancy can be explained by differences between the simulated and actual waveguide dimensions. The modes are high-Q ($\sim 130,000$ and $\sim 40,000$, respectively), suggesting that they are relatively well-isolated from any surface scattering or absorption compared to higher-order modes (Q $\sim 10^4$). Furthermore, both show large transmission contrast ($\sim 60\%$, $\sim 80\%$, respectively), demonstrating that the waveguide-loaded fiber can efficiently transfer energy to the high-index modes of the resonator. The lower Q-factor of the second mode may be due to coupling with other resonances in the system[139]. Although the ideality of coupling is difficult to demonstrate in this system, we do not observe a significant dependence on intrinsic quality factor as the coupling gap is decreased, suggesting small parasitic losses[164].

To further demonstrate the ability to transfer large amounts of power to the resonator, we op-

timized the coupling at a particular wavelength and then increased the power of our laser until we observed optical bistability[185]. As the laser is tuned across the resonance, the resonance peak is red-shifted until some critical detuning where the resonance transitions to a regime where it is no longer coupled (Fig. 2.7(e)), resulting in a characteristic "shark-fin" shape. Additionally, no degradation in coupling was observed for higher powers, up to the maximum laser power available.

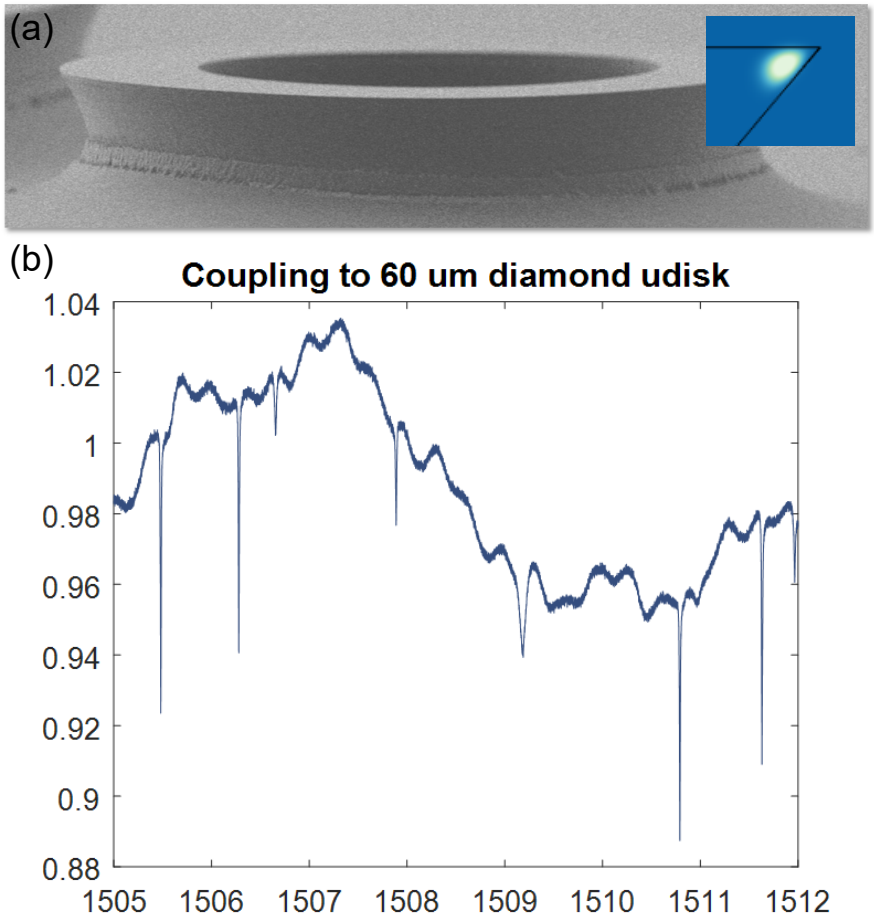

Figure 2.8: Diamond angle-etched microresonator. (a) The diamond microresonator has a diameter of 60 μm. The mode is confined to the corner, as shown in the inset. This design minimizes interaction with the surface, but needs to be rather large to avoid leakage into the substrate. (b) Spectrum measured with loaded coupler. Peak quality factor is measured to be 170,000. This spectrum is not visible with traditional means.

Fig. 2.8 shows a diamond microresonator fabricated with Faraday cage angle-etching. A thick alumina mask ($\sim$ 200nm thick) is used during etching to enable this large-scale structure. The mask has a donut topology to reduce the surface area of mask material, thereby eliminating micromasking due to re-sputtered mask. The mode sits in the corner of the microresonator, relatively deep within the diamond. This type of resonator design minimizes interaction with the surface and can potentially enable very high quality factors. However, the tradeoff with this approach is that the optical mode sits deeper inside the resonator, increasing the mode's effective index and minimizing coupling overlap. Using a loaded fiber coupler, the refractive indices are brought closer together and the modal spectrum of the resonator can be probed. In this case, peak quality factors of 170,000 were recorded for a device of 60μm diameter at telecom wavelengths. Improved fabrication techniques (for example, ion beam etching) should see this quality factor only increase.

In summary, we have demonstrated a versatile technique which enables coupling from an optical fiber to a free-standing high-index micro-resonator with peak efficiency approaching 80%. This result enables the rapid and large-scale optical probing of material systems useful in nonlinear and quantum optics such as diamond[25], chalcogenides[23], lithium niobate[177], or III-Vs[126]. Shorter operating wavelengths for quantum photonics with single-photon emitters like NV centers in diamond micro-cavities can be obtained by fabricating the loader waveguide with high-index materials transparent at visible wavelengths[134,39]. Arrays of devices can be accessed by modifying this approach to work with dimpled optical fibers[121], where the loader waveguides are fabricated with a radius of curvature matched to the dimple. Furthermore, the device can be tailored to the specific resonator by modifying the geometry of the coupling region, leading to highly efficient coupling directly to optical fibers.

3

Fabrication techniques

Single-crystal diamond is still an exotic newcomer to the pantheon of semiconductor materials. Unlike other materials, such as silicon, its growth process is notoriously slow and expensive, relying on a diamond seed crystal and a chemical vapor deposition process to grow the crystal atom-by-atom. Crucially for photonics, diamond lacks a reliable thin-film technology, where a high-quality layer of dielectric can be grown or transferred on top of another material. Such a thin-film technology is used in the popular silicon-on-insulator (SOI) and silicon-on-sapphire (SOS) platforms, where the

silicon device layer rests on top of a dissimilar dielectric (silicon dioxide and sapphire, respectively). The dielectric spacer is of a lower refractive index than the silicon, providing optical mode confinement for etched photonic structures. An analogous commercial platform does not exist in diamond, although some research groups have made progress applying an ion implantation and liftoff process similar to the one used to fabricate SOI wafers[141].

Beyond the basic requirement of a scalable material platform, diamond fabrication itself needs to be extensively developed in order to meet the device specifications. For optical devices in particular, etch sidewalls must be very smooth to minimize scattering losses, and critical dimensions must be precisely replicated in order to have reliable translation of design to reality. In other scenarios, the challenge becomes finding an etch mask which can withstand the unforgiving diamond etch chemistry and yield large diamond structures with minimal mask deposition. In this chapter, we will discuss various techniques which we have developed to meet the challenges of diamond processing, from etch optimization to application-specific solutions.

3.1 Faraday cage angle-etching

High-quality heteroepitaxially grown substrates, from III-V semiconductors such as GaAs/AlGaAs to the readily available silicon-on-insulator (SOI) platform, form but a short section of the list of materials with attractive electro-optical, mechanical, piezoelectric, or other properties useful in nano-devices. These thin film structures, on the other hand, enable the necessary three-dimensional architectural control required for the isolation of the active layer from the bulk substrate. For nanomechanics, access to the mechanical degrees of freedom often requires the physical separation of structures from the host, such as in suspended beams, cantilevers, and membranes[51]. Nanophotonics, likewise, requires light confinement due to either total internal reflection or distributed Bragg reflection[84]. Although a heterolayer structure composed of materials with distinct photonic or phononic

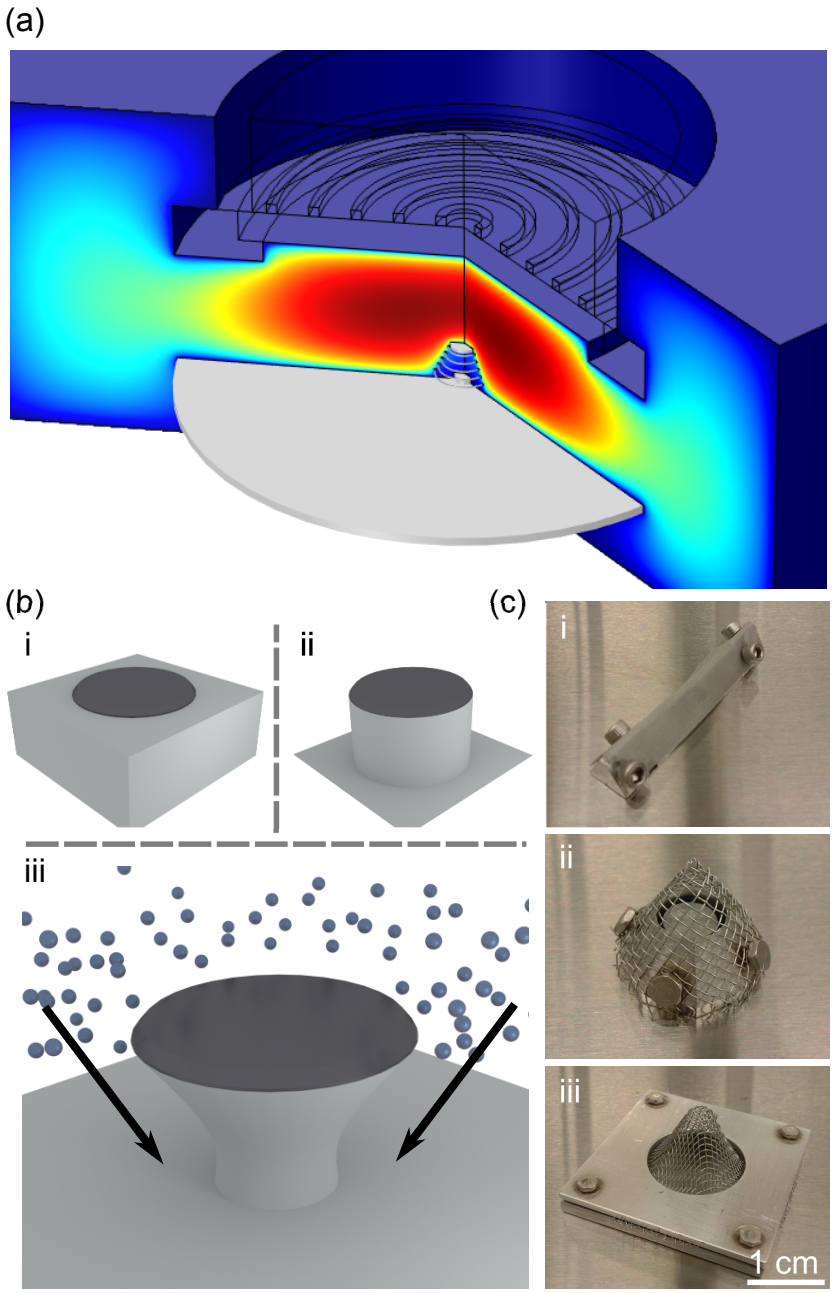

Figure 3.1: Overview of Faraday cage angled-etching (a) 3D simulation of reactor chamber with an argon plasma and a central Faraday cage (details in text). (b) Schematic of angled-etching. The ions, directed by the equipotential on the cage boundaries, are incident upon the sample at an angle. The etch mask defines the shape of the structure. (c) Examples of different cages used for angled-etching, including (i) a triangular cage with a fine mesh (ii) a wrapped cage with a coarse mesh and (iii) a molded cage with a medium mesh, with a shape maintained by retaining plates.

properties can suffice to make these devices, a common route is to go one step further and remove (sacrifice) the layer immediately underneath the device with a selective, isotropic etch, thereby suspending the structure in air. For the most part, this design constraint severely restricts the use of a wide range of materials, like the metal oxides ($LiNbO_3$, $BaTiO_3$), group IV (SiC or diamond), III-V (GaN), II-VI semiconductors (PbS, CdS), and other high quality single crystals (quartz, sapphire) due to the difficulty of growing thin films ($\sim$1 μm) on host substrates at requisite quality and crystallinity. Nonetheless, the push for chip-scale devices made out of these materials has not faltered, with recent demonstrations of thin film technologies (i.e. crystal-ion slicing for complex metal oxides[105,82,81] and diamond[196,108], alternative preparation from the bulk[66,183]) enabling the fabrication of nanophotonic devices[179,177,101], albeit at lower quality than commercially available bulk substrates.

Outside of the thin film paradigm, several pattern transfer techniques with three-dimensional control have been developed. A commonality across a number of these methods (which this work shares) is the modification of the ion angular distribution during etching. This includes techniques such as Reactive Ion Beam Etching (RIBE)[102,36], ion-sheath sculpting[170], passivation gas flow and DC bias control[176], and focused ion beam etching[9,14,197]. A second class of methods relies on an anisotropic etch followed by a second isotropic etch, where the layer of interest is shielded from etching by a protective coating[119]. Although this technique requires favorable etch chemistries, it has recently seen extension to new platforms like diamond[89,88].

In this chapter, we study the physics and ion dynamics of a new etching configuration - Faraday cage angled-etching (FCAE). Specifically, a Faraday cage is placed inside an ICP-RIE etch chamber during processing, with the sample placed therein. After the plasma is struck, an equipotential develops at the cage boundaries, resulting in a field-free region inside the cage with electric fields pointing normal to its faces[18,38,148,83,28]. Ions entering from the plasma then get directed toward the sample at an oblique angle of incidence, undercutting the structure and freeing it from the substrate upon completion (Fig. 3.1(b)), effectively fashioning a device from the bulk. Building on initial demon-

strations in diamond single crystals[28,15], free-standing mechanical cantilevers are shown in silicon and quartz. The effects of varying cage parameters on silicon angled-etching are observed, driving cage optimization. Finally, multiphysics simulations of the dynamics of FCAE elucidate the physical processes affecting etch performance.

3.1.1 Fabrication methods

Faraday cages built for etching typically consisted of an aluminum base plate with an aluminum mesh arranged around the sample etching region in the shape of a cone or a triangular prism, though in principle other geometries may be interesting[38,83]. Fig. 3.1(c) shows several different cage constructions. Meshes (TWP Inc.) used in the experiment had wires of ∼400 μm diameter at a ∼1.6 mm pitch (coarse), ∼250 μm diameter and ∼635 μm pitch (medium), or ∼50 μm diameter and ∼125 μm pitch (fine). Coarse and medium sized meshes could be pressed and molded into shape, while fine meshes required some underlying structure to provide support. The cages were constructed in a variety of sizes, with the only significant limitation being the height of the etching tool's load-lock (the smallest cage was ∼7 mm tall and the largest ∼18 mm).

Fig. 3.1(c,i) shows a simple, folded mesh over a base support, and Fig. 3.1(c,ii) shows a coarse meshed wrapped around a large aluminum base. Panel (iii) shows a molded-mesh cage design. The mesh is created by pressing it against two custom metal dies, although it can also be done by hand. Because of the modularity of this design it is easy to swap out meshes constructed for different etch angles. The retaining ring around the mesh (Fig. 3.1(c,iii)) provides stability and reduces movement during loading and unloading; likewise, directly affixing the mesh to the carrier wafer via an etch-compatible adhesive maintains angled-etching functionality and placement stability.

Angled-etching of devices follows a process flow similar to standard lithography[28]. First, a hard mask is defined either through photolithography or electron beam lithography (Fig. 3.1(b, i)). Once defined, the sample is etched vertically in order to ensure clearance from the substrate (Fig. 3.1(b, ii)).

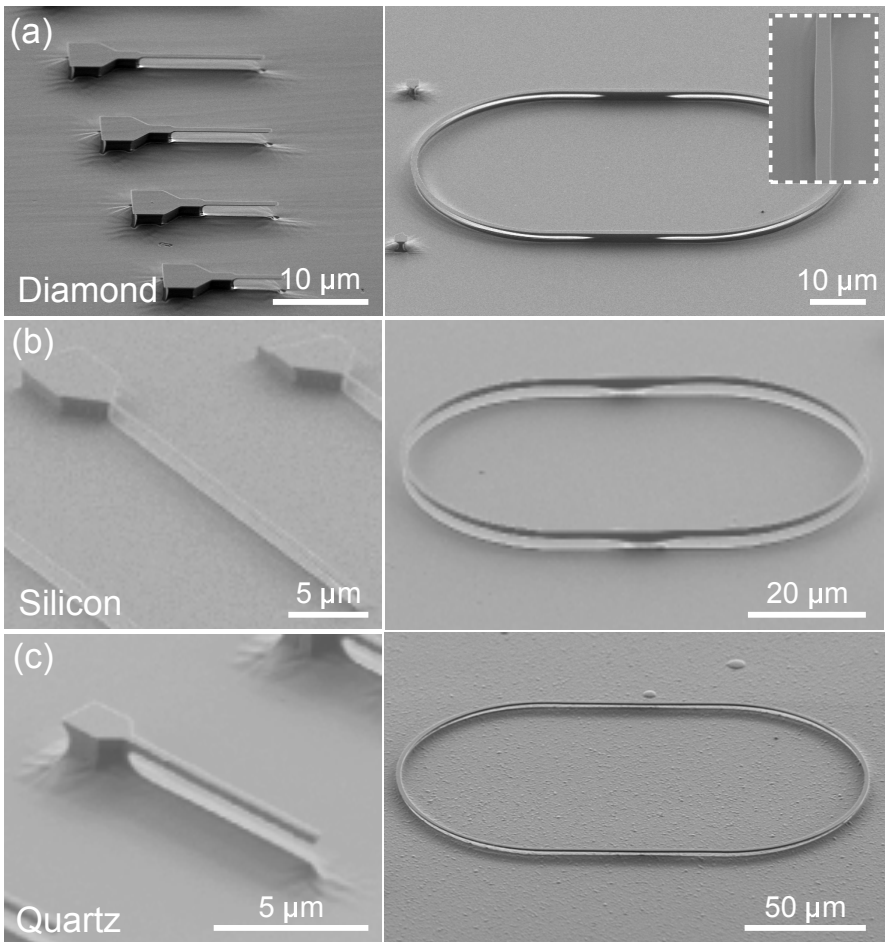

Figure 3.2: Angled-etched devices (left, nanomechanical cantilevers, right, optical resonators) in different materials. (a) Structures etched in diamond, following the recipe in the text. The lines in the substrate visible on the cantilever device are from the coarse cage used to etch, which impacts the ion trajectory. The optical resonators are supported by flared-out sections in the mask, which transfer to pedestals for an appropriately-timed etch.[25]. Silicon (b) and quartz (c) cantilevers, $\sim 40\mu\mathrm{m}$ and $\sim 10\mu\mathrm{m}$ long, respectively. All the optical resonators pictured support modes in the telecom wavelength range.

After this step, the sample is placed inside a Faraday cage and etched at an angle, often with the same recipe used to etch vertically (schematically shown in Fig. 3.1(b, iii)). Typically, the extent of etching is recorded via scanning electron microscopy (SEM) in order to time the release of the structure precisely. This becomes especially important for devices where thin supporting regions underneath the device need to be maintained [25,100]. After this step is run to completion, the mask is removed.

Diamond angle-etched devices were first demonstrated for quantum photonics applications [28,67,15] and have since shown promise as nanomechanical resonators [30] and high quality factor optical cavities [25]. For completeness, we describe the process here. A smooth, polished (<1 nm rms roughness) diamond surface is first cleaned in a refluxing mixture of equal parts perchloric, nitric, and sulfuric acid [8]. After this, it is placed in a boiling piranha mixture (3:1 sulfuric acid to hydrogen peroxide) before being rinsed and sonicated in solvent, then dried. A thin layer of titanium (∼15 nm) is then deposited on the surface. This helps with resist adhesion and charge compensation during the electron beam lithography. A negative-tone electron beam resist is spun on (FOx-16, hydrogen silsesquioxane, HSQ) to a thickness of ∼1 μm.

After the resist is exposed and developed, the titanium layer is first removed with a brief Ar/Cl_2 plasma etch (Unaxis Shuttleline) at forward power of 250 W and RF power of 400 W at 8 mTorr. This ensures the smooth removal of the titanium underlayer. The sample is then etched with an oxygen chemistry flowing at 50 sccm and maintained at 10 mTorr. The forward and RF powers are held at 100 W and 700 W, respectively. Once the vertical etch has been completed, the sample is placed in the appropriate Faraday cage. The same oxygen etch is run, except for an additional slow flow (2 sccm) of either Ar or Cl_2 to mitigate micromasking [28]. Once the etch reaches completion, the mask is removed in HF. For sensitive applications, critical point drying can increase the device yield.

Free-standing silicon devices are important in a number of applications and are already made possible by the readily available silicon-on-insulator platform. However, this approach may be inadequate for some select uses, including mid-infrared photonics [161,111,112] or mechanical resonators

with small vertical cross-section[114]. Angled-etching of silicon can circumvent the material thickness restrictions inherent with SOI technology, realizing structures with both lateral and vertical dimensions on the order of tens of nanometers[100]. Angled-etching of silicon cantilevers (Fig. 3.2(b)) was based on an SF_6/C_4F_8 plasma chemistry (STS MPX/LPX RIE). Two mask materials were shown to have sufficient selectivity: a lifted-off sputtered alumina mask, well known to be highly selective to this chemistry, and a HSQ mask. Fluorine chemistries are commonly used to etch silicon but the formation of a passivation layer must be well-controlled in order to prevent chemical undercut. The mask was removed with a standard hydrofluoric (HF) etch or a vapor HF etch, in the case of HSQ. The vertical etch flowed 130 sccm of C_4F_8 and 80 sccm of SF_6 at a pressure of 10 mTorr, RF power of 1000 W, and platen power of 12 W on a silicon carrier wafer.

The same plasma parameters were used for the angled etch but a quartz wafer was used as the carrier. This reduces the loading due to Si etching and effectively frees more SF_6 molecules for etching. Because the angled etch of SF_6 radicals is directional while passivation layer formation is isotropic, a shadowing effect arises from the blocking of etching ions by the structure. In order to compensate for this to allow for a reasonable etching rate, more SF_6 must remain in the plasma, either by increasing the flow or switching carrier wafers. Using this recipe, silicon waveguides and high quality factor optical resonators have recently been demonstrated[100].

Quartz is an interesting material platform, widely used as a mechanical resonator in micro-electromechanical systems (MEMS) and other technologies[31]. However, lacking a native thin-film technology, it is difficult to integrate into nano-scale systems. Motivated by its excellent material properties, we realized quartz nanobeam cantilever resonators via angled-etching[159] (Fig. 3.2(c)). A metal hard mask was patterned using a bi-layer poly(methyl methacrylate) (PMMA) liftoff procedure (cantilevers) or direct etching using an e-beam resist as a mask (optical resonators). The pattern was then transferred to the quartz using a CHF_3-based recipe at 10 mTorr pressure (STS MPX/LPX ICP-RIE). The ICP power was held at 600 W and the platen at 90 W. Argon, CF_4, CHF_3, and H_2 were all flowed at 6,

2, 50, and 15 sccm respectively. Due to the physical milling introduced by angled-etching and the small loadlock (∼8 mm) on this etcher, the angled-etching was completed in a different machine which could accommodate a larger cage design (Nexx ECR RIE). This etcher relies on a different mechanism to generate a plasma (electron-cyclotron resonance), yet was still shown to be suitable for angled-etching. The microwave power was set to 600 W while the platen RF power was set to 90 W and the pressure was held at 10 mTorr. The ratio of gasses remained the same, with the overall flow rate reduced by a factor of two to accommodate machine constraints.

3.1.2 Etching simulations

The effect of the Faraday cage on the reactor potential was investigated with COMSOL Multiphysics using the plasma physics module[75]. This simulation technique has been shown to yield results of reaction parameters with reasonable agreement to experiment[75,43] and can be extended to couple to models describing feature profile evolution[74]. The modeling was preformed using a standard Gaseous Electronic Conference (GEC) reference cell with argon gas, a chemistry which qualitatively captures the impact of the presence of the Faraday cage. The simulation was performed under axial symmetry, with the cage wires forming rings around the sample. No biasing was coupled into the simulation, so the forward power remained at effectively 0 W. Figs. 3.1(a) and 3.3(a) show the potential distribution of the plasma with a cage placed inside. The cage wires are maintained at ground while the plasma potential develops in the simulation. The etcher frequency used in the simulation is the standard 13.56 MHz. The coil current is maintained at 100 A, and the power within is monitored until convergence (typically measured ∼2000 W). As expected, the plasma potential decreases as the plasma gets closer to the cage, creating electric fields normal to the cage face. Towards the bottom of the cage, the plasma sheath starts to resemble the sheath at the carrier wafer, flattening out. As the sheath moves up the cage, it follows its contours more closely, resulting in electric fields that are more normal to the cage face.

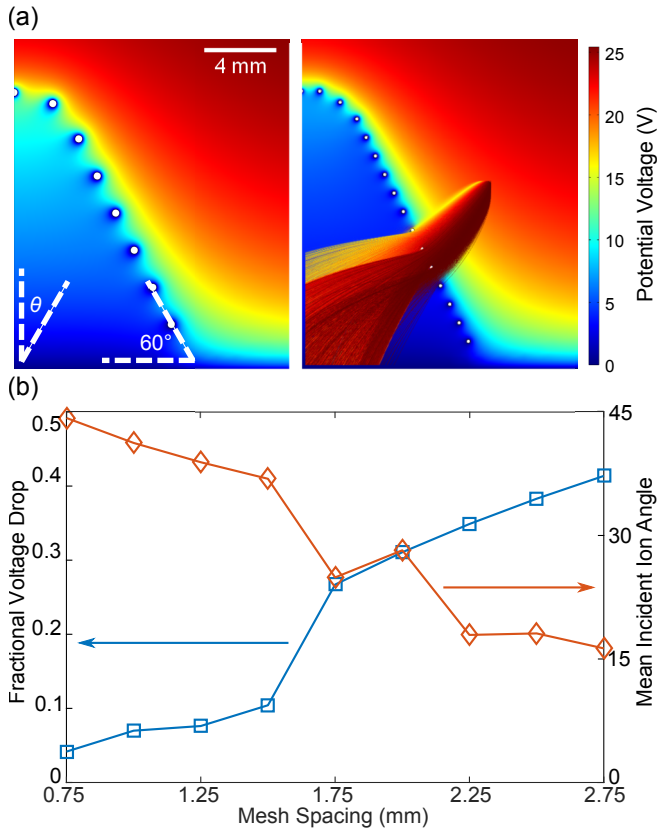

Figure 3.3: Simulated voltage drops inside the Faraday cage. (a) Simulated etching potentials for a coarse cage (left) and fine cage (right). The coarse cage has a larger potential gradient within, influencing the path the ion takes. The simulated ion trajectories from a single point outside the cage are overlayed on the right hand side, with coloring to provide visual contrast. All simulated cages have a cage angle of 60 degrees. (b) Voltage drop inside the cage relative to the etch potential (squares, left axis) and resultant mean incident ion etching angle, defined from the normal (diamonds, right axis).

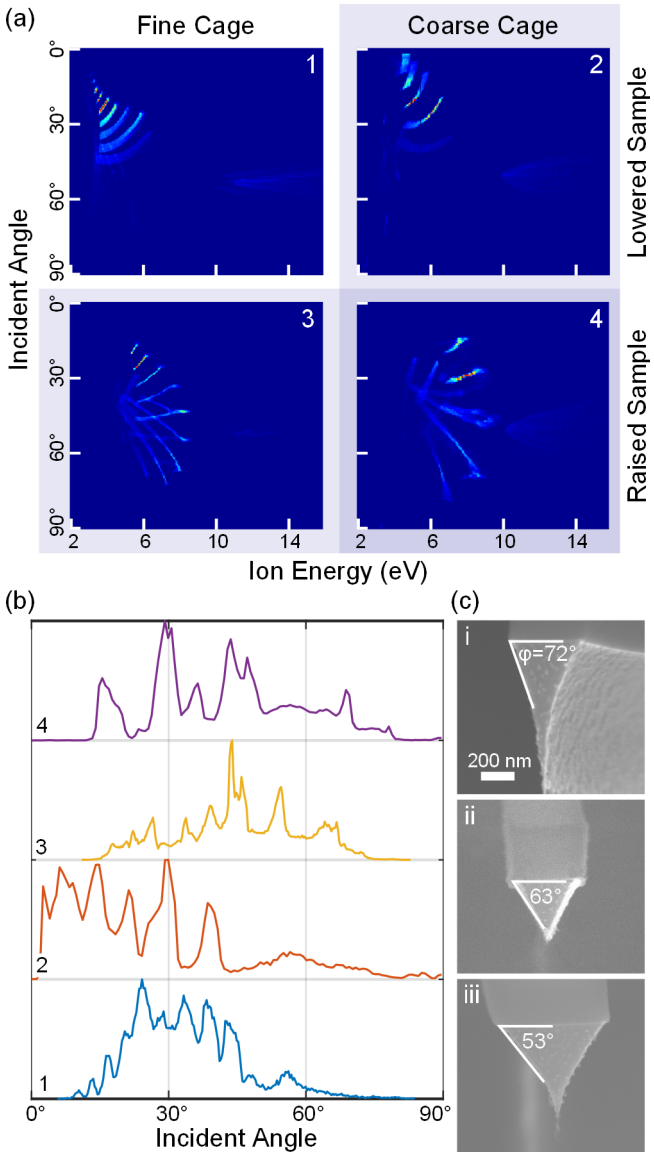

Figure 3.4: Detailed results from changing cage parameters. (a) Incident angle/ion energy histograms for simulated cage designs with fine (left) and coarse (right) mesh spacings and raised (bottom) and lowered (top) samples. The ion data are taken upon collision with the sample area. The angle-energy histograms show clear "banding" of the ions due to the effect of the Faraday cage wires. Notably, both the raised sample position and coarse mesh spacing increase the variability in angle/momentum space. Generally, higher ion energy (equivalently, momenta) can negatively effect mask selectivity as the etch becomes more physical. (b) Averaged ion incident angle for the cage designs in part (a), corresponding to their number. (c) Head-on SEM images of cantilevers etched in (i) a raised, coarse cage, (ii), a lowered, fine cage, and (iii) a raised, fine cage. The angle ϕ is defined as the observed etch angle. The SEMs show that ϕ varies in accordance with simulation. The visible sidewall roughness seems to be a characteristic of the particular etch and is not greatly affected by cage design.

By changing the simulated cage geometry, we can investigate the effects of different cage types on the potential distribution using straightforward simulation parameters. In particular, we change the mesh parameters by changing the thickness of the wires or their spacing. Another parameter of interest is the sample height within the Faraday cage. This is changed by placing a metallic (grounded) block within the cage, and then recording the final ion velocities upon hitting the top surface, where the sample presumably lies. Throughout the simulations, the mesh face angle is maintained at 60 degrees. The impact of the mesh angle is not explicitly studied, as the etch angle has been seen to be proportional to the mesh face angle previously[28].

To study the action of the ions under the cage potential, the motion of ions was simulated using a particle tracing module in COMSOL. After releasing the ions (argon) from the sheath with a Maxwellian velocity distribution (at a temperature of 400 K), their velocity and position is charted. Once the ion intersects with a defined sample region, its velocity vector and position is recorded to generate an angle/energy ion distribution map. This simulation procedure is repeated with the different cage geometries, showing stark differences in the incident angle distribution, as well as ion energies.

3.1.3 Faraday cage etching: lessons learned

Synthesizing the simulation results and incorporating experimental evidence, we can make definite statements on the plasma physics and ion dynamics, whose qualitative ideas transfer to the etching of many other material types in FCAE. To begin with, the simulations (Fig. 3.3) show that coarser meshes allow the leakage of potential past the mesh, resulting in a secondary ion sheath inside the cage. This implies the existence of a secondary electric field pointed normal to the carrier wafer, after the initial electric field which accelerates the ions and gives them their angle. This secondary field is detrimental to the function of the cage, adding a vertical component to the incident ion angle. In Fig. 3.3(a), the simulated cage potentials show the elevated internal potential in the coarsely-spaced

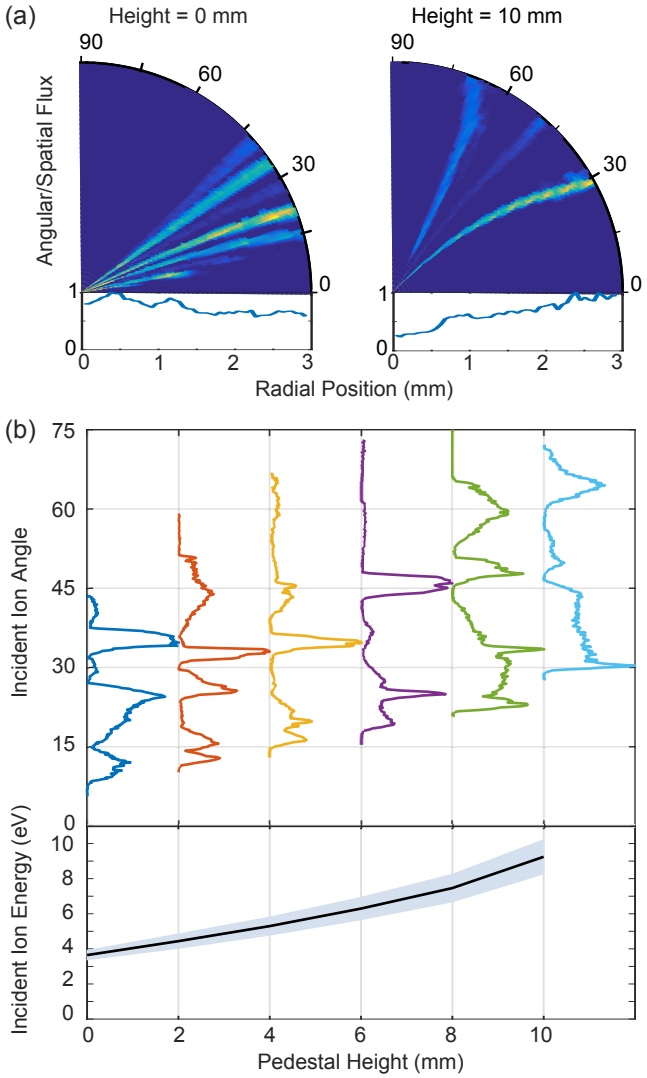

Figure 3.5: Impact on sample height and incident angle across a sample. (a) Incident ion angle as a function of radial position on the sample for samples on the carrier (left) and placed on a 10 mm pedestal (right). The bottom shows the relative flux of etching ions as a function of position. The sample placed on the pedestal shows a larger variation in ion flux and incident angle. (b) Incident ion angles for a cage as the sample is placed on higher pedestals (top). Mean ion energy and standard deviation for ions as the sample height is raised (bottom)

mesh. Fig. 3.3(b) makes the link between the internal secondary sheath and ion incident angle explicit. As the mesh spacing increases, the fractional voltage drop (defined as the maximum voltage within the cage divided by the maximum voltage outside the cage) increases to as high as 40%. Likewise, as this internal voltage drop increases, the average incident ion angle decreases, as the ions are deflected to a more vertical incidence (here defined as zero degrees).

By looking at the ion distribution in both incident angle and energy, as in Fig. 3.4(a), we can fine-tune our observations. An ideal Faraday cage design would show a narrow distribution of ions in both angle and energy space, implying that all incident ions are similar in both direction and momentum, providing a well-defined and controlled etch. In reality, we see that finer meshes are seen to create a larger incident etching angle on the surface, closer to the value prescribed by the cage geometry, although they also create tightly-defined "bands" of angles. These bands are a result of the local deflection of the trajectory of the ions due to the mesh wire itself[148]. Finer meshes have more wires per unit area, resulting in more of these bands. However, the larger wire sizes of coarser meshes tend to bend the trajectory for a longer time, increasing the glancing angle. To simplify the interpretation of these results, we project the ion distribution onto the angular axis, averaging over the different ion energies, as in Fig. 3.4(b). The corresponding distributions carry the same label as in part (a). These simplified charts show the decrease in standard deviation of incident angle afforded by finer meshes (labels 1 and 3). They also simultaneously show the increase in mean ion etch angle as the sample is raised within the Faraday cage, an effect which is further explored in Fig. 3.5. In a real chamber, a DC bias is expected to increase the electric potential gradient, resulting in a mean incident ion angle closer to that prescribed by the Faraday cage. The stronger electric fields should also reduce the perturbation on the ion trajectories caused by the cage wires, resulting in narrower ion distributions.

To support the results from these simulations, silicon cantilever samples were fabricated with the recipe outlined beforehand and imaged at head-on incidence to record the etch angle, defined as ϕ in

Fig. 3.4(c). A coarse mesh was observed to increase the verticality of the etch (larger ϕ, Fig. 3.4(c,i)) even though the sample was raised. Positioning a sample lower while using a fine mesh made the etch more vertical (Fig. 3.4(c,ii), while raising the sample within a fine mesh gave the most acute etch angle, Fig. 3.4(c,iii). This corroborates exactly the results from simulation, where the incident angle changes starkly based on sample positioning and mesh type (Fig. 3.4(a,b)).

The etch rate and etch angle at different points inside the conical Faraday cage was experimentally shown to vary dramatically due to the asymmetry of the cage itself and off-center placement of the sample within. In order to mitigate the effects of this so-called etch gradient, the sample was etched at intervals of 30 seconds, with rotations of the cage in-between [26]. Although this process significantly symmetrizes the etch angle across the device, the etch rate was still observed to be variable, with the highest etch rate in the middle of the device. This is supported by the simulations of the ion trajectories in Fig 3.5(a). Toward the center of the sample, there is a large flux of ions coming not only from the side of the cage (at an angle), but also from the top of the cage, where it curves to become horizontal (directing ions downward). Capping the top of the cage with a shield can prevent these stray ions from entering the cage. Experimentally, it was also observed that cages which resulted in more acute etch angles ϕ had larger etch gradients. This can also be observed in the simulation data, as the angle/ion energy histogram displays a large spread of incident ion angle for samples positioned higher and with tighter meshes. Looking across the sample, Fig. 3.5(a) shows just how sample position affects ion properties. For low-positioned samples, the angle distribution of ions is constant across the sample. There is some modulation of flux, as seen in the bottom chart, possibly due to the effects mentioned beforehand. Highly-positioned samples show a larger variation in flux and a large change in incident ion angle, especially towards the edge of the sample. In general, from Fig. 3.5(b), it can be seen that higher-positioned samples tend to not only have higher-energy incident ions, but a larger spread of incident angles and energies as well, possibly leading to less-controlled etching conditions.

It is a well-known result from ion-beam etching that physical mask milling changes with changes in etch angle. Within FCAE etching, the selectivity of the etch decreases as more acute etch angles (ϕ approaching 0) are attempted[18]. This effect is seen qualitatively in diamond devices using both HSQ and alumina as a mask. Furthermore, because of the milling action that occurs when etching ions are incident on the mask, acute etch angles are also associated with an increase in micromasking, roughening the surface of the etched device significantly. Likewise, increasing the mask coverage of the sample leads to an increase in micromasking as more mask particles are resputtered into the chamber. This can be mitigated by designing patterns with less mask area.

With this in mind as well as the results from Fig. 3.5(a), angled-etching may also be used to non-invasively engineer continuous etch profiles on a sample surface. This can be accomplished by appropriately shaping the design of the Faraday cage. In our simulations, we have shown that even a cage without bends (and therefore a nominally constant radius of curvature equal to zero) can create a spatially-varying angular and flux distribution in ions, thereby modifying the local etch rate. Adding curved components in the Faraday cage can create a lensing effect, in analogy with electron and ion optics. This can be used to three-dimensionally pattern a surface without any lithographic steps[96].

3.2 Reactive ion-beam undercut etching

Although an excellent method for prototyping on a small scale, Faraday cage etching suffers from severe uniformity and repeatability issues, limiting its use outside of research. We explore a new technique for creating freestanding photonic and mechanical nanostructures via reactive ion beam etching (RIBE), which provides excellent uniformity and wafer scale processing. RIBE is a derivative of ion beam etching (IBE) or ion beam milling, where a broad area ion source is used to collimate and direct a beam of high energy ions of a noble gas such as Ar, Xe or Ne towards a target to physi-

cally remove material via sputtering. Ions are extracted from a plasma source by a set of electrically biased grids. Voltage applied to these grids controls the ion energy and the divergence of the ions within the beam. The ion flux exiting the source, or the beam current, is determined by the ion density set by the plasma source. Therefore, in an IBE system, the ion energy, flux and divergence can be independently set over a wide operational range enabling a uniquely unrestricted ion etching process. The use of reactive gases allows for more versatile etching of an expansive list of materials, otherwise difficult to etch with traditional IBE relying on physical sputtering. Fundamentally, an RIBE system is an IBE system whose inert gases have been replaced by reactive gases, such as chlorine or oxygen. The distinction comes in how the plasma is generated. In typical Kaufman DC ion sources, a hot filament is used to create a plasma which is incompatible with reactive gases. For this reason, RF sources are utilized for reactive gases, either Electron Cyclotron Resonance (ECR) or Inductively Coupled Plasma (ICP). CAIBE is a version of reactive ion beam etching, however it still uses a DC ion source of inert gases (Ar, Xe or Ne), while injecting a directed spray of the reactive gas towards the etch target. The kinetic energy of the inert gas ions exiting the ion source creates collisional fragmentation with the reactive gas and the resulting ions are directed towards the sample to perform the etch. Both techniques can yield great results, however the etch mask selectivity and overall wafer scale uniformity are better with RIBE[156].

Fig. 3.6 depicts the fabrication procedure to create free- standing optical and mechanical nanostructures via reactive ion beam undercut etching, from now referred to as Reactive Ion Beam Undercut Etching (RIBUE). First, an etch mask is patterned onto the sample surface (Fig. 3.6(a)) followed by a top down etch with the sample mounted on a rotating sample stage perpendicular to the ion beam path as seen in Fig.3.6(b). Once the desired vertical depth is achieved, the sample is tilted to obtain an acute angle with respect to the collimated ion beam, as seen in Fig.3.6(c). Undercut etching is performed with the sample stage rotating, uniformly etching underneath the mask. Once the desired undercut has been achieved the etch mask is removed yielding a free-standing nanostructure

as seen in Fig.3.6(d).

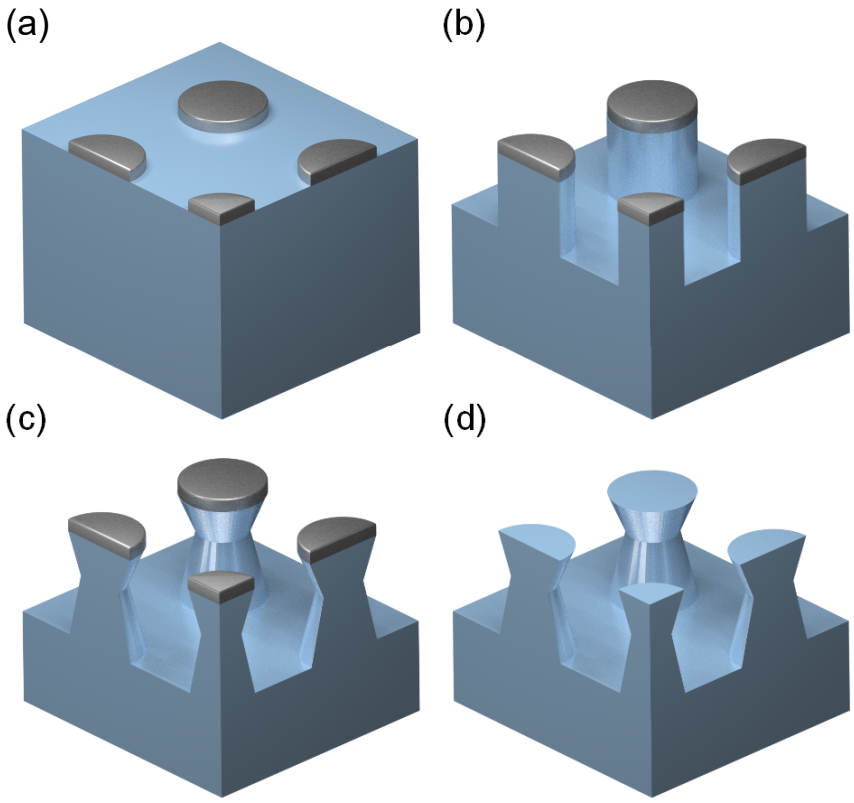

Figure 3.6: Schematic of reactive ion beam undercut etching process. (a) Mask (typically niobium) is defined via pattern transfer. (b) The sample is mounted perpendicular to the ion beam path and etched on a rotating sample state. (b) The sample is tilted to the desired etch angle and the stage continues rotating, uniformly etching underneath the etch mask. (d) The mask is removed, yielding an undercut structure. Etching may continue until the structure underneath is completely etched away. By varying the mask size, support regions may be defined, yielding freestanding nanostructures on a bulk substrate.

3.3 Thin film fabrication

The final diamond fabrication process in common use in the Lončar lab is termed "thin film fabrication", referring to the preparation of a thin film of diamond from bulk by painstaking polishing, etching, and transfer. This technique, although it has the lowest yields and requires the most labor,

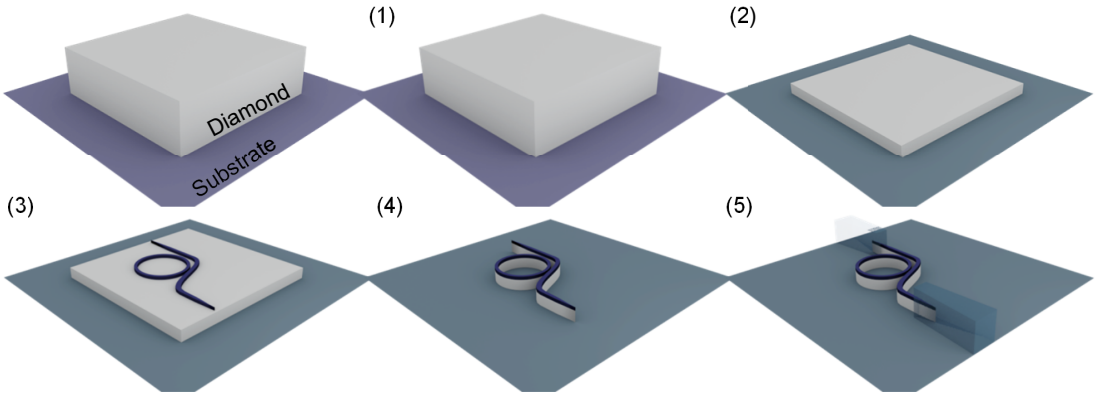

Figure 3.7: Schematic flow-chart of thin film fabrication. (1) The diamond is cleaned and placed on a sapphire substrate. (2) The diamond is thinned down via ICP etching and transferred to a new substrate. (3) Resist is patterned on the diamond. (4) The sample is etched. (5) The coupling waveguides which go to the end of the chip are defined and the sample is cleaved.

has resulted in the highest quality factor diamond microresonators, from visible to the telecom. The core motivation of this technique is to reproduce a diamond thin film architecture as if the diamond were grown as a single crystal on the substrate of choice - leading to a diamond-on-insulator (DOI) device. Unlike the angled-etching methodologies, DOI devices are attractive because they emulate photonic architectures which are already well-developed, from coupling waveguides to spot converters. However, they also present researchers with challenges, from fabrication to design.

For convenience, a summary of the fabrication process follows in 3.3, presented without detailed notes.

1. Diamond preparation
    (a) Clean received thin diamond plates in refluxing acid mixture (1:1:1 nitric:perchloric:sulfuric acid) for 2 hours
    (b) Retrieve thin diamond plates and rinse in DI water
    (c) Clean in piranha acid mixture for 15 minutes

(d) Retrieve and rinse in DI water, then in methanol

(e) Place diamond plates on sapphire carrier wafer, gently blowing dry with nitrogen gun

- If the sides of the plate are not perfectly vertical, place the plate so that the side with larger surface area is facing up
- Record which side is facing up

2. Diamond etching

(a) Complete one etch cycle

i. 30' Ar/Cl_2 etch at 400 W ICP power and 250 W forward power

ii. 25' O_2 etch with 700 W ICP power and 100 W forward power

(b) Remove the diamond from sapphire carrier wafer by placing drops of water on diamond

(c) While diamond is still wet, transfer via tweezers diamond from carrier wafer to prepared piranha solution

(d) Repeat steps 1c to 1e

- Before drying diamond on sapphire wafer, inspect and ensure diamond is facing opposite direction
- The side with less surface area should be facing up: if the diamond plate is situated such that there exists a slight overhang, resputtered material can concentrate underneath, which, upon becoming exposed to the etch, can be re-sputtered again and cause imperfections on the etched diamond surface

(e) Continue etch cycles until diamond is <2 μm thin

(f) Finish etching diamond to desired thickness with O_2 etch only

3. Substrate preparation

(a) Clean and sonicate substrate in solvents, blowing dry

(b) Immediately before diamond transfer, activate substrate surface with O_2 plasma (300 mT, 75 W, 1 minute)

4. Diamond Transfer

   (a) Place single drop of hydrofluoric acid on thinned diamond plate

   (b) Wait 2-5 minutes, ensuring hydrofluoric acid does not completely dry from film, and replenishing if need be

   (c) Drop water on diamond, which should easily debond from the sapphire carrier, if not, gently nudge film with tweezer while flowing water over it

   (d) Dilute the residual hydrofluoric acid out with water

   (e) Bring both diamond and receiving substrate to edge of wafer, using surface tension to steer the diamond

   (f) Bring receiving substrate into contact with water containing diamond, using gravity and surface tension to transfer diamond from sapphire to substrate

   (g) Pipette off excess water from substrate and position diamond film using tweezers or nitrogen from gun as guide

   (h) Finish blowing dry diamond from substrate, at which point the diamond should be firmly bonded

5. Device processing

   (a) Deposit 2-4 monolayers of ALD film, preferably silicon dioxide, for adhesion

   (b) Spin FOx-16 resist at 4000 RPM

   (c) Bake device for 5 minutes at 90 C

   (d) Spin E-Spacer at 4000 RPM

   (e) Expose device in electron-beam tool (dose 3000 $\mu C/cm^{-1}$, 125 keV accelerating voltage)

   (f) Develop in TMAH 1 minuted, followed by DI rinse, then leave in methanol for 3 min-

utes, finally blow drying with nitrogen

(g) Deposit 2nm of ALD oxide on device for adhesion

(h) Anneal device at 450 C for 1 hour in oxygen atmosphere to anneal ALD film

(i) Prepare doped-glass mode converters per section 3.5

(j) Spin protective photoresist layer

(k) Partially dice chip on backside and cleave

(l) Remove resist and clean substrate

Notable peculiarities of the fabrication process are the lack of a thick protective layer on top, as well as the preservation of the hard etch mask. In early iterations, a few capping materials were explored. These have the advantages of both protecting the nanostructures, as well as reducing the index contrast of the waveguide with the surrounding material, reducing scattering losses. PECVD oxide is a generally popular choice, but its use in these structures is forbidden because of its relatively high refractive index - in fact, due to its high silicon content it has a higher index than the doped-glass waveguides developed in the next section. Polydimethylsiloxane (PDMS) was also explored as a capping material, due to its ease of use, low refractive index, and transparency at visible, especially compared to other polymer materials. However, at the high pump powers necessary for nonlinear optics, devices capped with PDMS quickly failed, as the PDMS absorbed heat and began to burn. In the end, only a very thin ALD oxide layer is used as a capping layer, mostly for adhesion of the waveguides to the substrate.

The preservation of the hard mask is another interesting feature of this fabrication. This step is necessary because all attempts at removing the hard mask (HF, diluted HF, very slow vapor HF, other wet and dry etch chemistries, sacrificial underlayers) were unsuccessful, either roughening the waveguides or causing the waveguides to delaminate from the substrate. In the end, the preferred approach was to make sure that the remaining hard mask had sufficient optical quality, as done through annealing and surface treatment.

3.4 Diamond etch mask choice

Unlike the plethora of etch chemistry choices that face the silicon process engineer, the choice for diamond is relatively limited. Although an Ar/Cl_2-based etch is used to smooth diamond in-between oxygen etch cycles[8], the low selectivity of the etch makes using it in structure definition impractical (though it is often used to define the hard masks which we will discuss). The etch rate of this process is also relatively slow, etching 60 nm/minute. This leaves an O_2 plasma etch as the etch chemistry of choice for diamond nanostructure definition. This process was finely tuned to give straight sidewalls and good selectivity for a variety of masks, with typical parameters of 50 sccm O_2 in a 4 mT pressure chamber and ICP and forward powers of 700 and 100 W, respectively[65]. The etch rate of diamond is typically 160 nm/minute with good uniformity and run-to-run stability. This process is used across etch modes, for thin-film fabrication, nanostructure fabrication, and even undercut etching. In some instances, when micromasking or redeposition is an issue, a small amount (typically 2 sccm) of a more aggressive gas is flown into the chamber, typically argon though chlorine has been used in the past (chlorine sees limited modern use because sub-surface chlorine implantation during etching can cause charge instability in diamond color centers and even contaminate down-stream processing tools)[40].

With this constrained set of etch chemistries - oxygen with an optional flow of argon - the main flexibility in process design comes from judicious choice of mask. A review of the oxygen etch chemistry with gold, alumina, and oxide masks has been preformed previously[65]. Of these, alumina nanoparticles performed the best, but an easy deposition technique was lacking.

The following table summarizes new masking strategies tried over the course of this PhD. All masks were defined via etch-transfer, barring the flowable oxide (FOx) mask - they were first deposited using the denoted process, then a new mask was written on top in the denoted resist (via electron- or photo-lithography, as appropriate) and their qualitative results are summarized.

Mask	Dep.	Soft Mask	Definition	Results
FOx-16	Spin-coat	N/A	TMAH development	Most reproducible mask with acceptable sidewalls, used in all photonics experiments. Low selectivity. Adhesion issues remediated by depositing either thin (15 nm) titanium layer or thin ALD of silica or alumina.
Oxide	PECVD	MaN 2403	CHF_3 or C_4F_8 etch	Relatively rough due to poor pattern transfer from MaN to oxide. More selective than pure-FOx
Oxide	PECVD	FOx-16	CHF_3 or C_4F_8 etch	Roughness equivalent to FOx-16, but more selective
Alumina	Reactive Sputter	SPR 220 PR	Cl_2/BCL_3 alumina etch	Roughness equivalent to FOx-16, excellent selectivity, poor critical dimension transfer. Suitable in applications where critical dimension tolerance is loose
Gold	Electron-beam dep.	Shipley 1805 PR	Liftoff	Relatively rough, moderate selectivity. Only good in situations where liftoff is a convenient fabrication method and tolerances are loose, alumina masks are better in all other regards.
Niobium	Sputter	FOx-16	Ar/Cl_2 etch	Best selectivity in O_2 chemistry (>50:1). Sidewalls about as smooth as FOx-16 only, but redeposition of mask material is a larger concern.

Table 3.1: Summary of different fabrication approaches towards diamond thin-film devices

3.5 Mode-converter fabrication for high-power applications

Due to the small size of the diamond substrate, waveguides which extend to the end of the sample cannot be natively written and defined. To couple light efficiently onto the diamond sample, we developed a doped spin-on glass process which allows for directly-writeable optical waveguides. This represents an evolution of previous polymer-based waveguides [63], as it has improved power-handling ability at visible wavelengths.

This waveguide fabrication scheme has significant advantages for integration with delicate devices, as in the case of diamond. Because the process is entirely done without any damaging plasma etch steps, the quality of the waveguide surface is preserved. Furthermore, the reduction of steps reduces the fabrication overhead and possibility for errors to accumulate. However, it requires a careful choice of materials, as polymer waveguides (as used for telecom devices [63]) can easily burn up. To this end, an electron-beam sensitive resist must be chosen which has sufficiently high refractive index to provide the contrast necessary to guide optical modes. Although metal alkoxide-based sol-gels (such as those that result in alumina or titania structures when processed) have been used as high -selectivity masks [152,153], their stability in air is low and creating the requisite thick waveguides is difficult. For this reason, hydrogen silsesuioxane (HSQ), a common spin-on-glass, was used as the base chemistry for the waveguides. Because annealed HSQ tends to have lower refractive index than glass [13], titanium butanate added to the HSQ compensates. The ratio of HSQ to titanium butanate is controlled to maintain acceptable stability in air while still providing the requisite refractive index. Such sol-gel solutions are prepared immediately before use, having a typical lifetime of <30 minutes.

The exact fabrication process is as follows. After defining and etching a diamond device, alignment marks which were written with the diamond structures are first masked off with Kapton tape. A sol-gel solution is prepared by mixing FOx-16 and titanium butanate ($Ti(OBu)_4$) in a 4:1 ratio. The titanium butanate increases the refractive index of the final waveguides, ensuring the waveg-

uides have a higher index than the substrate. Next, the diamond sample is placed on a spinner and brought up to 1000 RPM. Then, 2 drops of the sol-gel solution are dropped on the diamond sample as it begins to accelerate to 3000 RPM. The sample spins at 3000 RPM for 30 seconds. The sample is not baked, but undergoes an O_2 plasma treatment (300 mT, 75 W, 30 seconds) after the Kapton tape is removed to increase adhesion of conductive polymer to the masked area. To protect the deposited sol-gel during plasma treatment, a dummy piece of silicon is placed on top of the sample, exposing only the area where the tape was removed. The sample is then taken back to a spinner, where E-Spacer (ShowaDenko) is spun on at 4000 RPM. The waveguides are written at a dose of 6000$\mu C/cm^2$. The sample develops for 10 seconds in TMAH, followed by a rinse in DI, then a rinse in methanol. To reduce absorption, the waveguides must be annealed. The anneal was limited to 460 C in O_2 for three hours to avoid burning the diamond. Ideally, a rapid thermal anneal to 1100 C in O_2 can completely oxidize the waveguides[73], though the performance improvement over the low temperature anneal was negligible.

Typically, waveguide cross-sections for these waveguides are defined to be 2μm width by 1.5μm in height. The height is set by the concentration of $Ti(OBu)_4$ and spin speed, while the width is defined during e-beam writing. This cross-section is optimized to ensure maximal coupling into a 710 nm wavelength laser beam focused on the facet. Fig. 3.8 demonstrates a waveguide-resonator system with light coupled in, showing that light is well-guided in the system. The quality factor of the resonator was 40,000, limited by bending losses due to the low index contrast. In a straight end-to-end measurement, the input/output coupling loss was limited to <3 dB/facet, while the loss due to the waveguide was estimated to be between 0.1 and 1 dB/cm.

As an outlook, sol-gel based photonics have potential for more than just passive devices. Here, they fill the role of an easy-to-fabricate interface between an integrated, on-chip photonic component and free-space optics. Sol-gel waveguides can easily serve as a drop-in replacement in commercial photonic systems where such an interface is a common requirement. They can also fulfill the

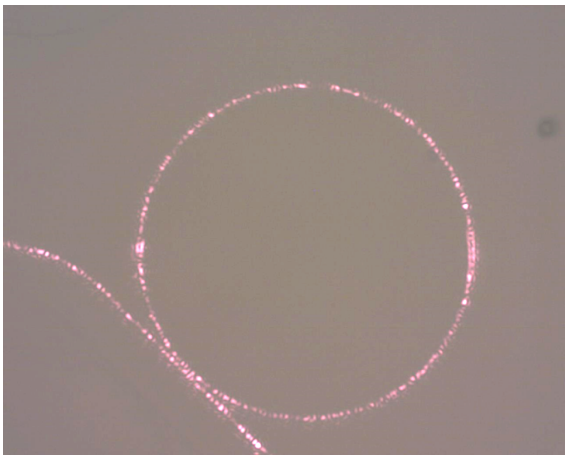

Figure 3.8: Micrograph of sol-gel based resonator with visible light coupled in. The scattered visible light shows little loss. The diameter of the microresonator pictures is 100μm and the recorded quality factor was 40,000.

functions of other passive photonic devices, such as arrayed waveguide gratings. However, the benefits of the platform are not limited there, as the material degree of freedom allows the incorporation of doping ions directly into the waveguide, specifically ions such as Er^{3+}, which are used in telecommunications for signal amplification. Such devices have been reported in the literature, typically of glass-based sol-gels[188,118,175,20], in microtoroid geometries[187,76] and directly integrated on silicon photonics[136] for on-chip lasers. The ease of fabrication and material versatility positions sol-gel photonics is an area ripe for future discovery.

4

Raman lasing in diamond

In contrast to passive or linear optical devices, nonlinear optics concerns itself with the generation of new colors. In this chapter, the focus rests on the Raman process, where a pump photon interacts with a phonon in the material, changing the photon's energy. This process can involve the creation of a phonon (resulting in a down-shifted, Stokes photon) or the absorption of a phonon (resulting in an up-shifted, anti-Stokes photon).

Phonon-mediated light scattering can be further typified into two separate processes: Brillouin

scattering and Raman scattering. The former involves phonons from the acoustic branch, while the latter involves phonons in the optical branch. Because Brillouin scattering relies on acoustic-branch phonons, the energy imparted from the phonon is relatively low (a typical acoustic phonon is on the order of GHz, as the name suggests relating to the speed of sound in the material) Furthermore, the energy of an acoustic phonon is strongly dependent on its wavevector, meaning that in a phonon-mediated scattering process both energy and momentum must be conserved. This additional constraint makes the construction of Brillouin lasers difficult, although there have been excellent demonstrations to-date[155,106,107,154].

The energy of the optical phonon, by contrast, is relatively independent of the wavevector. This means that Raman scattering is an inherently phase-matched process, where as long as the energy is conserved in the process, the momentum will be as well. Additionally, the high energy of the optical phonons means that the frequency shift imparted on the photon is orders of magnitude greater than that for Brillouin scattering (40 THz for diamond, 18 THz for silicon).

4.1 Raman lasers

Under certain conditions, the Raman scattering process can be be amplified, leading to a situation where scattering feeds back on itself in a process akin to lasing. In a material with refractive index at the pump n_P and at the Stokes n_S, we can derive a gain coefficient for such a process, given by[123]

$$g_s = \frac{4\pi^2 \omega_S N T_2}{n_S n_P c^2 m \Omega_R} \left(\frac{d\alpha}{dq}\right)^2 \tag{4.1}$$

where N is the density of atoms, m is the reduced mass, and Ω_R is the angular frequency of the Raman shift. The angular frequency of the Stokes output is given by ω_S. $d\alpha/dq$ is the polarizability of the optical phonon and tends to increase when the pump frequency approaches the bandgap. Looking at this equation, we can see why diamond has the highest gain coefficient of known mate-

rials: the carbon atoms which make up its lattice are densely-packed yet still light. These attributes also explain its high phonon frequency, or Raman shift[123].

We can derive the evolution of the intensity of a pump beam and a Stokes signal shifted exactly Ω_R away with the following set of coupled differential equations, where A_i describes the electric field in mode i and α_i the loss[113]:

$$\frac{\partial |A_p|^2}{\partial z} = -\alpha_p |A_p|^2 - (g_R(\omega)/A_{\text{eff}})|A_p|^2 |A_s|^2 \tag{4.2}$$

$$\frac{\partial |A_s|^2}{\partial z} = -\alpha_s |A_s|^2 + (g_R(\omega)/A_{\text{eff}})|A_p|^2 |A_s|^2 \tag{4.3}$$

Specifically, for a microresonator system like the ones we are dealing with, the threshold of the Raman laser $P_T \equiv |A_p|_T^2$ is given by[165]

$$P_T = \frac{\pi^2 n^2}{\lambda_P \lambda_R} \frac{V_{\text{eff}}}{\Gamma B g_s} \frac{Q_c^P}{(Q_T^P)^2 Q_T^S} \tag{4.4}$$

Where n_P is the effective index at the pump, λ_P and λ_S are the pump and stokes wavelengths, respectively, and V_{eff} is the effective mode volume. Γ refers to the spatial overlap between the pump and Stokes modes, with B a correctional factor for internal backscattering. The quality factors for the pump and stokes are denoted Q^P and Q^S, respectively, where Q_c distinguishes coupling quality factor from Q_T, or total quality factor. The threshold is proportional to Q^2/V, emphasizing the important role device processing has in attaining lasing behavior.

More generally, a Raman laser can occur in any cavity system where there is a buildup of optical energy in at least one of the optical modes, given a high enough quality factor. Tabletop systems (where V_{eff} is much larger) use mirrors to form the cavity system.

A final important consideration for Raman lasers is the direction of emission of the Stokes beam. Fig. 4.1 visualizes scattering efficiency of the Raman process in cubic crystals (in this case, diamond)

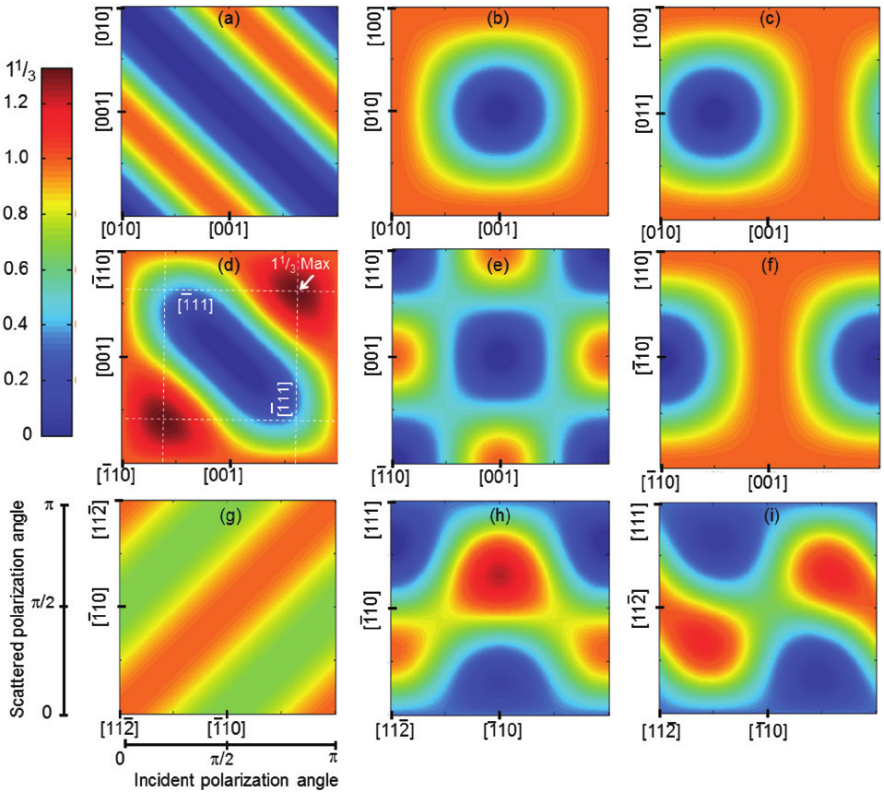

Figure 4.1: Reproduced from Mildren & Rabeau [123] with permission. Scattering efficiency $\left(\sum_j |\mathbf{e}_s \mathbf{R}_j \mathbf{e}_i|^2\right)/d^2$ as a function of input and Stokes beam polarizations for different beam orientations relative to the crystal axis. (a-c) For an incident beam along a <100> axis to (a) axially-directed scattering (b) perpendicular scattering in <100> and (c) perpendicular scattering in <110>; (d-f) For an incident beam along a <110> axis for (d) axially directed (e) perpendicular in <110> and (f) perpendicular in <100> directed scattering; (g-i) For an incident beam along a <111> axis for (g) axially directed, (h) perpendicular in <112>, and (i) perpendicular in <110> directed scattering.

along pump/Stokes directions for different crystal orientations. Depending on waveguide geometry and orientation, the Raman process can be suppressed or enhanced.

4.2 Raman in diamond

Diamond serves as a compelling material platform for Raman lasers operating over a wide spectrum due to its superlative Raman frequency shift (∼40 THz), large Raman gain (∼10 cm/GW @ ∼1 μm wavelength) and ultra-wide transparency window (from UV (>220 nm) all the way to THz, except for a slightly lossy window from $\sim\ 2.6\ -\ 6\,\mu$m due to multiphonon-induced absorption)[123,55]. Furthermore, the excellent thermal properties afforded by diamond (giant thermal conductivity of ∼1800 W/(m K) @ 300 K and low thermo-optic coefficient of $\sim\ 10^{-5}\mathrm{K}^{-1}$)[123,86] along with negligible birefringence[116,55] make it an ideal material for high-power Raman lasing with greatly reduced thermal lensing effects[123,116].

The availability of CVD-grown, high-quality polished, single-crystal diamond plates has enabled the development of bulk Raman lasers using macroscopic optical cavities across the UV[58], visible[124,125], near-infrared[131,94,149,150,54,182] and even mid-infrared[151] regions of the optical spectrum. Although showing great performance with large output powers (many Watts)[182] and near quantum-limited conversion efficiencies[125,149], most operate in pulsed mode in order to attain the very high pump powers required to exceed the Raman lasing threshold[125,58,54,182]. Demonstration of continuous-wave diamond Raman lasing has been challenging, with very few reports[116,131,94]. Bulk cavity systems also require precise alignment and maintenance of optical components for the laser to function robustly.

Translating Raman laser technology onto an integrated optics platform where the light is confined to nano-waveguides[19,145] and/or high quality-factor (Q) micro-resonators[165,93,171,147] can greatly reduce pump power requirements and enable stable continuous-wave (CW) operation without the

need for any complicated alignment of optical components. Such compact micro-resonator-based Raman lasers, especially if integrated on-chip, might be particularly useful for spectroscopy and sensing applications in harsh environments[135,133] as well as medical device technologies[135,122]. To date, chip-based Raman microlasers have been demonstrated in silicon racetracks[147,146] and photonic crystals[171], and silica microtoroids[93]. Such telecom-laser-pumped devices have shown CW lasing with low threshold powers ($\mu W - mW$), albeit at limited Stokes wavelengths around $\sim 1.6 - 1.7\,\mu m$, and cascaded operation out to $\sim 1.85\,\mu m$[147]. This is due to the relatively low value of the Raman frequency shift in silicon ($\sim$15.6 THz) and silica ($\sim$12.5 THz) compared to diamond ($\sim$40 THz). Moreover, the losses due to two-photon and free carrier absorption in silicon need to be mitigated via carrier extraction that complicates the device layout and fabrication process[145,171,147,146]. Silica-based devices require ultrahigh-Q cavities ($\sim 10^8$) to effectively compensate for the extremely low Raman gain coefficient (>100x smaller than silicon and diamond). Additionally, the broad Raman gain spectrum in silica ($\sim$10 THz) makes single-mode operation non-trivial[165,93]. These devices (microspheres, microtoroids) are also difficult to integrate into a compact, fully-integrated on-chip package, requiring careful alignment of a tapered fiber to evanescently couple light into the resonator[93], although recently-developed spiral waveguides and wedge resonator geometries are amenable to more robust coupling techniques[127]. Finally, both silica and silicon suffer from severe thermal management issues and absorption losses outside of their traditional operating windows, raising a question mark on high-power operation over a wide spectrum.

Diamond can potentially overcome these drawbacks and has recently emerged as a novel nanophotonics material with applications in integrated, on-chip quantum[66,52] and nonlinear optics[62]. Diamond's large bandgap of $\sim$5.5 eV and lack of Reststrahlen-related absorption at low frequencies affords it a wide space for creating high quality factor resonators. Here we demonstrate the first CW, tunable, on-chip Raman laser operating at $\sim 2\,\mu m$ wavelengths using telecom-laser-pumped high-Q, waveguide-integrated diamond racetrack resonators embedded in silica on a silicon chip. Further on

we will refine this technique and extend Raman lasing to near-visible operation.

4.3 Raman laser pumped at telecommunication wavelengths

The Raman process (Fig. 4.2(a)) involves scattering of a high energy pump photon at frequency ω_P, into a low energy Stokes photon at frequency ω_S, via the creation of an optical phonon of frequency Ω_R, such that $\omega_P - \omega_S = \Omega_R$. For diamond, $\Omega_R \sim 40$ THz, corresponding to high-energy optical phonons vibrating along the <111> direction[123,149]. For pump wavelengths in the telecom range ($\lambda_P \sim 1.6\,\mu$m), $\omega_P \sim 190$ THz, resulting in a Stokes wavelength λ_S near $\sim 2\,\mu$m ($\omega_S \sim 150$ THz). Our diamond waveguides, with $\sim 700 \times 800$ nm cross-section embedded in silica, support modes both at the pump and Stokes wavelengths with good spatial overlap (Fig. 4.2(b)). Raman scattering does not require any phase matching as it is an inelastic process. The efficiency of this process, however, is very low in bulk materials and can be significantly increased using optical cavities. In particular, if the cavity is resonant with the Stokes wavelength it can provide optical feedback needed to stimulate the scattering process, which can lead to lasing action. If the cavity is also resonant at the pump wavelength, it can boost up the pump intensity by a factor of the finesse and further enhance the stimulated process. The threshold for Raman lasing in such a doubly-resonant cavity is inversely proportional to the product of the Qs of the pump and Stokes modes[165,93]. The Raman gain spectrum in diamond is extremely narrow with a full-width at half-maximum (FWHM) of ~ 60 GHz[123,86]. To ensure that a resonator mode exists close to the gain maximum, long racetrack micro-resonators (path length $\sim 600\,\mu$m) are designed with free spectral range (FSR ~ 180 GHz) approaching the Raman scattering linewidth (Fig. 4.2(c),(d)).

The basic fabrication process was developed from the previously described approach for integrated diamond devices[66,63,62]. Initially, a $\sim 20\,\mu$m thick, type-IIa CVD, single-crystal diamond (Delaware Diamond Knives) was cleaned in a refluxing acid mixture of nitric, sulphuric, and per-

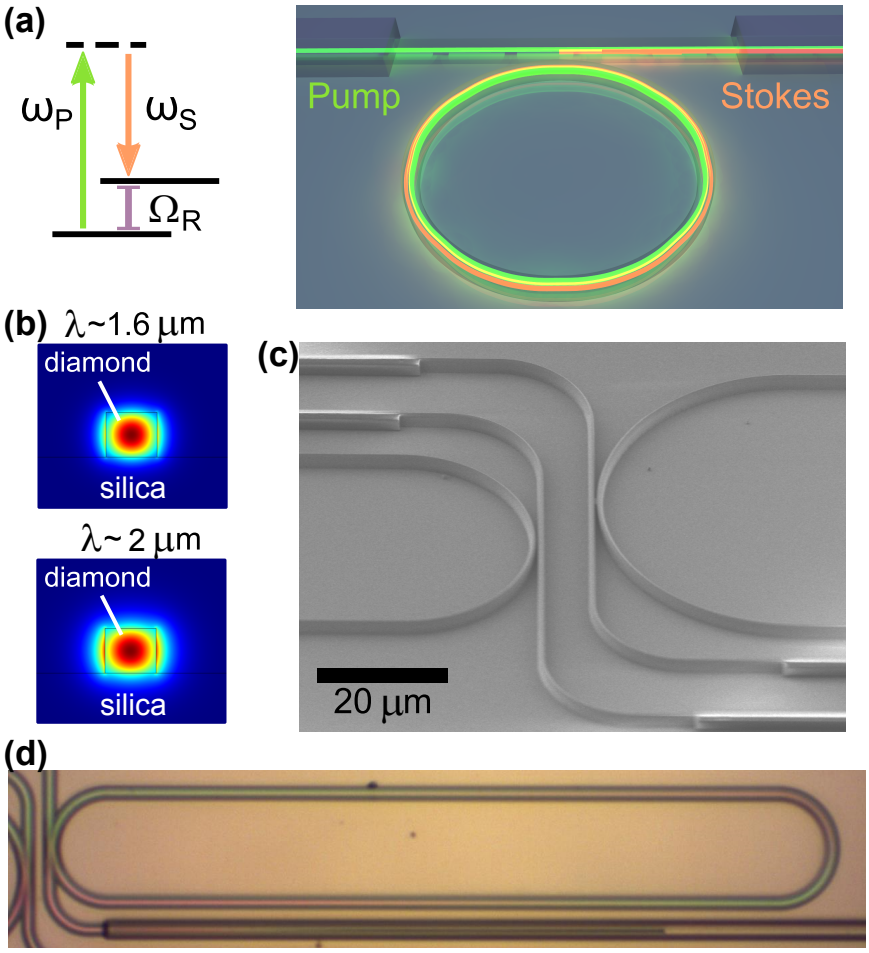

Figure 4.2: Diamond-microresonator based Raman laser design. (a) Energy level diagram of the Raman scattering process (left), wherein a high energy pump photon with frequency ω_P is scattered into a lower frequency Stokes photon, ω_S, and an optical phonon, Ω_R ($\sim$40 THz in diamond). We pump with telecom lasers ($\lambda_P \sim 1.6\,\mu$m) corresponding to $\omega_P \sim 190$ THz, resulting in a Stokes output at $\omega_S \sim 150$ THz i.e. $\lambda_S \sim 2\,\mu$m. A schematic illustrating the device principle (right) shows a pump wave (green) entering a high-Q microcavity, where it enables Stokes lasing (orange) via stimulated Raman scattering. (b) Simulated TE mode profiles of diamond waveguides with width 800 nm and height 700 nm fully embedded in silica, at the pump ($\lambda_P \sim 1.6\,\mu$m, top) and Stokes ($\lambda_S \sim 2\,\mu$m, bottom) wavelengths, showing good overlap. (c) Scanning-electron-microscopy image of the nano-fabricated diamond racetrack resonators on a SiO_2-on-Si substrate before cladding with PECVD silica, showing the bus-waveguide-coupling region (gap $\sim$500 nm) and transition to polymer (SU-8) waveguides for efficient coupling to lensed fibers. (d) Optical micrograph of a diamond racetrack micro-resonator with path length $\sim 600\,\mu$m and bending radius $\sim 20\,\mu$m, after a PECVD silica cladding layer is deposited on top

.

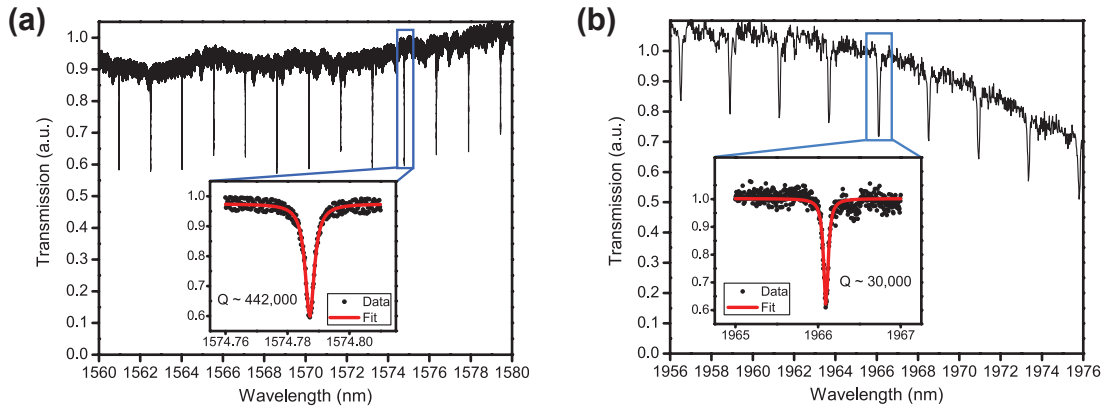

Figure 4.3: High-Q modes at pump and Stokes wavelengths. (a) Transmission spectrum of the diamond racetrack resonator at telecom (pump) wavelengths taken by sweeping a continuous-wave laser reveals high-Q transverse-electric (TE) modes with 30-40% extinction ratio (under-coupled resonances). The path length of the resonator is $\sim$600 μm, corresponding to an FSR of $\sim$1.5 nm ($\sim$180 GHz). Inset: A loaded Q of $\sim$440,000 is inferred from the Lorentzian fit to the mode at $\sim$1574.8 nm. (b) Transmission spectrum of the diamond resonator at the Stokes wavelength range near $\sim$2 μm ($\sim$40 THz red-shifted from the pump) taken using a broadband super-continuum source again reveals high-Q TE modes with 30-40% extinction ratio (under-coupled resonances). Inset: A loaded Q of $\sim$30,000 is inferred from the Lorentzian fit to the mode at $\sim$1966 nm, although this may be limited by the resolution ($\sim$0.056 nm) of our optical spectrum analyzer.

chloric in equal ratios. The device was then thinned to specification (<1 μm) by cycling Ar/Cl_2 and O_2 etching steps in a dedicated Plasmatherm inductively-coupled-plasma reactive-ion-etcher (ICP-RIE) while bonded via Van der Waals forces to a sapphire carrier wafer[66]. The diamond was etched on both sides to remove residual stress/strain from the polishing procedure. Afterwards, the thin diamond film was transferred to a SiO_2/Si substrate with a 2 μm thermal SiO_2 layer. To promote resist adhesion, a thin layer (<5 nm) of SiO_2 was deposited via atomic layer deposition on the diamond film. Afterwards, an etch mask was patterned using Fox 16 electron-beam resist (spin-on-glass, Dow Corning) in an electron-beam lithography tool (Elionix ELS-F125) under multipass exposure. The faces of the supplied thin diamond plates are non-parallel due to the polishing process, with a thickness wedging of $\sim$300 nm per $\sim$1 mm length. The pattern was aligned to the diamond thin film such that the polishing gradient ran parallel to the racetrack devices. This pattern was then etched into the diamond with a final oxygen etch. The Fox 16 resist was left on the diamond. The

completed waveguide had dimensions of ∼800 nm in width and ∼700 nm in height, while the coupling region had a gap of around ∼500 nm. The diamond bus waveguide tapered off over a length of ∼200 μm to an end width of ∼150 nm. Polymer coupling pads to the end of the substrate were written in SU-8 aligned to the adiabatically tapered diamond waveguides[63]. Finally, a layer of ∼3 μm of silica was deposited with plasma-enhanced chemical vapor deposition (PECVD) in order to cap the devices and aid in the polishing of the end facets.

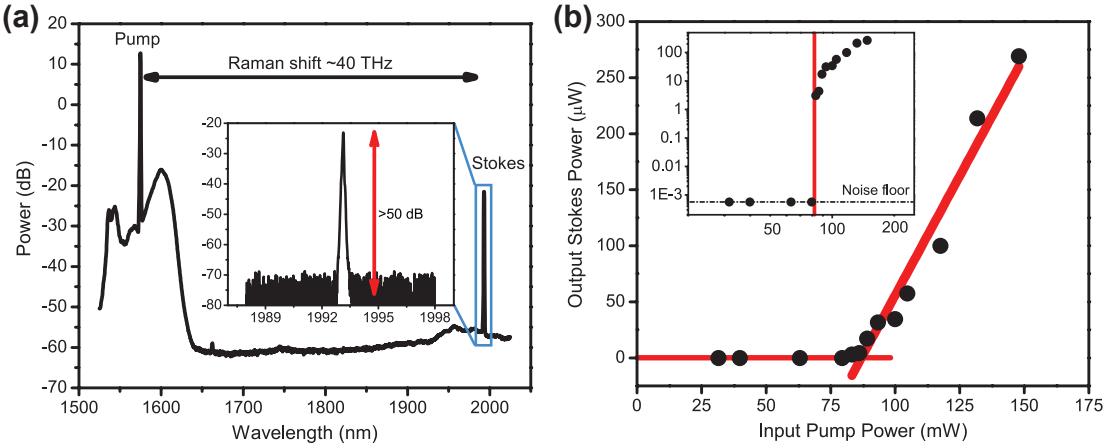

Figure 4.4: Observation of Raman lasing and threshold measurement. (a) Optical spectrum analyzer (OSA) signal when the pump is tuned into a resonance near ∼1575 nm with ∼100 mW power shows the emergence of the Raman line at the Stokes wavelength of ∼1993 nm, ∼40 THz red-shifted from the pump. Inset: A high-resolution scan zooming into the Stokes output reveals >50 dB sideband suppression ratio (>60 dB on-chip after correcting for out-coupling losses). (b) Output Stokes power at ∼1993 nm versus input pump power at ∼1575 nm (both estimated in the bus-waveguide), displaying a clear threshold for Raman lasing at ∼85 mW pump power. The external conversion slope efficiency is ∼ 0.43%, corresponding to an internal quantum efficiency of ∼ 12%. Inset: A log-log plot of the output Stokes power versus input pump power reveals a ∼40 dB jump above the noise floor in the output at threshold.

The on-chip diamond resonators are characterized using a lensed-fibre-based coupling setup[63,62]. Transmission measurements at telecom were taken by sweeping a continuous-wave laser (Santec TSL-510) across the resonances and sending the output to an amplified photodetector (Newport 1811). The insertion loss for the device was measured to be ∼5 dB per facet (∼10 dB total loss from input to output lensed fibre) for telecom wavelengths. In order to measure the resonator modes

around the Stokes wavelengths near 2 μm, a broadband supercontinuum source (NKT Photonics SuperK) was coupled into the device and the output spectrum was recorded on an optical spectrum analyzer (OSA, Yokogawa AQ6375) with a maximum resolution of 0.056 nm. The insertion loss for the device was measured to be ∼9.5 dB per facet (∼19 dB total loss from input to output lensed fibre) at these longer wavelengths, likely because the lensed fibers are designed for telecom wavelengths. Transmission measurements revealed that the diamond resonators support high-Q modes at both the telecom pump (Fig. 4.3(a)) and ∼2 μm Stokes wavelengths (Fig. 4.3(b)). The modes at telecom were found to be under-coupled with ∼30-40% transmission dips on-resonance and high loaded Qs around 400,000 (Fig. 4.3(a)). The higher-wavelength modes around 2 μm also showed under-coupling with ∼30-40% extinction ratios on-resonance and loaded Qs around 30,000, although this may have been limited by the resolution of our optical spectrum analyzer.

For Raman lasing measurements, high pump power was achieved by boosting the input laser power through either a C-band ($\sim 1535 - 1570$ nm) or L-band ($\sim 1570 - 1610$ nm) erbium-doped fiber amplifier (EDFA, Manlight). The pump laser was first set at a slightly blue-detuned position near a resonance before slowly being shifted into it. Power absorbed by the resonator and its host material causes a thermal redshift of the resonance resulting in a characteristic 'shark-fin' shape, allowing the pump to be slowly tuned towards the transmission minimum while stabilizing the power coupled into the resonator[147,62]. While tuning the pump, the Stokes output was monitored on the OSA. When the pump laser is tuned into a resonance with sufficient power, Raman lasing at the Stokes wavelength is observed (Fig. 4.4). After the onset of Raman lasing at a particular detuning, the pump was further fine-tuned to maximize the output.

Fig. 4.4(a) shows the measured optical spectrum with the Stokes line ∼40 THz away from the pump. A zoom into the Stokes line (inset of Fig. 4.4(a)) shows resolution-limited linewidth and >60 dB sideband suppression ratio after correcting for losses, characteristic of low-noise single-mode operation. Fig. 4.4(b) shows the measured output Stokes power as a function of input pump

power, displaying a clear threshold and onset of Raman lasing at ∼85 mW of CW pump power in the coupling waveguide. Stokes powers >250 μW are coupled into the output waveguide, corresponding to an external conversion slope efficiency above threshold of ∼ 0.43%. This is limited by the severely undercoupled nature of the resonances at both the pump and Stokes[165,93], and the internal quantum efficiency itself is estimated to be ∼ 12%. Knowing the Q-factor and mode volume of our device enables us to extract an effective Raman gain value of ∼2.5 cm/GW from the Raman lasing threshold formula[165,93]. This is comparable to, but lower than, previous estimates for diamond at these wavelengths (∼6 cm/GW)[123], suggesting that our Stokes mode is probably not positioned exactly on the Raman gain peak.

We also demonstrate discrete tuning of the Raman laser over a wide bandwidth by tuning the pump laser to separate adjacent resonances. Fig. 4.5(a) shows the result of 14 separate measurements which show a Raman signal spanning from <1950 nm to >2050 nm. The discrete tuning range is >100 nm, or ∼7.5 THz, which corresponds to ∼ 5% of the center frequency and was limited by the operation bandwidth of our pump amplifiers. Within this range, over 40 uniformly spaced longitudinal modes can be individually addressed, each separated by the cavity FSR of ∼180 GHz (Fig. 4.5(b)). Continuous, mode-hop free tuning of the Stokes output over ∼7.5 GHz is also achieved (Fig. 4.5(c)) by tuning the pump within a single thermally red-shifted resonance. As the pump detuning from resonance is decreased, the intra-cavity power increases and the pump and lasing modes are both shifted to the red[147]. Beyond the resonance (sharp edge of the 'shark-fin') the mode is no longer pumped and the cavity begins to cool down, shifting the resonance back to its original position. In order to create a Raman laser which can be tuned over the entire output range continuously, it should suffice to create a resonator with a sufficiently small FSR on the order of the thermal shift (this would require a resonator path length ∼10x our current device which should be possible via a winding spiral resonator design). Then, by tuning into a mode and using its redshift (or, alternatively, an external heater), it should be possible to sweep across one resonance and carry the Stokes

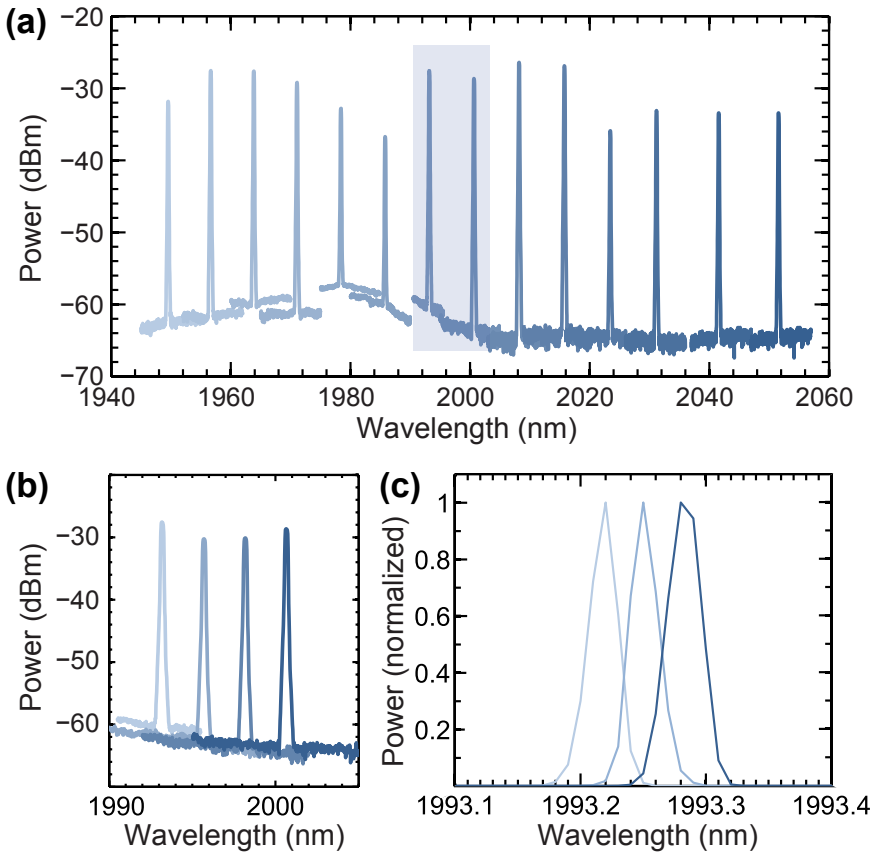

Figure 4.5: Discrete and continuous tuning of Raman laser output wavelength. (a) Discrete tuning of the Stokes wavelength over a range $>100\ \mathrm{nm}$ ($\sim 7.5\ \mathrm{THz}$ or $\sim\ 5\%$ of the center frequency). The pump is tuned to 14 separate resonances, each spaced by $3\times$FSR ($\sim 550\ \mathrm{GHz}$), and the Raman line is recorded with an optical spectrum analyzer (OSA) at each pump wavelength. (b) Stokes output of adjacent modes. Here the pump is tuned to neighboring resonances (one FSR apart) within the highlighted region of Fig. 4.5(a). The output modes are also spaced by an FSR or $\sim 180\ \mathrm{GHz}$. Thus, over 40 individual longitudinal modes can be accessed over the entire demonstrated tuning range. (c) Mode-hop-free tuning of the Stokes wavelength over $\sim 0.1\ \mathrm{nm}$ or $\sim 7.5\ \mathrm{GHz}$. The pump frequency is tuned within a thermally red-shifted resonance ('shark-fin' shape), thus tuning the output Stokes wavelength in a continuous fashion. The output power is normalized to the peak emission at each pump wavelength. The linewidth of the Stokes mode is limited by the minimum resolution of our OSA ($\sim 0.05\ \mathrm{nm}$).

from one longitudinal mode of the resonator to the next continuously[147].

In conclusion, we have demonstrated a CW, low-threshold, tunable, on-chip Raman laser operating at ∼2 μm wavelengths based on waveguide-integrated diamond racetrack microresonators. Our results firstly introduce diamond as a viable material for compact, on-chip Raman lasers over a wide spectrum, and secondly present a new laser source in the technologically-exciting 2 μm region[162]. The threshold power in our current device, although the lowest demonstrated in any kind of diamond Raman laser by a few orders of magnitude, is still limited by the severe under-coupling of the bus waveguide to the resonator and could be further reduced by moving to near critically-coupled modes for the pump[165,93]. This can be easily achieved, for example, by slightly reducing the coupling-gap between the bus-waveguide and resonator. The external conversion efficiency can also be drastically increased by having over-coupled resonances for the Stokes in addition to critical-coupling for the pump[165,93], and this should naturally happen in the current design if the intrinsic Qs of the pump and Stokes mode are of the same order. Longer coupling sections and other coupling designs can also be investigated[147]. Further improvement can be made by having higher intrinsic Q[62] and/or smaller FSR (to ensure maximum Raman gain) i.e. longer path-length resonators[147]. Another limiting factor comes from the orientation of the diamond itself. Our devices are fabricated in [100]-oriented diamond, and the pump and Stokes mode are both TE polarized, where Raman gain is sub-optimal and there is no polarization preference for the Stokes[123,149]. By ensuring that the light polarization is parallel to <111>, for example using angle-etched resonators[28,25] in thick [111]-diamond plates, the efficiency of the Raman process can be enhanced[123,149]. Further, by moving to such an all-diamond structure, the resonator should be able to support more circulating power and reach higher output powers while also offering a route toward longer-wavelength/cascaded Raman lasers, where the absorption of silica would limit performance otherwise. Nonetheless, the current platform already offers a large amount of flexibility, with the option to fabricate devices at visible wavelengths, where the Raman gain is $\sim 20\times$ higher[123]. Operation in the visible could also enable

integration of classical nonlinear optics technologies (Raman lasing, Kerr frequency combs) with the quantum optics of color centers [66,52,62]

4.4 Raman lasing at visible wavelengths

Integrated photonic devices operating at visible wavelengths must overcome additional challenges relative to their telecommunications-band counterparts. Apart from relatively immature instrumentation, there is a dearth of qualified low-loss material platforms. This requirement is even more stringent in nonlinear photonic systems, where high visible pump powers and non-negligible linear and nonlinear absorption can lead to device failure. To enable pumping of the diamond microresonator system at sufficient powers above threshold, we develop a directly-writable doped-glass waveguide and integrate it with our diamond waveguides. This technique creates an end-fire interface from the edge of the sample to the diamond, whose physical size is limited to <1 mm. Per these developments, we demonstrate the first integrated Raman laser operating at near-visible, studying in particular polarization conversion between pump and Stokes and low-threshold lasing.

Device fabrication broadly follows the same process reported previously [66,101] and discussed in detail in Chapter 3.3. A 1x1 mm, ~30 μm thick electronic-grade single-crystal diamond (Element Six) was cleaned in a refluxing acid mixture. After rinsing in water, the diamond was then placed from methanol directly on a sapphire carrier wafer. This promotes a loose adhesion to the carrier. The diamond is etched on one side, then flipped, cleaned, and etched (Ar/Cl_2 cycled with O_2). The thinned diamond is then transferred to fused silica by first de-bonding from the carrier wafer with a drop of hydrofluoric acid, which is then gently washed away and diluted. Immediately before re-bonding to the silica substrate, the silica surface is activated with an O_2 plasma (300 mT, 100 W, 1 minute). Once the diamond is bonded, a monolayer of Al_2O_3 is deposited via atomic layer deposition (ALD), promoting adhesion of the FOx-16 electron-beam resist (spin-on glass, Dow Corning).

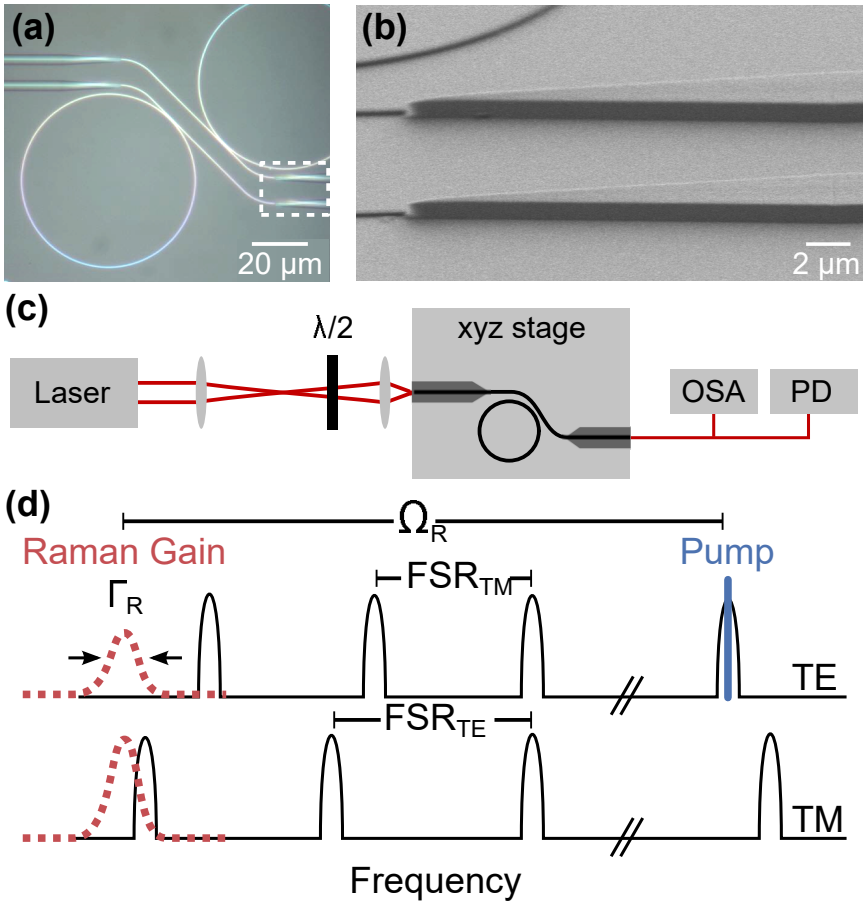

Figure 4.6: Microresonator design and experimental concept. (a) Micrograph of diamond ring resonators, with coupling region highlighted. (b) Scanning electron micrograph of doped-glass to diamond coupling region. (c) Experimental setup. The light from a CW Ti:Sapphire laser is focused onto a waveguide facet after passing through a half-wave plate. A lensed fiber collects light and directs it to an optical spectrum analyzer and a photodetector. (d) Experimental concept. A CW pump is tuned to resonance, creating a Raman gain region at $\Omega_R \sim 39.99$ THz Stokes shift for both TE and TM polarizations. The strength of the Raman gain for these polarizations is dependent on the specific microresonator parameters. The difference from the Stokes Raman gain to the nearest resonance mode is termed the Stokes-mode mismatch.

The resist is written under multi-pass exposure (Elionix F-125) with waveguides aligned parallel to the diamond thickness gradient[101]. The finished resonators (Fig. 4.6(a)) are 60 μm in diameter and have cross-sectional dimensions of 300 nm in width and ∼300 nm in height. The diamond waveguides are adiabatically tapered down in width across a length of 200 μm to enable high coupling efficiency.

The diamond sample is prepared for subsequent processing by depositing ∼2 nm of ALD oxide[71]. This adheres the diamond to the substrate and prevents de-bonding. To prevent out-gassing when under high optical power, the sample is annealed at 460 C for 1 hour in O_2. This drives impurities out of the ALD film. Directly-written doped-glass waveguides to be used as spot-size converters (Fig. 4.6(b)) are then defined with an electron beam exposure and developed.

Afterwards, a layer of photoresist is spun and the sample is diced, cleaved, and cleaned. Finally, the sample is annealed at 460 C for 3 hours in an oxygen atmosphere to drive out residual impurities in the doped-glass waveguides and ALD film, as well as to etch away any final graphitized carbon and terminate the diamond surface in oxygen bonds[130]. Apart from the thin ALD film, the structures are left uncapped.

Optical measurements of the devices are performed in an end-fire testing setup, schematically shown in Fig. 4.6(c). An aspheric singlet lens focuses a CW Ti:Sapphire laser (M2, Solstis) onto the facet of a waveguide. Free-space coupling has several advantages over a lensed optical fiber, as was used in previous experiments[101], by reducing insertion loss and maintaining stable polarization even at high powers.

After passing through the device, the light is collected on the other end with a lensed fiber (OZ Optics). A typical device showed losses of ∼5 dB/facet. An optical fiber splitter is then used to send the output to a photodetector for resonance measurements, or an optical spectrum analyzer (OSA, Yokogawa AQ6370) for Raman measurements.

Two separate devices are reported. The first device under test is characterized in Fig. 4.7. Its

waveguide cross-sections are shown in Fig. 4.7(a). Both TE and TM spectra are shown in Fig. 4.7(b), where the insets show the TE pump resonance used and its corresponding Stokes resonance. The pump and Stokes quality factors are 301,000 and 85,000, respectively. No higher-order Stokes or anti-Stokes processes were observed. With an observed Raman threshold of 20 mW, we calculate[165] an effective Raman gain value of 3.2 $\mathrm{cm\,GW^{-1}}$. This is lower than previous values for diamond at these wavelengths $>$ 10 $\mathrm{cm\,GW^{-1}}$[123], and can be attributed to imperfect confinement (74% of the optical power is in the diamond core) and Raman gain dependence on the orientation of the crystal axes with respect to the propagation/polarization direction[123,53]. The near-threshold data implies an external conversion efficiency of 1.7%, corresponding to an internal efficiency of 85%, close to the theoretical maximum of 90%.

Fig. 4.8(a) shows the Stokes-mode mismatch and recorded Stokes power at the various pump wavelengths. Because the TE-TE transition was well-matched throughout the device, Stokes lasing was observed with an output over a bandwidth of >150 nm, or 60 THz, corresponding to 17.5% of the center frequency, from 800 $-$ 950 nm. Threshold measurements in 4.8(b) were taken at the pump resonance from Fig. 4.7.

4.5 Polarization conversion with Raman lasing

Fig. 4.9 shows the Stokes Raman gain-resonator mode mismatch for both TE (a) and TM (b) pump wavelengths to TE Stokes (blue circles) and TM Stokes (red triangles) for the second device under test. The locations of the resonances are recorded at low optical power, so thermal shifts are not taken into account. The bottom ribbons show where Raman lasing was observed at high pump powers ($\sim$200 mW in the waveguide), with the pump wavelengths corresponding to the well-matched areas ("doubly-resonant" operation). The correspondence between the Stokes-mode mismatch and the observed Stokes output reflects in the wider output bandwidth for the TE pump,

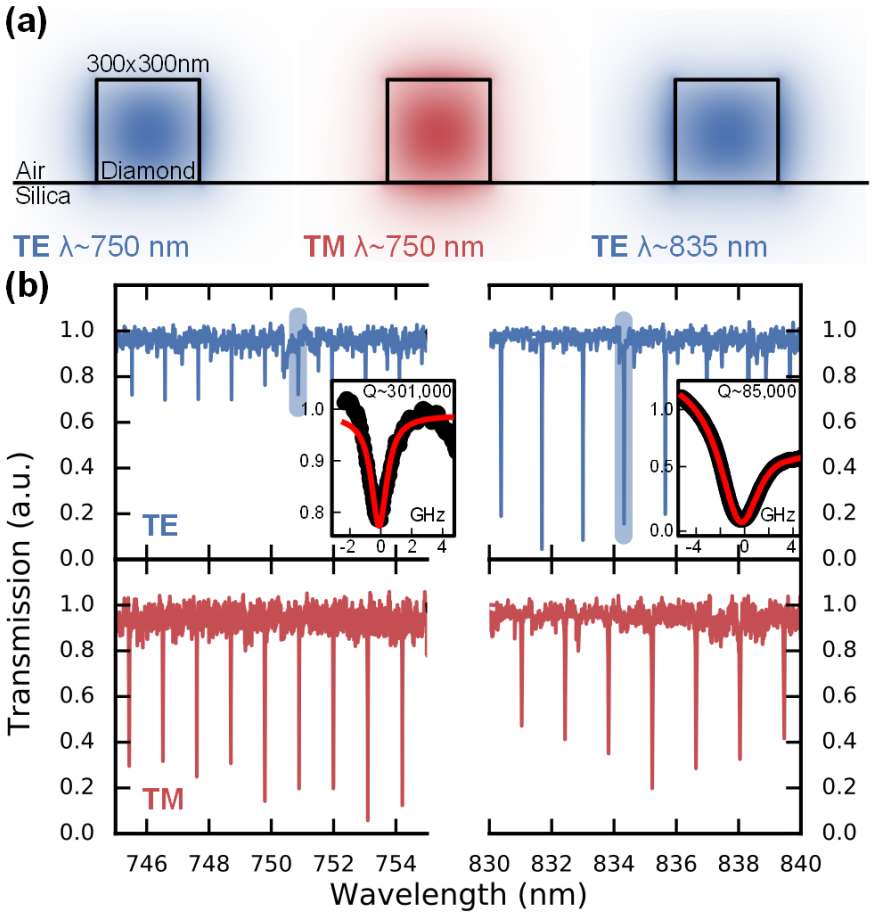

Figure 4.7: Low-threshold device measurements at low power. (a) Cross-sections of simulated mode intensities at TE pump (left, blue), TM pump (center, red) and TE Stokes (right, blue). (b) Transmission data for studied resonator at TE (top, blue) and TM (bottom, red) probe polarizations. Insets show zoomed-in highlighted regions of the TE spectra. The studied pump resonance shows a quality factor Q ∼301,000 at $750.88\ \mathrm{nm}$, corresponding to a Stokes resonance of $Q \sim$ 85,000 at $834.34\ \mathrm{nm}$

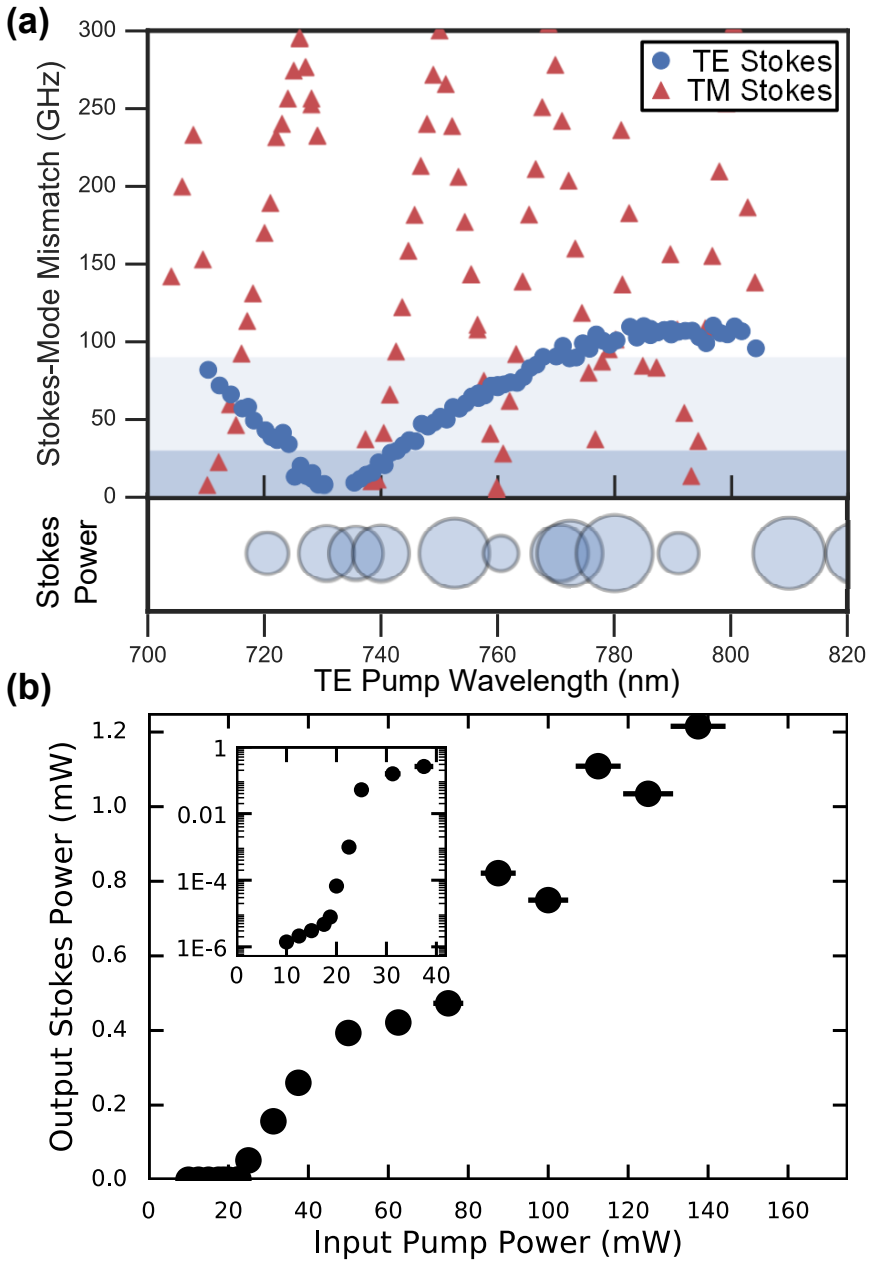

Figure 4.8: Low-threshold device Raman measurements. (a) Stokes-mode mismatch plot (top) shows the mismatch between the Raman gain from a TE pump resonance and the nearest TE (blue, circles) and TM (red, triangles) resonance. The blue and light blue bands show the mode mismatch lying within one half-maximum and three half-maximums of the Raman gain peak, respectively. The ribbon (bottom) qualitatively shows (by circle size) recorded Stokes power from Raman lasing at the respective pump wavelengths. The TE-TE transition is well-matched across the measurement, and Stokes output is observed over $60\ \mathrm{THz}$ of tuning (bottom). (b) Threshold measurements for TE pump at the resonance from Fig. 4.7, or $750.88\ \mathrm{nm}$. The inset shows a semi-log plot of the near-threshold data.

and the two distinct lasing regions for the TM pump. Only TE-TE and TM-TE Raman lasing were observed. A TM Stokes output was not seen, likely due to the lower quality factors for TM resonances and forbidden TM-TM transition for [001]-oriented diamond[123,53].

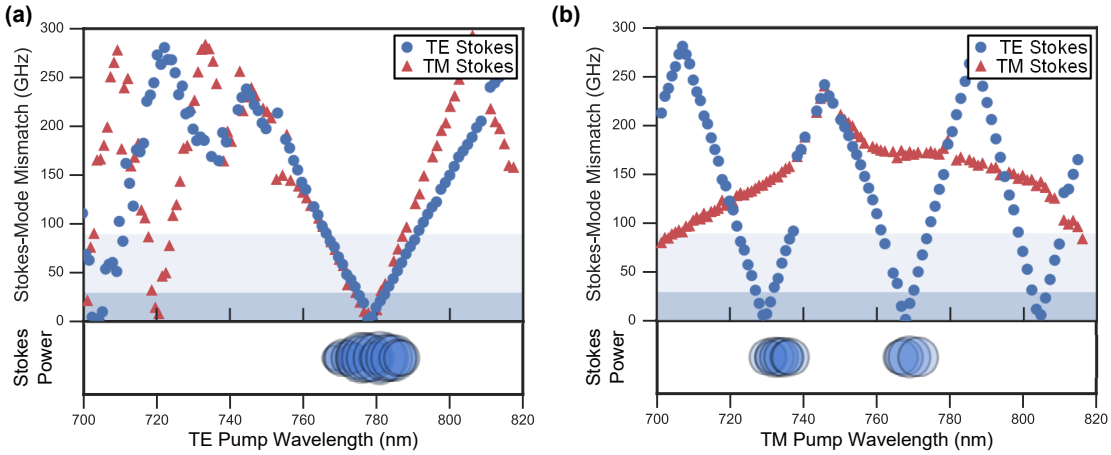

Figure 4.9: Raman lasing and polarization conversion dependence on Stokes-mode mismatch, same nominal dimensions as Fig. 4.7. (a) The mismatch between the Raman gain created by a TE pump resonance and the nearest TE (blue circles) and TM (red triangles) mode is plotted (top). Though both TE and TM resonances are well-matched, the output is likely TE, based on the higher quality factor of these resonances. (b) A similar plot as (a), but for TM pump resonance. Two regions of Stokes output are recorded at roughly $\sim$735 nm and $\sim$770 nm. The TM pump is converted to a TE Stokes, and the neighborhood of modes for which Raman lasing is observed is narrower, as expected from the Stokes-mode mismatch plot.

In conclusion, we have demonstrated an integrated CW Raman laser operating in the near-visible region. Owing to directly-written doped-glass waveguides, we were able to show lasing at high pump powers across 60 THz of tuning. The mismatch between the Raman gain peak and the nearest resonator mode was shown to be an important parameter, and polarization conversion from TM-TE was also observed. We studied the threshold behavior for TE-TE Raman lasing for a given Stokes output resonance, finding low threshold power $\sim$20 mW.

These results strongly suggest the difficulty of suppressing Raman processes over a broad bandwidth in large-FSR resonators, as is required for generating visible Kerr frequency combs in diamond[128]. Nonetheless, the improvements in diamond quality factor and coupling-waveguide power

handling are broadly applicable to integrated optics in the visible, including quantum, single-emitter based devices. The silicon-vacancy center, for instance, has an optical transition very close to the wavelengths studied here, and co-integration can be used for hybrid nonlinear-quantum optics devices. By translating these results to an angle-etched platform [25,100], diamond's full potential can be realized.

4.6 Future platforms for Raman lasing

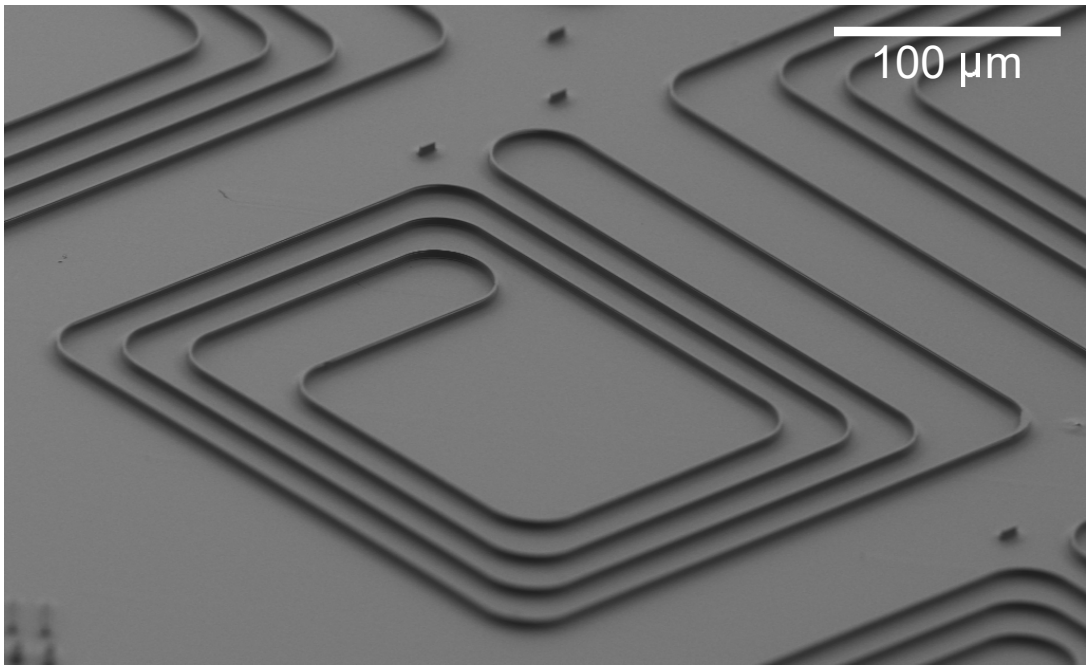

Figure 4.10: Candidate resonator for next-generation Raman lasers. The long path-length of the device leads to a very small FSR (<60 GHz), ensuring that every mode lies within one FWHM of the Raman gain peak. In addition, this angle-etched device is made of monolithic diamond, which has advantages for high-power handling.

To date, Raman lasing in diamond been demonstrated exclusively in the thin-film platform. However, due to persistent issues with creating long path-length resonators in this platform, it is very difficult to obtain resonators with small FSR. Although a whispering gallery-mode resonator in

diamond has been demonstrated[79] with azimuthal mode FSR of 20 GHz, no lasing was reported.

Angled-etching, as discussed previously, is a robust platform for fabricating scalable diamond photonics. Because all processing steps occur from a planar substrate, there are no wedging issues as in the thin-film platform. To this end, a diamond microresonator defined via angle-etching can be used as a small-FSR resonator for angle-etching. Furthermore, because the crystal axis for the diamond is well-known in bulk substrates, the waveguide orientation can be easily aligned with a preferred crystal axis.

Fig. 4.10 shows such a fabricated microresonator. In addition to its attractive FSR and optimized Raman gain, this platform also has advantages for high-power handling, being composed of monolithic diamond. The difficulty in fabrication stems mostly from irregularities in angled-etching, which is remedied in current process flows using ion beam etching. The last difficulty is in coupling sufficient light into the resonator. Namely, the coupling coefficient required to couple into a microresonator is linearly proportional to the length of the resonator, all else being equal. A technique such as loaded fiber coupling[100], discussed in section 2.4.3, can be used to enhance the coupling into the resonator and provide the requisite input powers for Raman lasing.

5

Kerr nonlinearities in diamond

Before beginning our discussion of Kerr nonlinearities in diamond, it is useful to take a step back and do a proper accounting of all nonlinear optical physics in diamond, expanding outside the scope we restricted ourselves to in the last chapter. Although the high-level description of Raman lasing was sufficient to phenomologically explain experimental results, by diving deeper we can begin to understand how to design optical structures with a mature understanding of diamond's material properties. Broadly speaking, this treatment also applies to other crystalline materials such as sili-

con (which possesses the same crystal symmetry as diamond) and, with slight modifications, other materials like barium fluoride or even lithium niobate.

To begin, we start with the nonlinear Schrodinger equation as it appears in most textbooks in nonlinear optics:

$$\frac{\partial A}{\partial z} = -i\frac{\beta_2}{2}\frac{\partial^2 A}{\partial t^2} + i\gamma|A|^2 A \tag{5.1}$$

This bare-bones partial differential equation is sufficient to describe a wealth of nonlinear phenomena, and even permits analytical solutions in certain cases, namely solitons[195]. Here, β_2 is the second-order dispersion and γ is the nonlinear coefficient. However, as hinted by the discussion in the previous chapter, it does not include any discussion of Raman scattering, nor does it including higher-order dispersive terms or what is known as the shock term. Under a full accounting of all of these terms and incorporating the tensorial nature of third-order nonlinear effect, we arrive at the following generalized nonlinear Schroedinger equation[113,44]:

$$\frac{\partial A_i}{\partial z} = \sum_{k=2} \frac{i^{k+1}\beta_k}{k!}\frac{\partial^k A_i}{\partial t^k} + i\left(1 + \frac{i}{\omega_0}\frac{\partial}{\partial t}\right) P_i^{\mathrm{NL}} \tag{5.2}$$

where $\beta_k = \frac{\partial^k \beta}{\partial \omega^k}$. The shock term is given by $i/\omega_0 \partial/\partial t$. The nonlinear polarization $P_i^{\mathrm{NL}}(z, t)$ contains the third-order response from both the electronic and Raman contributions. More concretely,

$$P_i^{\mathrm{NL}}(z, t) = A_j(z, t)\left[\int_{-\infty}^{\infty} R_{ijkl}^{(3)}(t-\tau)A_k^*(z, \tau)A_l(z, t-\tau)\,\mathrm{d}\tau + i\Gamma_R(z, t)\right] \tag{5.3}$$

Here we introduce $\Gamma_R(z, t)$, a multiplicative stochastic variable describing the noise from the Raman process. The variable has frequency-domain correlations described by $\langle \Gamma_R(\Omega, z)\Gamma^*(\Omega', z')\rangle = (2f_R\hbar\omega_0)|\mathrm{Im}[\tilde{h}_R(\Omega)]|[n_{\mathrm{th}}(|\Omega|) + U(-\Omega)]\delta(z - z')\delta(\Omega - \Omega')$, where Ω describes the phonon

modes and n_{th} gives the Bose distribution function, $\tilde{h}_R$ the Raman response in frequency space, and $U(\cdot)$ the Heaviside step function. The noise distribution is maximal at the Raman shift frequencies, with more noise in the Stokes process. The third order susceptibility is a given by the tensor

$$R^{(3)}_{ijkl}(\tau) = \delta(\tau)\chi^{e}_{ijkl} + h_R(\tau)\chi^{R}_{ijkl} \tag{5.4}$$

Where χ^{e}_{ijkl} is the electronic nonlinear susceptibility and χ^{R}_{ijkl} is the Raman nonlinear susceptibility. This equation is telling us that we treat the electronic susceptibility as instantaneous and the Raman susceptibility as having a delayed response $h_R(\tau)$. Note that in principle $R^{(3)}_{ijkl}(t)$ can also vary as a function of z. In the special case of a cubic lattice like diamond, we can derive a concrete form for both the electronic and Raman contributions, accounting for the electronic nonlinear anisotropy by defining $\rho \equiv 3\chi^{e}_{1122}/\chi^{e}_{1111}$:

$$R^{(3)}_{ijkl}(\tau) = \gamma^{e}(\omega_0)\delta(\tau)\left[\frac{\rho}{3}(\delta_{ij}\delta_{kl} + \delta_{ik}\delta_{jl} + \delta_{il}\delta_{jk}) + (1-\rho)\delta_{ijkl}\right] + \gamma^{R}(\omega_0)h_R(\tau)(\delta_{ik}\delta_{jl} + \delta_{il}\delta_{jk} - 2\delta_{ijkl}) \tag{5.5}$$

and the Raman response $h_R(t)$ is[113]

$$h_R(t) = \Omega_R^2 \tau_1 e^{-t/\tau_2}\sin(t/\tau_1) \tag{5.6}$$

where $\tau_2 = 1/\Gamma_R$ is the inverse of the gain bandwidth and $\tau_1 = 1/\sqrt{\Omega_R^2 - \Gamma_R^2} \approx 1/\Omega^R$ is approximately the inverse of the Raman frequency shift Ω_R.

5.1 Four-wave mixing

By ignoring the Raman term and considering only the interaction between three different modes (A_p, A_s, and A_i, where the pump A_p dominates), we can arrive at the following system of equations

which describes the nonlinear-mediated frequency mixing and amplification (here β_m is the propagation constant for mode m)[113].

$$\frac{\partial A_p}{\partial z} = (i\beta_p - \alpha_p/2)A_p + i\gamma^e_{pppp}|A_p|^2 A_p \tag{5.7}$$

$$\frac{\partial A_s}{\partial z} = (i\beta_s - \alpha_s/2)A_s + 2i\gamma^e_{spps}|A_p|^2 A_s + i\gamma^e_{spip}A_p^2 A_i^* \tag{5.8}$$

$$\frac{\partial A_i}{\partial z} = (i\beta_i - \alpha_i/2)A_i + 2i\gamma^e_{ippi}|A_p|^2 A_i + i\gamma^e_{ipsp}A_p^2 A_s^* \tag{5.9}$$

Relevant to our discussion, by making our system a resonator we can enhance the conversion efficiency η of our frequency generation process into the idler by[1] $\eta \sim \gamma|A_p|^4|A_s|^2$.

To demonstrate four-wave mixing in a device, we employ the same fabrication process discussed in the previous chapter for visible integrated Raman lasers. In addition to a strong CW pump from the Ti:sapphire laser, we add an additional signal beam into an optical fiber which carries both the pump and the signal. The setup is shown schematically in Fig. 5.1(a).

Having measured the dispersion of our device and knowing it is both anomalous and relatively low, we pump the resonator at different powers and observe the generated idler. The modes chosen are nearly adjacent, located at wavelengths of 732, 733, and 734 nm, with the pump parked centrally. These modes showed excellent quality factors, with peak Q of 300,000. A close up spectrum of the modes is shown in Fig. 5.1(b).

Fig. 5.1 shows the measured output at the OSA of the generated light with two different pump wavelengths. Although the process did not reach large enough powers to show cascading of four-wave mixing, the expected output behaved according to theoretical expectations.

To our knowledge, this is the shortest wavelength four-wave mixing ever demonstrated in an integrated optics platform. Generation of new idler colors could occur under stronger CW pumping, or, alternatively, with a higher-quality factor device. Using the attained data, we can estimate the

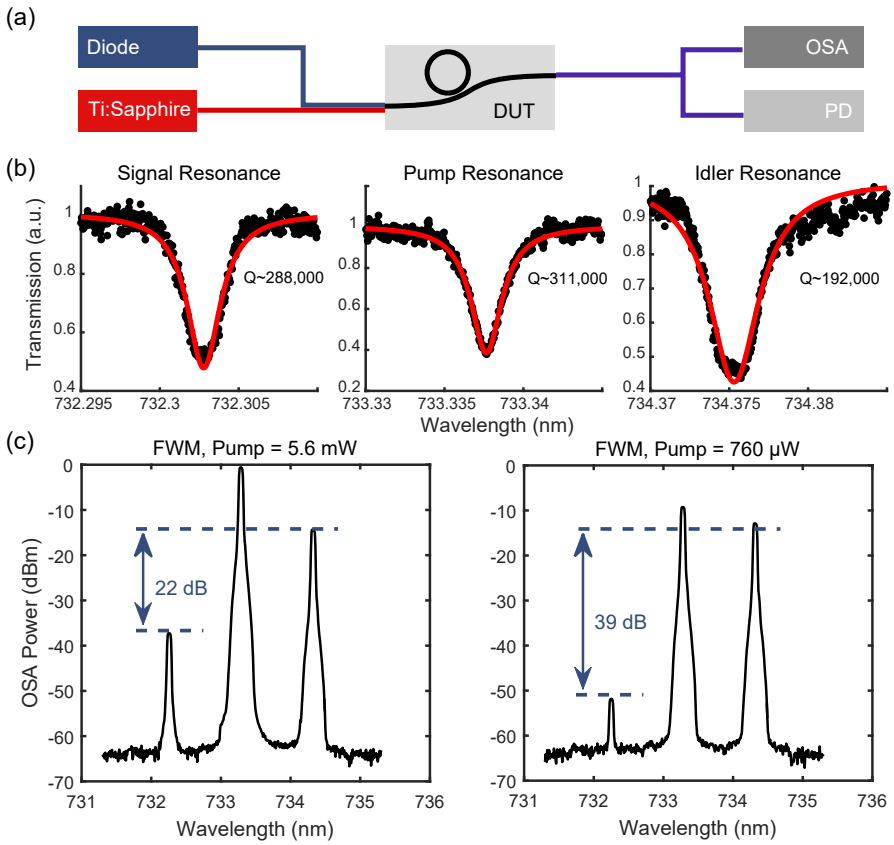

Figure 5.1: Four-wave mixing experiment. (a) Schematic of experimental setup. The light from the pump and probe laser is combined into the same optical fiber and focused on the chip end facet with a lensed fiber. The light is collected and measured via OSA, while a photodetector monitors the power. (b) Measured resonances of the signal, pump and idler. The highest-Q resonance (>300,000) was chosen as the pump. (c) Measured output of four-wave mixing experiment, showing good agreement with theoretical expectations. A pump of 0.76 (5.6) mW and and idler of 0.3 mW yields a signal of 0.04 (1.2) μW, as shown in the left (right) plot.

threshold for optical parametric oscillation - where a single, intense pump beam can generate new idler colors on its own - as around 60 mW in the waveguide, or roughly ten times the power that was able to be coupled in.

5.2 Supercontinuum generation

For short, high-energy pulses traveling in an anomalously dispersive medium, the pulse can interact with itself through the GNLSE and generate new colors and undergo incredibly complex temporal dynamics. This process is known as supercontinuum generation, and it sees widespread use in various fields of optics, especially as a bright, broadband, and coherent source of light. Although supercontinuum generation has been known since the 1970s[3,109], it began to see widespread use upon the invention of the microstructured fiber[140,49]. This fiber was able to create an anomalously dispersive region in the pump wavelength of interest (typically 800nm, as pumped by a Ti:Sapphire, but available for other wavelengths as well).

Within the laboratory, waveguides or fibers which are able to provide broad supercontinua see enormous use as components of larger systems. They aid in exoplanet searches[34,57], optical coherence tomography[78], among other technologies. Recently, waveguides on an integrated photonics platform have been able to show results where traditional glass fibers had not been able to reach, including generation in the mid-IR[157,69,70] and even up to the UV[192]. Diamond's promise for supercontinuum generation encompasses these regions as well, due to its broadband transparency and superlative material properties.

Because supercontinuum-generating devices consist of a simple waveguide with an in- and out-coupling facet, we choose to fabricate our design in the angle-etched platform. This approach offers a number of advantages over the DOI approach in this case. First, because the diamond waveguide is suspended, it tends to have a zero-dispersion wavelength that is bluer than that for DOI devices.

Second, the angle-etched platform allows us a way to write a large device (∼4mm) on the diamond without worrying about film wedge angle or fabrication. With a judicious choice of facet design, we can fabricate a waveguide which approaches the edge of the diamond and couples light in, transporting it to the other side.

A schematic of the design is show in Fig. 5.2(a). The waveguide is supported throughout with a thin support structure of diamond, no more than 100 nm wide. The loss expected from the support structure is shown in Fig. 5.2(c), where both the height off the substrate and the support width were taken into account. For the design chosen (support width < 100 nm, support height 2.4 μm), the simulated loss was less than 0.1 cm^{-1} for the wavelength and polarization of interest (800 nm, TE). The simulated dispersion of the device is shown in Fig. 5.2(d), demonstrating anomalous dispersion for a few different waveguide parameters.

The devices were fabricated via RIBUE in bulk diamond with beveled edges to facilitate coupling. Measurements of the sample showed that these first-generation devices had relatively high insertion loss (around 10 dB/facet) as well as relatively high loss in general (around 20 dB/cm). Fig. 5.3(a) shows SEMs of the fabricated device. The waveguide was fabricated along the 100 axis of diamond, which ensured that there was no TE-TE or TM-TM Raman transition, and pumped with TE polarization. Although we expect a TE-TM Raman transition as per previous work[101], because the TM mode is significantly more lossy, there would not be significant build-up of TM energy to encourage stimulated Raman scattering.

Simulated results are shown in Fig. 5.3(a), and initial experimental results are shown in Fig. 5.3(b). Although the broadening is limited to >100 nm, this early demonstration of an entirely new integrated photonics design shows that, if steady progress can be made, the simulated supercontinuum results should one day be attainable.

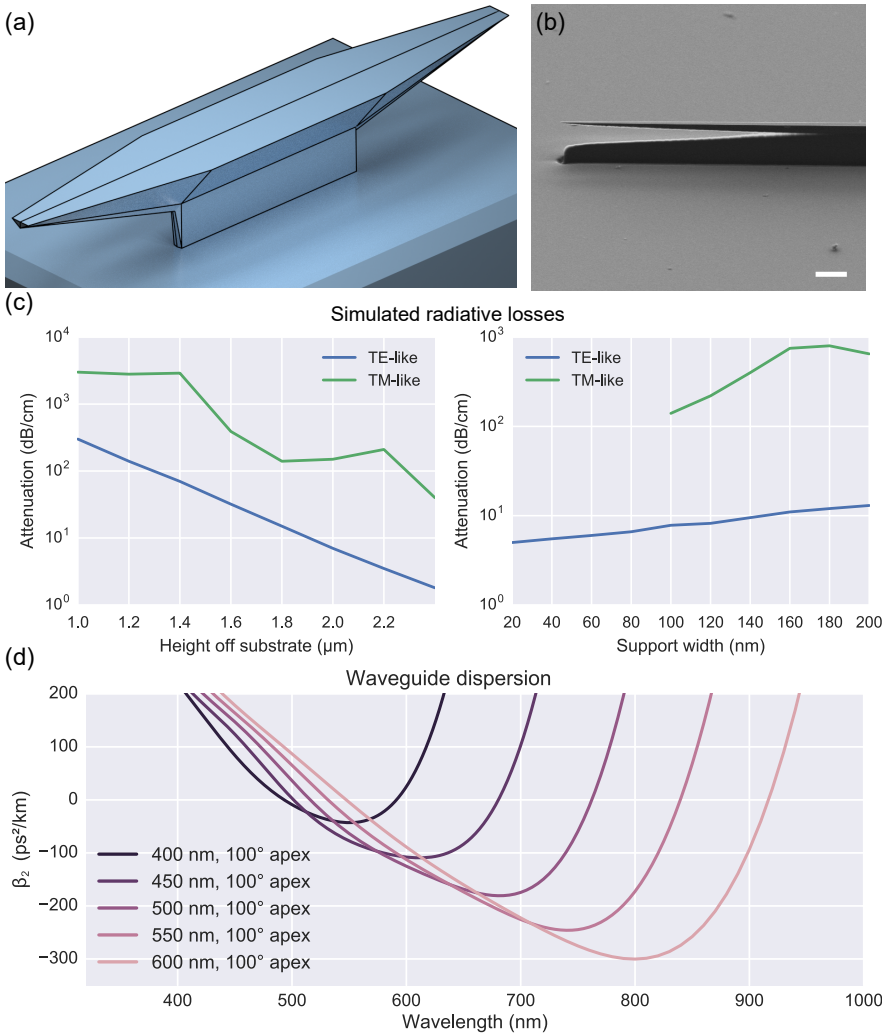

Figure 5.2: Supercontinuum diamond waveguide design and simulation. (a) Device architecture. The waveguide is tapered on either end to support coupling into the waveguide mode, and the bulk of the waveguide is supported by a thin (<100 nm) structure above the substrate. (b) Close-up SEM of fabricated tapered region. The scale bar is two microns. (c) Simulated losses from height of waveguide off substrate, holding width constant, (left) and width of supporting structure, holding height constant at 2 μm (right). Simulations performed at 800 nm pump wavelength, with devices of cross-section 500 nm and apex angle 100 degrees. (d) Anomalous dispersion for waveguides in triangular geometric cross-section of varying widths. The apex angle is kept constant at 100 degrees.

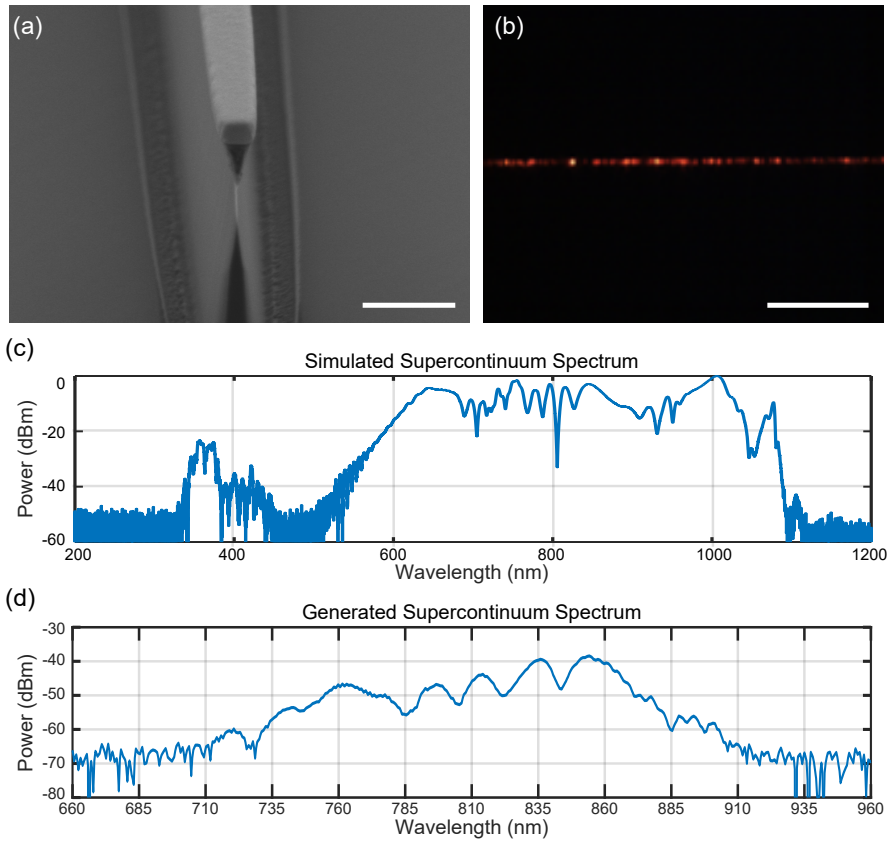

Figure 5.3: Supercontinuum diamond waveguide experimental results. (a) Device SEM of head-on waveguide, showing angular etch. The mask is left on. Scale bar $1\mu m$ (b) Device micrograph with in-coupled visible light. The light is clearly waveguided. Scale bar is 100 μm. (c) Simulated supercontinuum spectrum at 550 nominal width, 100 degree apex angle, 80 MHz repetition rate and 6 mW average power. (c) Experimental supercontinuum spectrum for waveguides at 550 nm nominal width, 100 fs pulses at 80 MHz repetition rate and 50 mW average power. The large insertion losses for this device (>10 dB) required higher pump powers. The fabricated device width may be smaller than the nominal width, affecting supercontinuum generation.

5.3 Frequency comb generation

Integrated nonlinear optics has been seeing an incredible driving force from the prospect of coherent, low noise, broadband spectrum generation on a chip. Similarly to the case in supercontinuum generation, here an optical field builds up and can turn into solitons and broadband light. The main point of contrast is that in this case, the light circles around in a resonator, and may be CW to start with (instead of a short pulse, as in supercontinuum generation).

The first experiments of frequency comb generation were done in microtoroidal glass resonators [46], but recent progress has seen demonstrations in a number of material platforms, including silicon nitride [92,21], silicon [194,193], aluminum nitride [85], silica microdisks [186,191], fluorites [80], and even diamond at telecom wavelength [62]. A limitation of a number of these platforms is the inability to attain a usable supercontinuum in visible wavelengths, though in some cases harmonic generation is able to create visible light from the circulating telecom pulse through $\chi^{(2)}$ and $\chi^{(3)}$ processes.

To appropriately model the a resonator within the context of the GNLSE (Eq. 5.2, all one has to do is apply the Ikeda map every round trip [60]:

$$A_i^{m+1}(t,0) = \sqrt{\theta_i}E_{\text{in}} + \sqrt{1-\theta_i}e^{-i\delta_{i0}}A_i^m(t,L) \tag{5.10}$$

In other words, one propagates the field A_i^m from 0 to the resonator path length L via Eq. 5.2, then adds Eq. 5.10 to the solution and repeats. The term θ_i appearing in the Ikeda map denotes the coupling parameter between the resonator mode and the pumping mode with field E_{in}. δ_{0i} describes the detuning of the pump field to the cavity resonance, related to its phase shift. This completely general solution uncovers a wealth of physics, even in simplifying approximations of the GNLSE [60].

However, for high-Q resonators the change in the field after each successive trip can be very small, so it may take millions of time steps to reach a steady state or solution. To simplify the Ikeda map,

we can perform a time average [41] and arrive at the so-called Lugiato-Lefever equation [117]

$$\frac{1}{\mathrm{FSR}}\frac{\partial A_i(t,\tau)}{\partial \tau} = \sum_{k=2}\frac{i^{k+1}\beta_k}{k!}\frac{\partial^k A_i}{\partial t^k} + iLP_i^{\mathrm{NL}} - [(\alpha_i L + \theta_i)/2 + i\delta_{0i}]A_i(t,\tau) + \sqrt{\theta}E_{\mathrm{in}} \qquad (5.11)$$

We omit the shock term as in most Lugiato-Lefever based models. Note that because we performed a time average, we no longer have z-dependence on our differential equation. Therefore, we lose some of our ability to model spatially-varying nonlinear index or Raman gain, unless we incorporate it back in via an averaged term.

Ignoring Raman and polarization effects and using the simulated dispersions for the visible-wavelength structures studied previously, we can simulate the frequency comb expected for such devices. As an example, see the top portion of Fig. 5.5. These initial simulation show incredibly promising indications of the potential of frequency combs in diamond. However, because diamond has a particularly strong Raman component, we will see that this can actually significantly disrupt frequency comb formation in simulations where it is accounted for.

5.4 Kerr-Raman competition in microresonators

We can theoretically investigate the combined effects of Raman and parametric gain on Kerr comb generation in diamond microresonators. Since the Raman gain is large in diamond, we can expect it to play an integral part in the nonlinear dynamics in the microresonator. We show that it is possible to minimize the strength of the Raman effect by designing the device geometry. Our analytic and numerical studies show that tuning the free-spectral range (FSR) of the microresonator can be used to achieve a frequency comb.

In Fig. 5.5 we show our numerical results for Kerr comb generation in diamond. For Kerr comb generation, the FSR must be chosen such that the peak of the Raman gain is far detuned from any

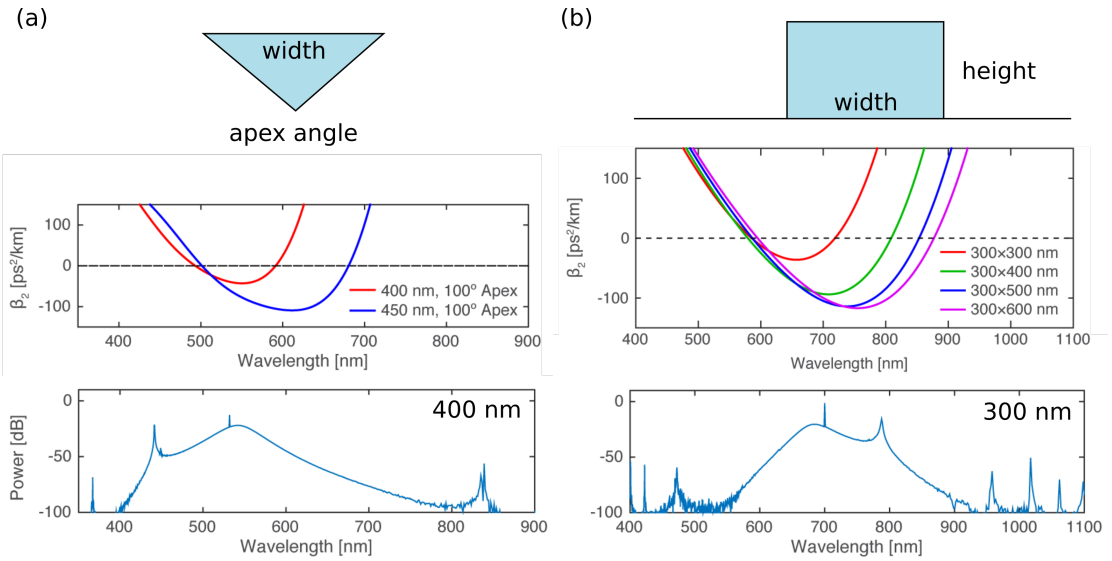

Figure 5.4: Simulated frequency combs and dispersion diagrams from collaboration with Gaeta group. (a) Simulated dispersion diagrams for diamond waveguides of triangular cross-section (top), showing anomalous dispersion well into visible wavelengths (middle) and simulated frequency combs results for 400 nm wide triangular waveguides (bottom) with over 2/3 octave generation. (b) Dispersion and frequency comb simulations for thin film-type diamond waveguide devices, for devices of 300 nm in height and various width. The bottom simulation shows a 300 nm wide waveguide.

cavity resonance in order to suppress the Raman gain as compared to the parametric (electronic) gain. This condition posts a maximum limit of the microresonator's dimension (or round-trip path length), especially for diamond microresonators operating in the visible regime where the Raman gain is typically much larger. In addition, we investigate the competing effects near the FSR limit in silicon microresonators. Our analysis offers a novel route to achieving Kerr comb generation in the presence of strong Raman effects and can be generalized to other nonlinear crystalline materials with a narrow Raman gain feature.

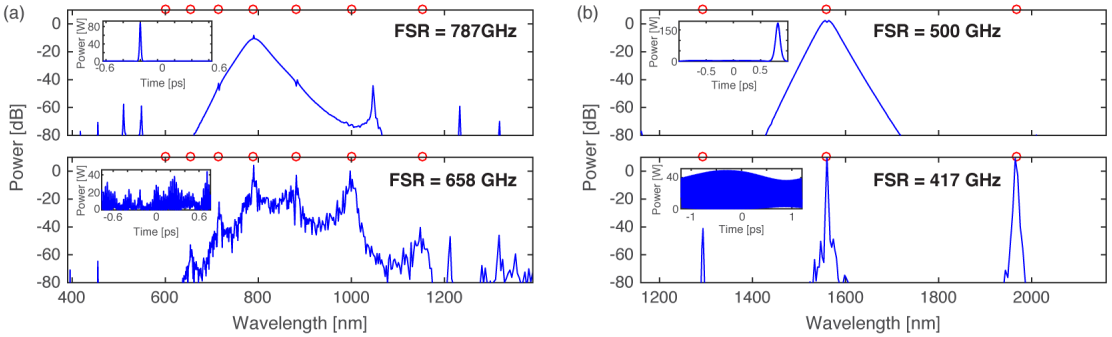

Figure 5.5: (a) Simulated spectrum for comb generation in diamond microresonators pumping at 790 nm for FSR's of 787 GHz (top) and 658 GHz (bottom). (b) Simulated spectrum a 1560 nm pump for FSR's of 500 GHz (top) and 417 GHz (bottom). The insets show the corresponding temporal profile. The red circles indicate the Raman frequency shift.

Although the results from Fig. 5.5 are strongly suggestive towards the difficulty of frequency comb generation, especially for small-FSR resonators, they are still not the final word. The simulations which Fig. 5.5 is based off of are only inclusive of a single mode with a Raman interaction. Although this is true for certain waveguide orientations in diamond, looking back at Fig. 4.1, we see that Raman transitions can be suppressed, especially if the pump beam is polarized in the TM direction. This full problem is significantly more challenging, however, because in a resonator the waveguide orientation is a function of position z.

5.5 Coupled GNLSE equations with Raman interaction

Looking back to Eq. 5.2, we see that the term $P_i^{\mathrm{NL}}(z, t)$ (via Eq. 5.3) contains implicitly within it a coupling term, describing a nonlinear interaction-mediated mixing of modes with possibly orthogonal polarizations. To go a step further, we can generalize Eq. 5.2 to include an overlap factor in the definition for the nonlinear coefficient γ_{ijkl}, denoted η_{ijkl} given by the spatial distribution of the mode $\mathbf{E}(x, y, \omega)$ [113]

$$\eta_{ijkl} = \frac{\int \mathbf{E}_i^* \mathbf{E}_j \mathbf{E}_k^* \mathbf{E}_l \, \mathrm{d}x \, \mathrm{d}y}{\left[\prod_v \int |\mathbf{E}_v|^4 \, \mathrm{d}x \, \mathrm{d}y\right]^{1/4}} \tag{5.12}$$

More details are contained in Lin et al. [113]. For the simplification of two orthogonally polarized modes with unit overlap, we can derive a Raman-mediated coupled GNLSE equation which describes the evolution of the two modes.

An important consideration is the noise term Γ_R. Consider two orthogonally polarized modes - one which is seeded by an input pulse $A_p(0, t) = U_0(t)$ and the other which is left empty $A_s(0, t) = 0$. Although a Raman- or Kerr- mediated transition is allowed between the two, it can only occur when the electric field in both modes is nonzero due to the cross-multiplied terms in Eq. 5.3. To this end, in order to start the transition into the opposite mode, it must be initially seeded by the Raman noise term. By correctly modeling this spontaneous Raman scattering, we can fully simulate a coupled GNLSE in a crystalline resonator or waveguide.

5.5.1 Numerical details and simulation results

To study the problem of coupled-GNLSE, we developed a simulation package with a general framework and easy-to-use interface. The aims of this package are to provide researchers with the functionality to quickly design and understand nonlinear optical devices, especially in crystalline ma-

terial where the direction of propagation may change over the course of the waveguide, as is the case for diamond ring or racetrack resonators. The code is freely available at the following address: https://github.com/platawiec/NonlinearOptics.jl

The package is provided in Julia, a high-performance programming language with natural extensibility and excellent support for rapid prototyping. Beyond the performance of the language, Julia was chosen due to the well-supported ecosystem for differential equation solving via the DifferentialEquations.jl framework[137], which these problems are a subset of. This framework allows us to easily extend problems - for instance, by generating an Ikeda map from a GNLSE or easily running a Monte Carlo simulation to study coherence. Furthermore, by recasting the GNLSE into the so-called interaction picture, we effectively perform a stable discretization of the PDE, allowing us to take advantage of a number of well-developed ODE solver algorithms.

Because of the excellent ecosystem, we can also rigorously extend our models to include stochastic noise, specifically Raman noise. Although this process has been well-explored, the reports in the literature for how it is treated numerically are remarkably thin[48,44]. NonlinearOptics.jl, by contrast, models the noise as a multiplicative stochastic variable, diagonal in the frequency and spatial domains, and can solve the resulting stochastic differential equation via a number of built-in adaptive time-stepping algorithms.

We show a comparison of simulation results under a number of simplifying assumptions in Fig. 5.6. shows the results of propagation of two modes along 100 axis of diamond which are orthogonally polarized. These illustrative results highlight the subtleties implicit in modeling and are meant to inspire and guide future investigations. With the development of advanced modeling tools for nonlinear optics, device engineers can start to explore and optimize structure to affect desired results. As this chapter demonstrated, a careful accounting of Raman effects and Kerr nonlinearities is required to gain a complete picture of device function, especially for the short, intense pulses of light that can occur in supercontinuum and frequency comb generation. With a coherent understand-

ing of the opportunities and issues, the full parameter space of devices can be simulated, including spatially-varying dispersion profiles, changing crystal orientations, and even time-dependent variations or connections to other nonlinear-optical phenomena, such as the electro-optic effect.

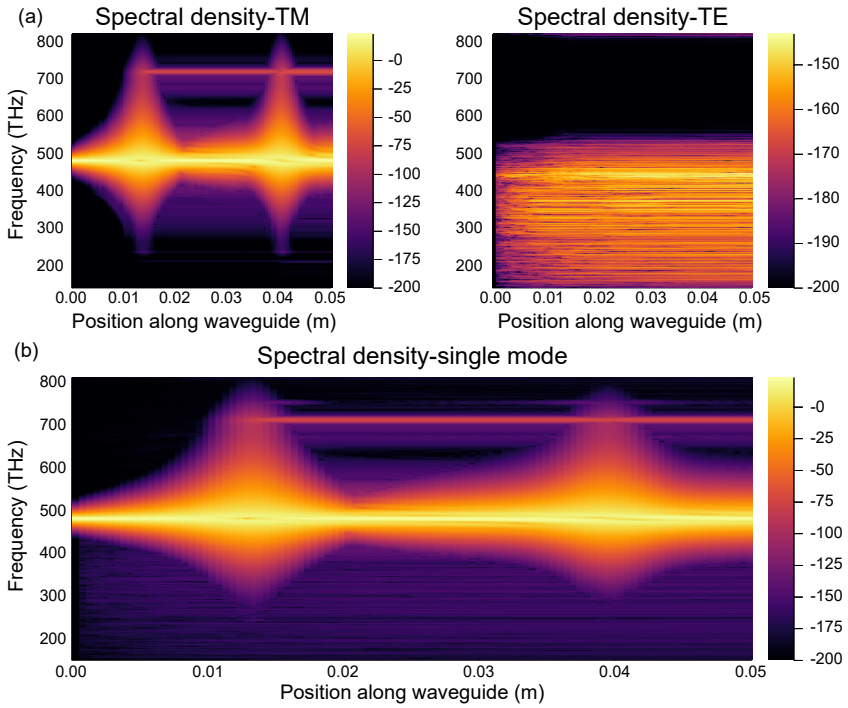

Figure 5.6: (a) Simulated spectrum for supercontinuum generation in 5 cm long-diamond waveguides oriented along the $\langle 100 \rangle$ direction, accounting for TE-TM mode coupling and the Raman noise term. The inital pulse is 100 fs long with 10 W peak power at 475 THz. Because the TE mode is initially unseeded, the only field accumulated is due to the noise term. (b) The same simulation, for the case of an isotropic material with otherwise equal material parameters. The generated supercontinuum is slightly different. This figure corresponds to standard GNLSE simulations which do not properly account for the tensorial nature of the material or mode coupling.

6

High-power metasurfaces

Using the nanofabrication techniques developed during this thesis, we can functionalize the diamond surface by decorating it with periodic or aperiodic structures. This idea, known as metasurface optics, has garnered intense commercial, industrial, and academic interest over the past decade. Diamond, being broadly transparent, possessing a relatively high refractive index, and maintaining its structural integrity even in adverse environments, can provide an excellent platform for bringing ideas in the metasurface community to niche areas. In harsh environments or for high-power laser

processes, where traditional optics can fail catastrophically, optics fashioned purely from diamond, as described in this chapter, lower system maintenance and failure modes by eliminating the use of fragile materials. Furthermore, of all the mid- and far-infrared materials, diamond is relatively easy to fashion, meaning it can form the base for advanced optical elements otherwise impossible to fabricate. In this chapter, we will discuss several optical designs and concepts and their experimental demonstrations in diamond.

6.1 Metasurface anti-reflective coating

Anti-reflective coatings on diamond are typically made by depositing a material with a refractive index equal to $n_{\text{ARC}} = \sqrt{n_{\text{air}} \cdot n_{\text{diamond}}} \approx 1.55$ with a thickness of $\lambda/(4n_{\text{ARC}})$ to ensure proper destructive interference. Such heterostructures display a decreased laser-induced damage threshold compared to bulk diamond. However, bulk diamond can be engineered to display an *effective* index equal to the required anti-reflective index condition. This can be done by etching a periodic structure in the diamond with periodicity a satisfying

$$a < \lambda_0 / n_{\text{diamond}} \tag{6.1}$$

where λ_0 is the wavelength of interest. For a wavelength of 10.6 μm, this corresponds to a periodicity of 4.4 μm. Seen from the point of view of diffractive optics, so long as any structures fabricated on the diamond satisfy this requirement, incident light will only diffract into the zeroth order[87].

In real structures, such as the truncated cones illustrated in Fig. 6.1(a), the sidewalls are not necessarily straight. This is advantageous in two ways, by reducing the index contrast at the cone apex and bottom and by providing a spatially-varying refractive index which minimizes back-reflections[61,72]. The refractive index profile along such a cone is shown schematically in Fig. 6.1(a). To estimate the refractive index at a particular z-value, we can employ the Bruggeman approximation[24,178]

$$\frac{(1-f)(\varepsilon_{\text{air}}-\varepsilon_{\text{eff}})}{\varepsilon_{\text{eff}}+g(\varepsilon_{\text{air}}-\varepsilon_{\text{eff}})}=\frac{f(\varepsilon_{\text{diamond}}-\varepsilon_{\text{eff}})}{\varepsilon_{\text{eff}}+g(\varepsilon_{\text{diamond}}-\varepsilon_{\text{eff}})} \tag{6.2}$$

Here, f denotes the filling fraction of the diamond in the substrate, and g denotes a geometry-dependent depolarization factor which we approximate as 0.5, for infinitely long ellipsoids[178]. Using this formula, we see that the ideal filling fraction is $f = 0.5$, meaning that half of etched surface is occupied by diamond.

Because of the spatially varying refractive index, this approximation is not strictly correct. Although a solution using an effective medium approach can be derived by using a transfer matrix method or rigorous-coupled-wave analysis (RCWA)[142], full finite-difference time-domain (FDTD) simulations of the etched substrate can establish the transmission of the structure with complete rigor. Using LUMERICAL, the design parameters were varied to derive the optimal diamond shape. A representative parameter sweep is shown in Fig. 6.1, where the pillar height is swept while holding the pillar top and bottom radius constant. This simulation shows an optimal design height of ~ 1.75 μm. FDTD simulations guided the fabrication and design of these antireflective, moth-eye surfaces.

To fabricate our monolithic diamond antireflective coatings, we used standard photolithographic techniques to ensure scalability. Fig. 6.2 shows the resulting etched diamond structures. First, a polycrystalline diamond substrate is cleaned in a boiling sulfuric acid and hydrogen peroxide mixture (3:1), then rinsed and sonicated in solvents (Fig. 6.2(b-i)). Then, in Fig. 6.2(b-ii), a layer of lift-off resist (LOR-3A, Shipley) is spun on the diamond substrate, followed by a layer of photoresist (S1805, Shipley). The sample is then exposed and developed (Fig. 6.2(b-iii)) such that the LOR is etched more than the top photoresist, allowing for the clean deposition of a 15 nm titanium/200 nm gold reist layer (Fig. 6.2(b-iv)) vie electron-beam deposition. The resist stack is then lifted off (Fig. 6.2(b-v)) and the sample is etched with an oxygen plasma (Unaxis Plasmatherm, 100 W substrate

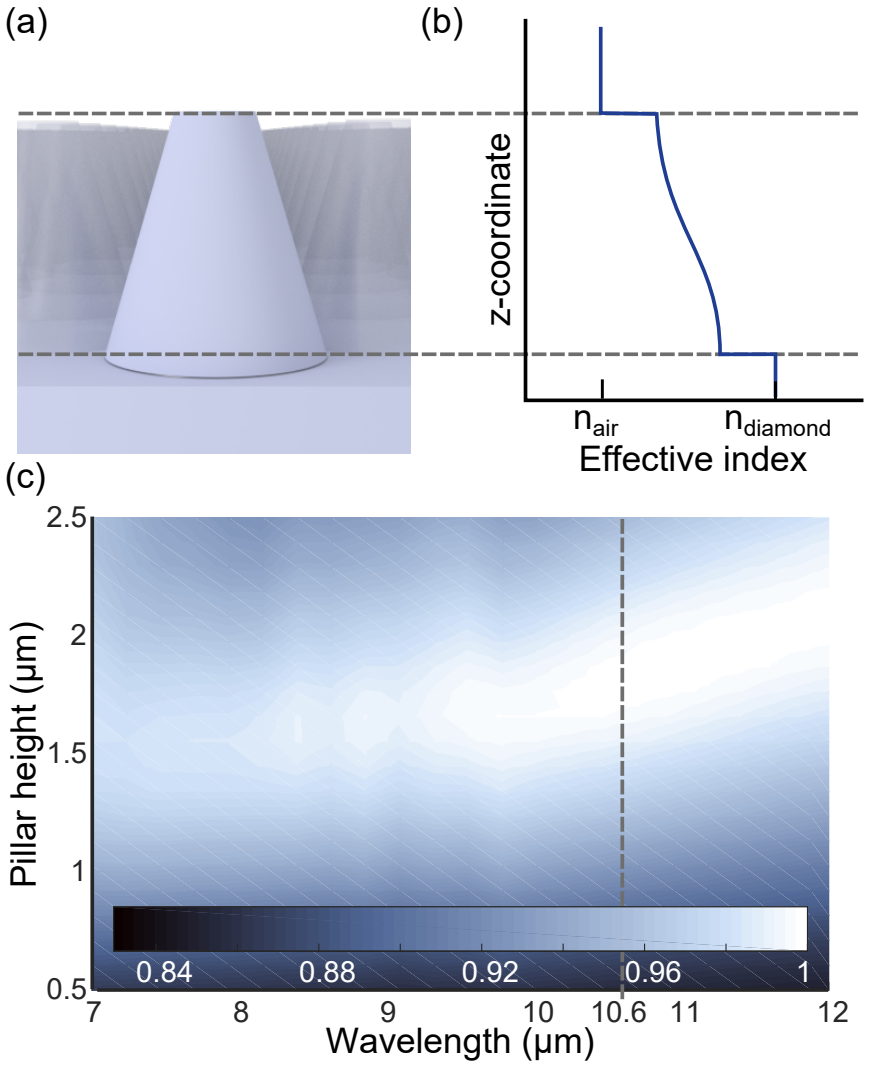

Figure 6.1: Design principle for monolithic anti-reflective coatings (a) Subwavelength dielectric pillars are fabricated in diamond with a slightly sloped sidewall. As the cross-section of the pillar varies slightly with z-coordinate (defined along the optic axis), there is an effective variation in refractive index, shown schematically in (b). The abrupt changes in refractive index at the cone-air and cone-substrate boundaries result in reflections, which can be compensated for by tuning the height of the structure for optimal destructive interference. (c) Transmission simulations shows the result of a full finite-difference time-domain simulation for the pillars shown in Fig. 2, with a top diameter of 2.8 μm and bottom diameter of 3.2 μm. The height of the simulated pillar is varied from 0.5 to 2.5 μm. The 10.6 μm wavelength result is traced with the dashed line, showing near-unity transmission for a design height of ∼1.75 μm for these dimensions.

power, 700 W ICP power). This results in a slightly sloped sidewall ($\sim$ 7 deg) which aids in the antireflective properties of the coating, as discussed previously. The sample is then cleaned to remove the resist layer, the final pattern illustrated in Fig. 6.2(b-vi).

For shorter wavelengths, targeting 1064 nm applications, a motheye anti-reflective coating was fabricated via electron beam. The pitch was around 400 nm, with conical pillars around 2 μm heigh. Transmission of over 98% was reported.

6.2 Metasurface quarter-wave plate

Diamond's excellent material properties position it to be a prime player in next-generation optical and mechanical systems. Processing techniques such as Faraday cage-angled etching[28,96,98], masked oxidation[120], quasi-isotropic plasma etching[89], and two-photon assisted etching[103] may enable new paradigms for diamond-based devices. Outside of applications in free-space mid-infrared optics like antireflective coatings[61] and wave plates[45], enormous potential remains to be tapped in the shorter side of the electromagnetic spectrum, provided that fabrication challenges stemming from the required smaller structures can be met. Beyond free-space optics, diamond may play a larger role in niche integrated optics devices, where optical components are fabricated directly on a silicon chip. Examples include integrated frequency combs[62], filters[64], Raman lasers[101], and quantum optical networks[66].

Using similar processing techniques as before, we have demonstrated precise mid-infrared quarter wave plates in diamond, with consistent phase shifts within three degrees of optimal for wavelengths of 10.6 μm. Briefly, these wave plates are fabricated by etching parallel, closely spaced ribs on the diamond surface. The anisotropy of the diamond structure causes different effective indices to be experienced for light with different polarizations. This is known as form birefringence[91,142,45]. With this technique, we have measured phase shifts of 90.42 deg in polycrystalline samples at 10.6 μm.

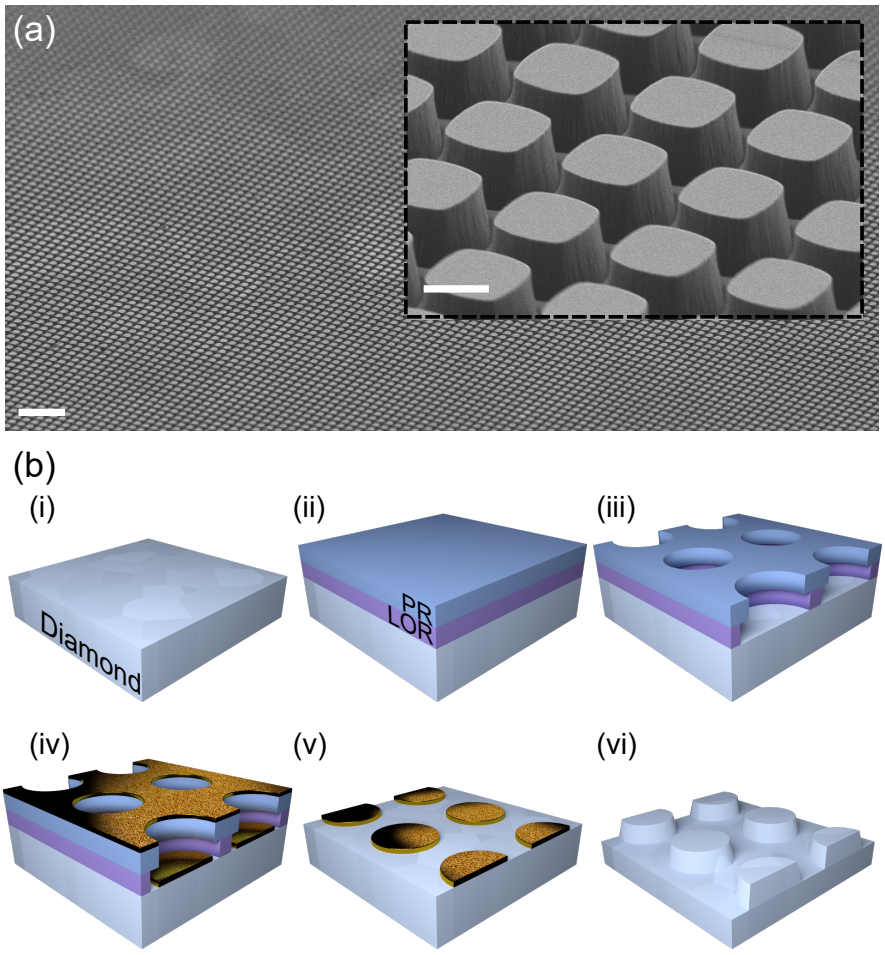

Figure 6.2: Fabrication of monolithic anti-reflective coatings in diamond (a) Scanning electron micrograph (SEM) of fabricated anti-reflective coating (scale bar 20 μm). The inset shows a zoomed-in image of the sample (scale bar 2 μm), with detail visible. The diamond pillars have a rounded-rectangular cross-section with a top side length of ∼2.8 μm and a bottom side length of ∼3.2 μm. The etch is timed to reach an optimal anti-reflective condition. (b) Process flow for fabrication of diamond anti-reflective coatings. (i) a polycrystalline diamond substrate is cleaned and then coated (ii) in a lift-off resist (LOR) and photoresist (PR) layer. (iii) The resist is then exposed and developed, with the LOR dissolving at a faster rate than the PR. (iv) A thin titanium (15 nm) is deposited, followed by a thicker (∼200 nm) gold layer via electron-beam deposition, and then (v) lifted off. The gold a mask during etching, which results in a final profile with slightly sloped sidewalls. (vi) The sample is then stripped of its mask and cleaned.

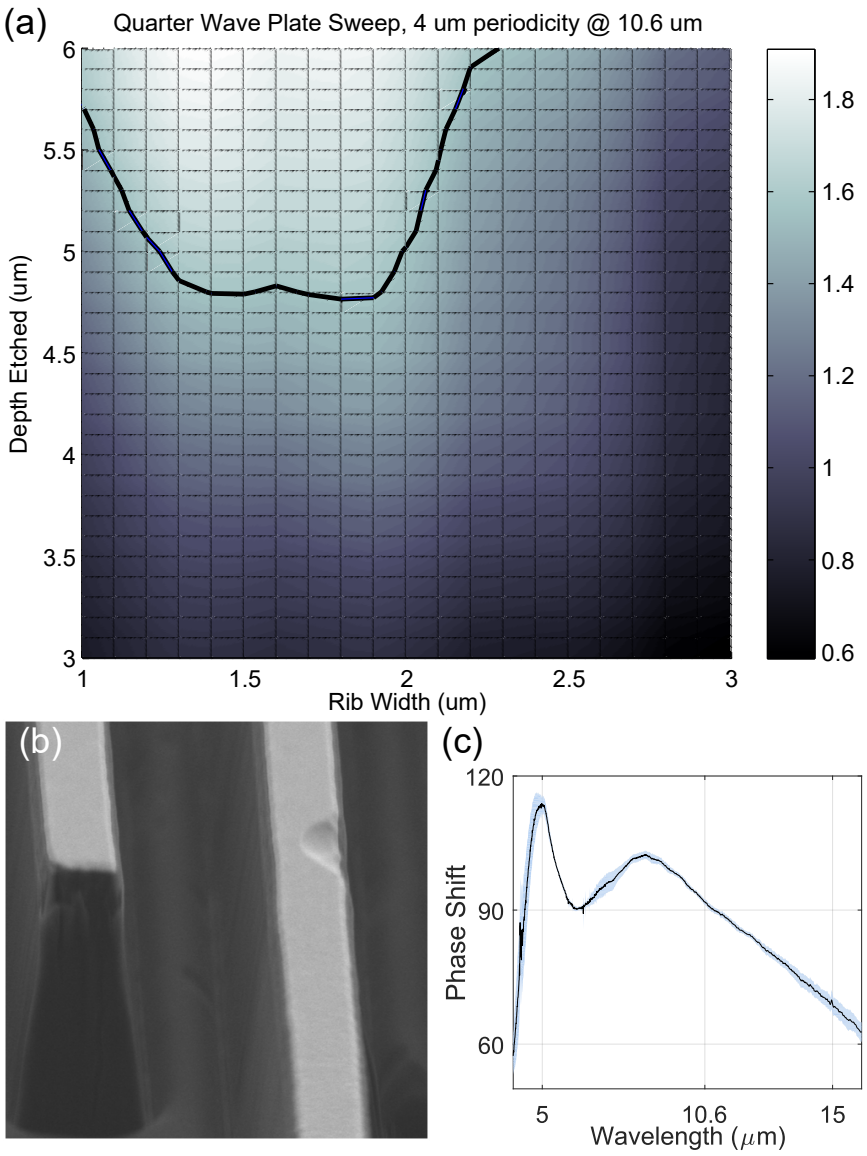

Figure 6.3: Quarter wave plate metasurface device (a) Parameter sweep of quarter wave plate design, where the depth etched and the width of the rib is varied. The thick black line describes the region where the QWP imparts a 90 degree phase shift. (b) SEM of etched QWP structure, showing ribs protruding almost 5 μm above the substrate. (c) Experimental results of phase shift imparted by QWP as a function of wavelength. The target of 90 degrees +/- 3 degrees at 10.6 μm wavelength was reached.

The sample and its measurement results are shown in Fig. 6.3. The structure was designed via a standard optimization sweep and fabricated with the liftoff procedure described in the preceding section. To obtain phase shift measurements at mid-IR wavelengths, an fourier-transform infrared spectrometer (FTIR) was used. A polarizer was placed both before and after the measured QWP, and any imperfections in the polarizers were calibrated and accounted for. These excellent results were repeatable across the substrate and over a number of devices.

6.3 Metasurface focusing lens

Finally, we take these ideas a step further and design a completely monolithic diamond metasurface lens. Borrowing from the concepts discussed earlier, we can engineer the phase shift of individual elements by changing their diameters. Such metalens designs, spurred by developments in material processing and the tantalizing applications that they can offer, have seen an explosion of research interest, with demonstrations in silicon[180], visible wavelengths[47,90,35], and plasmonic structure[110,56]. More advanced designs can offer new functionality, including polarization-functionalized imaging[4,5], retro-reflectors[6], and recently achromatic lenses[7,35]

Taking advantage of diamond's mid-IR transparency yet again, we design a metalens which functions in at 10.6 μm wavelengths, shown in Fig. 6.4(a). First, we find the phase shift a single diamond pillar imparts on an incident wavefront via FDTD simulations, the results of which are shown in 6.4(b). With this knowledge, we can choose a target pillar height and arrange the phase delay φ across the metasurface to emulate a lens, or $\varphi(r) \sim ar^2$ for some coefficient a which sets the focal length. The focal length was set to 2 mm for this design.

To characterize the lens, we perform a knife-edge measurement at different points along the beam path, with a power meter recording the incident energy as the knife is brought across the beam. This measurement gives an integrated power across the beam radius, which we fit to a double-gaussian

curve. By comparing the amount of energy in the broad gaussian to the amount of energy in the focused gaussian, we obtain a focusing efficiency of 57%, a relatively good value for a monolithic lens. The recovered beam profile is shown in Fig. 6.4(c). The beam waist at the focus can then be estimated by taking knife-edge measurements along the beam length, as shown in Fig. 6.4(d). By fitting the beam waist formula $w(z) = w_0\sqrt{1 + z^2/z_R^2}$, where z_R is the Rayleigh length given as $z_R = (\pi w_0^2)/(\lambda M^2)$, where M^2 is a measure of beam quality, we find that our minimal beam waist is 48.0 μm - far from diffraction limited. However, our fitted value for z_R reveals a poor beam quality (M^2 = 5.6, where M^2 = 1 for an ideal gaussian beam). This poor beam quality may come from the metasurface lens itself, or from the CO_2 laser used in the experiment.

In conclusion, we have demonstrated a metasurface-based diamond lens optic for focusing mid-IR light, potentially at high powers. By fashioning the device from polycrystalline diamond, we show that there is route toward a scalable process which can find use in industry in, for instance, extreme ultraviolet (EUV) photolithography systems, where high-power CO_2 lasers seed the ultraviolet generation. Overall, by combining a few simple concepts in metasurface optics to our diamond fabrication techniques, we have demonstrated a variety of different functional elements which can be used in experiment today.

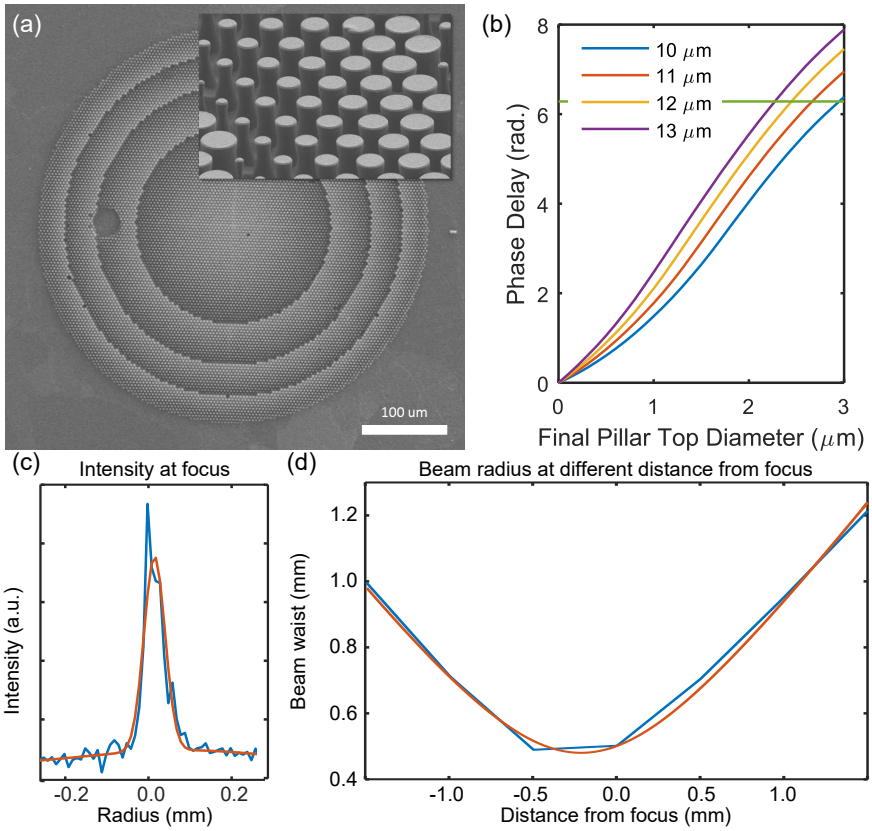

Figure 6.4: Design and fabrication of monolithic diamond metasurface lens for 10.6 μm wavelength focusing. (a) SEM of fabricated metalens. Inset shows close-up of pillars used as phase building blocks. (b) Design parameter sweep of different height pillars, assuming a 7 degree etch angle slope, as per fabrication results. A minimum pillar height of 10 μm is required to create a full 2π phase shift, though in practice taller pillars are used in order to relax the fabrication constraint. The design pitch is 4μm. (c) Results of knife-edge measurements on focused CO_2 laser beam. The lens shows focusing efficiency of 57%. (d) By repeating the knife-edge measurements over the beam path, we can fit the measured beam radius as a function of distance. The minimal focused spot radius is 48.0μm, indicating the beam is not diffraction limited. We can estimate the M^2 value to get an idea of the beam quality by comparing it to the ideal Rayleigh length. This gives us an M^2 of 5.6, suggesting imperfections in the lens or the probe beam itself.

7

Conclusion

7.1 Summary and outlook

Throughout this thesis we have explored a number of new and interesting physics brought to light because of diamond's superlative material and optical properties. Beginning with a discussion of passive photonic devices in Chapter 2, we enumerated the many possible loss mechanisms in a waveguide. This exercise helped to focus our fabrication efforts and fine-tune techniques and processes

when we found they were lacking. Furthermore, we explored entirely new materials and techniques for coupling photonics to the outside world. After briefly discussing mode converters, we described in detail a tapered-fiber technique which was used to access the spectrum of high-index microresonators.

Chapter 3 saw a discussion and analysis of fabrication techniques, how they applied to diamond photonics, and more generally how they could be applied to other material systems for three-dimensional fabrication. Fabrication via Faraday cage-angled etching proved to be an intricate subject with many subtleties in its process. Nonetheless, it could be applied successfully towards single crystal materials like diamond and quartz. Reactive ion beam undercut etching is an evolution and improvement of this technique, provided that a reactive ion beam etching tool with a tilting stage is available. This technique showed improved uniformity and was applied specifically to diamond photonic devices. Finally, we discussed the particulars of diamond fabrication used in the rest of the thesis, namely the diamond-on-insulator platform. A detailed description of the process was included, and fabrication techniques which did not improve the performance of the devices were overviewed in order to help guide the reader in their own efforts. Doped-glass mode converters and their fabrication methodology were also presented. These saw dramatic improvement in high-power handling capacity and helped to enable nonlinear optics in visible wavelengths.

Next, in Chapter 4, we touched on the theory and presented two experimental demonstrations of Raman lasers in an integrated diamond platform. The first, telecom-pumped device showed excellent performance and a wide tuning bandwidth of over 100 nm. Next, the improvements in integrated photonics for visible light were put together to demonstrate Raman lasing at these wavelengths. Not only was excellent performance demonstrated, but Raman-mediated polarization transitions were studied. Furthermore, as the Raman gain was studied over the bandwidth of the resonator, we could clearly observe hallmarks of Stokes-mode matching and enhancement only in certain regions.

Chapter 5 included a technical discussion of the Generalized Nonlinear Schroedinger Equation as it is relevant for diamond, including crystalline anisotropy and the tensorial nature of the third-order nonlinear susceptibility. From this foundation, we demonstrate four-wave mixing at visible wavelengths in a diamond platform, the shortest-wavelength demonstration to our knowledge. We also design, simulate, fabricate, and test a diamond waveguide device for supercontinuum generation at visible wavelengths in a novel architecture. These early experimental results show incredible promise for future exploration. Finally, we discuss frequency comb generation within the context of diamond microresonators. Presenting early simulations which neglected some terms in the GNLSE, we see promising early results. However, upon further investigation, we show that resonator design must be finely tuned in order to allow comb generation. Lastly, we make suggestions on how the model could be further improved, including simulations of coupled, cross-polarized modes in a nonlinear, crystalline waveguide.

The final chapter, Chapter 6, includes a brief survey of efforts towards metasurface optics made from monolithic diamond. We reviewed and presented results based on anti-reflective coatings for CO_2 laser wavelengths (10.6 μm) and extended them to quarter wave plate designs. Finally, we discuss a metasurface focusing lens design made of monolithic diamond. This lens was demonstrated experimentally and could be further improved.

In conclusion, the areas spanned by this PhD, although confined to the material of diamond, begin the process of answering emerging questions in photonics, and form the foundational groundwork for researchers beginning to look deeper into their systems. The insights gained in studying the nonlinear frequency generation processes - either Raman- or Kerr-mediated, are going to be important to apply in a variety of systems, especially in those where the Raman response is strong or the material is crystalline.

From the standpoint of diamond-based devices, be they quantum or photonic, considerable effort has been poured into reinforcing the foundations for next-generation and advanced systems and

architectures. With the careful accounting for fabrication done in this thesis, we hope that future researchers in the domain of diamond will have an easier time getting started and be able to apply some of the knowledge gleaned here into their own work. After all, though this thesis may be coming to an end, integrated optics in diamond has not seen its sunset. A number of devices can still be explored, from better-performing Raman lasers, to improved supercontinuum-generating architectures, to more interesting metasurface optics. Ultimately, we hold out hope that there may one day be a functioning frequency comb in the visible for diamond. The area of diamond photonics, and integrated nonlinear optics in general, will only see more and more activity as concepts advance and new application areas develop. With the conclusion of this thesis, we are well on our way.

References

[1] Absil, P. P., Hryniewicz, J. V., Little, B. E., Cho, P. S., Wilson, R. a., Joneckis, L. G., & Ho, P. T. (2000). Wavelength conversion in GaAs micro-ring resonators. *Optics letters*, 25(8), 554–556.

[2] Aharonovich, I., Castelletto, S., Simpson, D. a., Su, C.-H., Greentree, a. D., & Prawer, S. (2011). Diamond-based single-photon emitters. *Reports on Progress in Physics*, 74(7), 076501.

[3] Alfano, R. R. & Shapiro, S. L. (1970). Observation of self-phase modulation and small-scale filaments in crystals and glasses. *Physical Review Letters*, 24(11), 592–594.

[4] Arbabi, A., Arbabi, E., Kamali, S. M., Horie, Y., Han, S., & Faraon, A. (2016). Miniature optical planar camera based on a wide-angle metasurface doublet corrected for monochromatic aberrations. *Nature Communications*, 7, 1–29.

[5] Arbabi, A. & Faraon, A. (2017). Fundamental limits of ultrathin metasurfaces. *Scientific Reports*, 7, 1–6.

[6] Arbabi, A., Horie, Y., Ball, A. J., Bagheri, M., & Faraon, A. (2015). Subwavelength-thick lenses with high numerical apertures and large efficiency based on high-contrast transmitarrays. *Nature Communications*, 6(May), 7069.

[7] Arbabi, E., Arbabi, A., Kamali, S. M., Horie, Y., & Faraon, A. (2017). Controlling the sign of chromatic dispersion in diffractive optics with dielectric metasurfaces. *Optica*, 4(6).

[8] Atikian, H. A., Eftekharian, A., Jafari Salim, A., Burek, M. J., Choy, J. T., Hamed Majedi, A., & Lončar, M. (2014). Superconducting nanowire single photon detector on diamond. *Appl. Phys. Lett.*, 104(12), 22–26.

[9] Babinec, T. M., Hausmann, B. J. M., Khan, M., Zhang, Y., Maze, J. R., Hemmer, P. R., & Loncar, M. (2010). A diamond nanowire single-photon source. *Nature nanotechnology*, 5(3), 195–9.

[10] Barclay, P. E., Srinivasan, K., Borselli, M., & Painter, O. (2004a). Efficient input and output fiber coupling to a photonic crystal waveguide. *Optics Letters*, 29(7), 697–699.

[11] Barclay, P. E., Srinivasan, K., Borselli, M., & Painter, O. (2004b). Probing the dispersive and spatial properties of photonic crystal waveguides via highly efficient coupling from fiber tapers. *Applied Physics Letters*, 85(1), 4–6.

[12] Barwicz, T. & Haus, H. (2005). Three-dimensional analysis of scattering losses due to sidewall roughness in microphotonic waveguides. *Lightwave Technology, Journal of*, 23(9), 2719–2732.

[13] Barwicz, T., Watts, M. R., Popović, M. a., Rakich, P. T., Socci, L., Kärtner, F. X., Ippen, E. P., & Smith, H. I. (2007). Polarization-transparent microphotonic devices in the strong confinement limit. *Nature Photonics*, 1(1), 57–60.

[14] Bayn, I., Meyler, B., Salzman, J., & Kalish, R. (2011). Triangular nanobeam photonic cavities in single-crystal diamond. *New J. Phys.*, 13(2), 025018.

[15] Bayn, I., Mouradian, S., Li, L., Goldstein, J. A., Schröder, T., Zheng, J., Chen, E. H., Gaathon, O., Lu, M., Stein, A., Ruggiero, C. A., Salzman, J., Kalish, R., & Englund, D. (2014). Fabrication of triangular nanobeam waveguide networks in bulk diamond using single-crystal silicon hard masks. *Appl. Phys. Lett.*, 105(21), 211101.

[16] Bhaskar, M. K., Sukachev, D. D., Sipahigil, A., Evans, R. E., Burek, M. J., Nguyen, C. T., Rogers, L. J., Siyushev, P., Metsch, M. H., Park, H., Jelezko, F., Lončar, M., & Lukin, M. D. (2017). Quantum Nonlinear Optics with a Germanium-Vacancy Color Center in a Nanoscale Diamond Waveguide. *Physical Review Letters*, 118(22), 1–6.

[17] Borselli, M., Johnson, T., & Painter, O. (2005). Beyond the Rayleigh scattering limit in high-Q silicon microdisks: theory and experiment. *Opt. Express*, 13(5), 1515–30.

[18] Boyd, G. D., Coldren, L. A., & Storz, F. G. (1980). Directional reactive ion etching at oblique angles. *Appl. Phys. Lett.*, 36(7), 583.

[19] Boyraz, O. & Jalali, B. (2004). Demonstration of a silicon Raman laser. *Opt. Express*, 12(21), 5269–5273.

[20] Bradley, J. D. B. & Pollnau, M. (2011). Erbium-doped integrated waveguide amplifiers and lasers. *Laser and Photonics Reviews*, 5(3), 368–403.

[21] Brasch, V., Herr, T., Geiselmann, M., Lihachev, G., Pfeiffer, M. H. P., Gorodetsky, M. L., & Kippenberg, T. J. (2016). Photonic chip–based optical frequency comb using soliton Cherenkov radiation. *Science*, 351(6271), 357–360.

[22] Brenner, T., Bachmann, M., & Melchior, H. (1994). Vertically tapered InGaAsP/InP waveguides for highly efficient coupling to flat-end single-mode fibers. *Applied Physics Letters*, 65(7), 798.

[23] Broaddus, D. H., Foster, M. A., Agha, I. H., Robinson, J. T., Lipson, M., & Gaeta, A. L. (2009). Silicon-waveguide-coupled high-Q chalcogenide microspheres. *Optics Express*, 17(8), 5998–6003.

[24] Bruggeman, D. A. G. (1935). Dielectric constant and conductivity of mixtures in isotropic materials. *Ann. Phys.*, 24, 636–679.

[25] Burek, M. J., Chu, Y., Liddy, M. S. Z., Patel, P., Rochman, J., Meesala, S., Hong, W., Quan, Q., Lukin, M. D., & Lončar, M. (2014). High quality-factor optical nanocavities in bulk single-crystal diamond. *Nat. Comm.*, 5(May), 5718.

[26] Burek, M. J., Cohen, J. D., Meenehan, S. M., El-Sawah, N., Chia, C., Ruelle, T., Meesala, S., Rochman, J., Atikian, H. A., Markham, M., Twitchen, D. J., Lukin, M. D., Painter, O., & Lončar, M. (2016). Diamond optomechanical crystals. *Optica*, 3(12), 1404.

[27] Burek, M. J., Cohen, J. D., Meenehan, S. M., Ruelle, T., Meesala, S., Rochman, J., Atikian, H. A., Markham, M., Twitchen, D. J., Lukin, M. D., Painter, O., & Lončar, M. (2015). Diamond optomechanical crystals. *arXiv:*, 3(12), 1512.04166.

[28] Burek, M. J., De Leon, N. P., Shields, B. J., Hausmann, B. J. M., Chu, Y., Quan, Q., Zibrov, A. S., Park, H., Lukin, M. D., & Lončar, M. (2012). Free-standing mechanical and photonic nanostructures in single-crystal diamond. *Nano Lett.*, 12(12), 6084–6089.

[29] Burek, M. J., Meuwly, C., Evans, R. E., Bhaskar, M. K., Sipahigil, A., Meesala, S., MacHielse, B., Sukachev, D. D., Nguyen, C. T., Pacheco, J. L., Bielejec, E., Lukin, M. D., & Lončar, M. (2017). Fiber-coupled diamond quantum nanophotonic interface. *Physical Review Applied*, 8(2), 1–10.

[30] Burek, M. J., Ramos, D., Patel, P., Frank, I. W., & Lončar, M. (2013). Nanomechanical resonant structures in single-crystal diamond. *Appl. Phys. Lett.*, 103(13).

[31] Cady, W. (1922). The Piezo-Electric Resonator. *P. IRE*, 10(2), 83–114.

[32] Cai, M., Painter, O., & Vahala, K. J. (2000). Observation of critical coupling in a fiber taper to a silica-microsphere whispering-gallery mode system. *Physical Review Letters*, 85(1), 74–77.

[33] Cardenas, J., Zhang, M., Phare, C. T., Shah, S. Y., Poitras, C. B., Guha, B., & Lipson, M. (2013). High Q SiC microresonators. *Optics Express*, 21(14), 16882–16887.

[34] Chang, G., Li, C.-H., Phillips, D. F., Walsworth, R. L., & Kärtner, F. X. (2010). Toward a broadband astro-comb: effects of nonlinear spectral broadening in optical fibers. *Optics express*, 18(2002), 12736–12747.

[35] Chen, W. T., Zhu, A. Y., Sanjeev, V., Shi, Z., Lee, E., & Capasso, F. (2017). Phase and dispersion engineering of metalenses : broadband achromatic focusing and imaging in the visible. *ArXiv*, (pp. 1–30).

[36] Cheng, C. C., Scherer, A., & Yablonovitch, E. (1996). Nanofabricated Three Dimensional Photonic Crystals Operating at Optical Wavelengths. *Physica*, 17(111), 17–20.

[37] Chin, M. K. & Ho, S. T. (1998). Design and modeling of waveguide-coupled single-mode microring resonators. *Journal of Lightwave Technology*, 16(8), 1433–1446.

[38] Cho, B.-O., Hwang, S.-W., Ryu, J.-H., Kim, I.-W., & Moon, S.-H. (1999). Fabrication Method for Surface Gratings Using a Faraday Cage in a Conventional Plasma Etching Apparatus. *Electrochem. Solid St.*, 2(3), 129–130.

[39] Choy, J. T., Bradley, J. D. B., Deotare, P. B., Burgess, I. B., Evans, C. C., Mazur, E., & Lončar, M. (2012). Integrated TiO2 resonators for visible photonics. *Optics letters*, 37(4), 539–41.

[40] Chu, Y., De Leon, N. P., Shields, B. J., Hausmann, B., Evans, R., Togan, E., Burek, M. J., Markham, M., Stacey, A., Zibrov, A. S., Yacoby, A., Twitchen, D. J., Loncar, M., Park, H., Maletinsky, P., & Lukin, M. D. (2014). Coherent optical transitions in implanted nitrogen vacancy centers. *Nano Letters*, 14(4), 1982–1986.

[41] Coen, S., Randle, H. G., Sylvestre, T., & Erkintalo, M. (2013). Modeling of octave-spanning Kerr frequency combs using a generalized mean-field Lugiato–Lefever model. *Optics Letters*, 38(1), 37.

[42] Colbran, S. B. (2016). *The Element Six CVD Diamond Handbook*. Technical report, Element Six.

[43] Corr, C. S., Despiau-Pujo, E., Chabert, P., Graham, W. G., Marro, F. G., & Graves, D. B. (2008). Comparison between fluid simulations and experiments in inductively coupled argon/chlorine plasmas. *J. Phys. D Appl. Phys.*, 41(18), 185202.

[44] Corwin, K. L., Newbury, N. R., Dudley, J. M., Coen, S., Diddams, S. A., Weber, K., & Windeier, R. S. (2003). Fundamental noise limitations to supercontinuum generation in microstructure fiber. *Phys. Rev. Lett.*, 90(March), 203.

[45] Delacroix, C., Forsberg, P., Karlsson, M., Mawet, D., Absil, O., Hanot, C., Surdej, J., & Habraken, S. (2012). Design, manufacturing, and performance analysis of mid-infrared achromatic half-wave plates with diamond subwavelength gratings. *Applied optics*, 51(24), 5897–902.

[46] Del'Haye, P., Schliesser, a., Arcizet, O., Wilken, T., Holzwarth, R., & Kippenberg, T. J. (2007). Optical frequency comb generation from a monolithic microresonator. *Nature*, 450(7173), 1214–7.

[47] Devlin, R. C., Khorasaninejad, M., Chen, W.-T., Oh, J., & Capasso, F. (2016). High efficiency dielectric metasurfaces at visible wavelengths. *ArXiv*, (pp. 1–18).

[48] Drummond, P. D. & Corney, J. F. (2001). Quantum noise in optical fibers I Stochastic equations. *Journal of the Optical Society of America B*, 18(2), 139.

[49] Dudley, J. M. & Coen, S. (2002). Coherence properties of supercontinuum spectra generated in photonic crystal and tapered optical fibers. *Optics letters*, 27(13), 1180–1182.

[50] Eggleton, B. J., Luther-Davies, B., & Richardson, K. (2011). Chalcogenide photonics. *Nature Photonics*, 5(3), 141–148.

[51] Ekinci, K. L. & Roukes, M. L. (2005). Nanoelectromechanical systems. *Rev. Sci. Instrum.*, 76(6), 061101.

[52] Faraon, A., Barclay, P. E., Santori, C., Fu, K.-M. C., & Beausoleil, R. G. (2010). Resonant enhancement of the zero-phonon emission from a color center in a diamond cavity. *Nat. Photonics*, 5(April), 5.

[53] Feigel, B., Thienpont, H., & Vermeulen, N. (2016). Design of infrared and ultraviolet Raman lasers based on grating-coupled integrated diamond ring resonators. *J. Opt. Soc. Am. B*, 33(3), B5.

[54] Feve, J.-P. M., Shortoff, K. E., Bohn, M. J., & Brasseur, J. K. (2011). High average power diamond Raman laser. *Opt. Express*, 19(2), 913–922.

[55] Friel, I., Geoghegan, S. L., Twitchen, D. J., & Scarsbrook, G. a. (2010). Development of high quality single crystal diamond for novel laser applications. In C. Lewis, D. Burgess, R. Zamboni, F. Kajzar, & E. M. Heckman (Eds.), *Proc. SPIE*, volume 7838 (pp. 783819).

[56] Genevet, P., Capasso, F., Aieta, F., Khorasaninejad, M., & Devlin, R. (2017). Recent advances in planar optics: from plasmonic to dielectric metasurfaces. *Optica*, 4(1), 139.

[57] Glenday, A. G., Li, C.-H., Langellier, N., Chang, G., Chen, L.-J., Furesz, G., Zibrov, A. A., Kärtner, F., Phillips, D. F., Sasselov, D., Szentgyorgyi, A., & Walsworth, R. L. (2015). Operation of a broadband visible-wavelength astro-comb with a high-resolution astrophysical spectrograph. *Optica*, 2(3), 250.

[58] Granados, E., Spence, D. J., & Mildren, R. P. (2011). Deep ultraviolet diamond Raman laser. *Opt. Express*, 19(11), 10857–10863.

[59] Gröblacher, S., Hill, J. T., Safavi-Naeini, A. H., Chan, J., & Painter, O. (2013). Highly efficient coupling from an optical fiber to a nanoscale silicon optomechanical cavity. *Applied Physics Letters*, 103(18).

[60] Hansson, T. & Wabnitz, S. (2015). Frequency comb generation beyond the Lugiato-Lefever equation: multi-stability and super cavity solitons. *ArXiv*, (pp. 1–8).

[61] Harker, A. B. & Denatale, J. F. (1992). Diamond gradient index "moth-eye" antireflection surfaces for LWIR windows. *Proc. SPIE*, 1760, 261–267.

[62] Hausmann, B. J. M., Bulu, I., Venkataraman, V., Deotare, P., & Lončar, M. (2014). Diamond nonlinear photonics. *Nat. Photonics*, 8(April), 369–374.

[63] Hausmann, B. J. M., Bulu, I. B., Deotare, P. B., McCutcheon, M., Venkataraman, V., Markham, M. L., Twitchen, D. J., & Lončar, M. (2013a). Integrated high-quality factor optical resonators in diamond. *Nano Lett.*, 13(5), 1898–902.

[64] Hausmann, B. J. M., Bulu, I. B., Deotare, P. B., McCutcheon, M., Venkataraman, V., Markham, M. L., Twitchen, D. J., & Lončar, M. (2013b). Integrated high-quality factor optical resonators in diamond. *Nano Lett.*, 13(5), 1898–1902.

[65] Hausmann, B. J. M., Khan, M., Zhang, Y., Babinec, T. M., Martinick, K., McCutcheon, M., Hemmer, P. R., & Lončar, M. (2010). Fabrication of diamond nanowires for quantum information processing applications. *Diam. Relat. Mater.*, 19(5-6), 621–629.

[66] Hausmann, B. J. M., Shields, B., Quan, Q., Maletinsky, P., McCutcheon, M., Choy, J. T., Babinec, T. M., Kubanek, A., Yacoby, A., Lukin, M. D., & Lončar, M. (2012). Integrated diamond networks for quantum nanophotonics. *Nano Lett.*, 12(3), 1578–1582.

[67] Hausmann, B. J. M., Shields, B. J., Quan, Q., Chu, Y., De Leon, N. P., Evans, R., Burek, M. J., Zibrov, A. S., Markham, M., Twitchen, D. J., Park, H., Lukin, M. D., & Lončar, M. (2013c). Coupling of NV centers to photonic crystal nanobeams in diamond. *Nano Lett.*, 13(12), 5791–5796.

[68] Hensen, B., Bernien, H., Dreaú, A. E., Reiserer, A., Kalb, N., Blok, M. S., Ruitenberg, J., Vermeulen, R. F., Schouten, R. N., Abellán, C., Amaya, W., Pruneri, V., Mitchell, M. W., Markham, M., Twitchen, D. J., Elkouss, D., Wehner, S., Taminiau, T. H., & Hanson, R. (2015). Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature*, 526(7575), 682–686.

[69] Herkommer, C., Billat, A., Guo, H., Grassani, D., Zhang, C., Pfeiffer, M. H. P., Bres, C.-S., & Kippenberg, T. J. (2017). Mid-infrared frequency comb generation with silicon nitride nano-photonic waveguides. *ArXiv*, (pp. 1–7).

[70] Hickstein, D. D., Jung, H., Carlson, D. R., Lind, A., Coddington, I., Srinivasan, K., Ycas, G. G., Cole, D. C., Kowligy, A., Fredrick, C., Droste, S., Lamb, E. S., Newbury, N. R., Tang, H. X., Diddams, S. A., & Papp, S. B. (2017). Ultrabroadband Supercontinuum Generation and Frequency-Comb Stabilization Using On-Chip Waveguides with Both Cubic and Quadratic Nonlinearities. *Physical Review Applied*, 8(1), 1–9.

[71] Hiller, D., Zierold, R., Bachmann, J., Alexe, M., Yang, Y., Gerlach, J. W., Stesmans, A., Jivanescu, M., Müller, U., Vogt, J., Hilmer, H., Löper, P., Künle, M., Munnik, F., Nielsch, K., & Zacharias, M. (2010). Low temperature silicon dioxide by thermal atomic layer deposition: Investigation of material properties. *J. Appl. Phys.*, 107(6), 1–10.

[72] Hobbs, D. S. & MacLeod, B. D. (2005). Design, fabrication, and measured performance of anti-reflecting surface textures in infrared transmitting materials. *Proc. SPIE*, 5786(40), 349–364.

[73] Holzwarth, C. W., Barwicz, T., & Smith, H. I. (2007). Optimization of hydrogen silsesquioxane for photonic applications. *J. Vac. Sci. Technol. B*, 25(6), 2658.

[74] Hsu, C.-C., Hoang, J., Le, V., & Chang, J. P. (2008). Feature profile evolution during shallow trench isolation etch in chlorine-based plasmas. II. Coupling reactor and feature scale models. *J. Vac. Sci. Technol. B*, 26(6), 1919.

[75] Hsu, C.-C., Nierode, M. A., Coburn, J. W., & Graves, D. B. (2006). Comparison of model and experiment for Ar, Ar/O2, and Ar/O2/Cl2 inductively coupled plasmas. *J. Phys. D Appl. Phys.*, 39(15), 3272–3284.

[76] Hsu, H.-S., Cai, C., & Armani, A. M. (2009). Ultra-low-threshold Er:Yb sol-gel microlaser on silicon. *Optics express*, 17(25), 23265–23271.

[77] Huang, W.-P. (1994). Coupled-mode theory for optical waveguides: an overview. *Journal of the Optical Society of America A*, 11(3), 963.

[78] Humbert, G., Wadsworth, W., Leon-Saval, S., Knight, J., Birks, T., St. J. Russell, P., Lederer, M., Kopf, D., Wiesauer, K., Breuer, E., & Stifter, D. (2006). Supercontinuum generation system for optical coherence tomography based on tapered photonic crystal fibre. *Optics Express*, 14(4), 1596.

[79] Ilchenko, V. S., Bennett, A. M., Santini, P., Savchenkov, a. a., Matsko, a. B., & Maleki, L. (2013). Whispering gallery mode diamond resonator. *Opt. Lett.*, 38(21), 4320–3.

[80] Ilchenko, V. S., Savchenkov, a. a., Matsko, a. B., & Maleki, L. (2014). Crystalline whispering gallery mode resonators: in search of the optimal material. *SPIE Photonics West*, 8960, 896013.

[81] Izuhara, T., Gheorma, I. L., Osgood, R. M., Roy, A. N., Bakhru, H., Tesfu, Y. M., & Reeves, M. E. (2003). Single-crystal barium titanate thin films by ion slicing. *Appl. Phys. Lett.*, 82(4), 616–618.

[82] Izuhara, T., Osgood, R. M., Levy, M., Reeves, M. E., Wang, Y. G., Roy, A. N., & Bakhru, H. (2002). Low-loss crystal-ion-sliced single-crystal potassium tantalate films. *Appl. Phys. Lett.*, 80(6), 1046–1048.

[83] Jeong, H. E., Lee, J.-K., Kim, H. N., Moon, S. H., & Suh, K. Y. (2009). A nontransferring dry adhesive with hierarchical polymer nanohairs. *P. Natl. Acad. Sci. USA*, 106(14), 5639–5644.

[84] Joannopoulos, J. D., Johnson, S. G., Winn, J. N., & Meade, R. D. (2008). *Photonic Crystals: Molding the Flow of Light*. Princeton, NJ: Princeton University Press, 2 edition.

[85] Jung, H., Xiong, C., Fong, K. Y., Zhang, X., & Tang, H. X. (2013). Optical frequency comb generation from aluminum nitride microring resonator. *Optics Letters*, 38(15), 2810.

[86] Kaminskii, a. a., Ralchenko, V. G., & Konov, V. I. (2006). CVD-diamond - A novel chi(3)-nonlinear active crystalline material for SRS generation in very wide spectral range. *Laser Phys. Lett.*, 3(4), 171–177.

[87] Karlsson, M. & Nikolajeff, F. (2003). Diamond micro-optics: microlenses and antireflection structured surfaces for the infrared spectral region. *Optics express*, 11(5), 502–507.

[88] Khanaliloo, B., Jayakumar, H., Hryciw, A. C., Lake, D. P., Kaviani, H., & Barclay, P. E. (2015a). Single-Crystal Diamond Nanobeam Waveguide Optomechanics. *Phys. Rev. X*, 5(4), 041051.

[89] Khanaliloo, B., Mitchell, M., Hryciw, A. C., & Barclay, P. E. (2015b). High-Q/V monolithic diamond microdisks fabricated with quasi-isotropic etching. *Nano Lett.*, 15(8), 5131–5136.

[90] Khorasaninejad, M., Chen, W. T., Devlin, R. C., Oh, J., Zhu, A. Y., & Capasso, F. (2016). Metalenses at visible wavelengths: Diffraction-limited focusing and subwavelength resolution imaging. *Science*, 352(6290), 1190–1194.

[91] Kikuta, H., Ohira, Y., & Iwata, K. (1997). Achromatic quarter-wave plates using the dispersion of form birefringence. *Applied Optics*, 36(7), 1566–72.

[92] Kippenberg, T. J., Holzwarth, R., & Diddams, S. a. (2011). Microresonator-based optical frequency combs. *Science*, 332(6029), 555–559.

[93] Kippenberg, T. J., Spillane, S. M., Armani, D. K., & Vahala, K. J. (2004). Ultralow-threshold microcavity Raman laser on a microelectronic chip. *Opt. Lett.*, 29(11), 1224–6.

[94] Kitzler, O., McKay, A., & Mildren, R. P. (2012). Continuous-wave wavelength conversion for high-power applications using an external cavity diamond Raman laser. *Opt. Lett.*, 37(14), 2790–2.

[95] Knight, J. C., Cheung, G., Jacques, F., & Birks, T. A. (1997). Phase-matched excitation of whispering-gallery-mode resonances by a fiber taper. *Optics Letters*, 22(15), 1129–1131.

[96] Kovi, K. K., Balmer, R. S., & Isberg, J. (2015). Semi-isotropic surface etching of diamond using a Faraday cage. *Diam. Relat. Mater.*, 58, 185–189.

[97] Lake, D. P., Mitchell, M., Jayakumar, H., Santos, L. F. d., Curic, D., & Barclay, P. E. (2015). Efficient telecom to visible wavelength conversion in doubly resonant GaP microdisks. *arXiv*, (pp. 1–6).

[98] Latawiec, P., Burek, M. J., Sohn, Y.-I., & Lončar, M. (2016a). Faraday cage angled-etching of nanostructures in bulk dielectrics. *J. Vac. Sci. Technol. B*, 34, 041801.

[99] Latawiec, P., Burek, M. J., Venkataraman, V., & Lončar, M. (2015a). Waveguide-loaded silica fibers for coupling to high-index micro-resonators. *ArXiv*.

[100] Latawiec, P., Burek, M. J., Venkataraman, V., & Lončar, M. (2016b). Waveguide-loaded silica fibers for coupling to high-index micro-resonators. *Appl. Phys. Lett.*, 108, 031103.

[101] Latawiec, P., Venkataraman, V., Burek, M. J., Hausmann, B. J. M., Bulu, I., & Lončar, M. (2015b). On-chip diamond Raman laser. *Optica*, 2(11), 924–928.

[102] Lee, R. E. (1979). Microfabrication by ion-beam etching. *J. Vac. Sci. Technol. B*, 16(2), 164.

[103] Lehmann, a., Bradac, C., & Mildren, R. P. (2014). Two-photon polarization-selective etching of emergent nano-structures on diamond surfaces. *Nature Comms.*, 5, 3341.

[104] Levenson, M. D. & Bloembergen, N. (1974). Dispersion of the nonlinear optical susceptibility tensor in centrosymmetric media. *Physical Review B*, 10(10), 4447–4463.

[105] Levy, M., Osgood, R. M., Liu, R., Cross, L. E., Cargill, G. S., Kumar, A., & Bakhru, H. (1998). Fabrication of single-crystal lithium niobate films by crystal ion slicing. *Appl. Phys. Lett.*, 73(16), 2293–2295.

[106] Li, J., Lee, H., Chen, T., & Vahala, K. (2012). Characterization of a high coherence, Brillouin microcavity laser on silicon. *Optics Express*, 20(18), 369–373.

[107] Li, J., Lee, H., & Vahala, K. (2013). Microwave synthesizer using an on-chip Brillouin oscillator. *Nature communications*, 4, 1–7.

[108] Liao, M., Hishita, S., Watanabe, E., Koizumi, S., & Koide, Y. (2010). Suspended Single-Crystal Diamond Nanowires for High-Performance Nanoelectromechanical Switches. *Adv. Mater.*, 22(47), 5393–5397.

[109] Lin, C., Nguyen, V., & French, W. (1978). Wideband near-i.r. continuum (0.7–2.1 μm) generated in low-loss optical fibres. *Electronics Letters*, 14(25), 822.

[110] Lin, J., Genevet, P., Kats, M. a., Antoniou, N., & Capasso, F. (2013a). Nanostructured holograms for broadband manipulation of vector beams. *Nano Letters*, 13(9), 4269–4274.

[111] Lin, P. T., Singh, V., Cai, Y., Kimerling, L. C., & Agarwal, A. (2013b). Air-clad silicon pedestal structures for broadband mid-infrared microphotonics. *Opt. Lett.*, 38(7), 1031–3.

[112] Lin, P. T., Singh, V., Hu, J., Richardson, K., Musgraves, J. D., Luzinov, I., Hensley, J., Kimerling, L. C., & Agarwal, A. (2013c). Chip-scale Mid-Infrared chemical sensors using air-clad pedestal silicon waveguides. *Lab Chip*, 13(11), 2161–6.

[113] Lin, Q., Painter, O. J., & Agrawal, G. P. (2007). Nonlinear optical phenomena in silicon waveguides: modeling and applications. *Optics Express*, 15(25), 16604.

[114] Llobet, J., Sansa, M., Gerbolés, M., Mestres, N., Arbiol, J., Borrisé, X., & Pérez-Murano, F. (2014). Enabling electromechanical transduction in silicon nanowire mechanical resonators fabricated by focused ion beam implantation. *Nanotechnology*, 25(13), 135302.

[115] Lu, X., Lee, J. Y., Rogers, S., & Lin, Q. (2014). Optical Kerr nonlinearity in a high-Q silicon carbide microresonator. *Optics Express*, 22(25), 30826.

[116] Lubeigt, W., Bonner, G. M., Hastie, J. E., Dawson, M. D., Burns, D., & Kemp, A. J. (2010). Continuous-wave diamond Raman laser. *Opt. Lett.*, 35(17), 2994–2996.

[117] Lugiato, L. & Lefever, R. (1987). Spatial Dissipative Structures in Passive Optical Systems. *Physical Review Letters*, 58(21), 2209–2211.

[118] Lukowiak, A., Chiappini, A., Chiasera, A., Ristic, D., Vasilchenko, I., Armellini, C., Carpentiero, A., Varas, S., Speranza, G., Taccheo, S., Pelli, S., Battisha, I. K., Righini, G. C., Strek, W., & Ferrari, M. (2015). Sol–gel-derived photonic structures handling erbium ions luminescence. *Optical and Quantum Electronics*, 47(1), 117–124.

[119] MacDonald, N. C. (1996). SCREAM MicroElectroMechanical Systems. *Microelectron. Eng.*, 32, 49–73.

[120] McGray, C. D., Allen, R. A., Cangemi, M., & Geist, J. (2011). Rectangular scale-similar etch pits in monocrystalline diamond. *Diam. Relat. Mater.*, 20(10), 1363–1365.

[121] Michael, C. P., Borselli, M., Johnson, T. J., Chrystal, C., & Painter, O. (2007). An optical fiber-taper probe for wafer-scale microphotonic device characterization. *Optics Express*, 15(8), 4745–4752.

[122] Mildren, R., Convery, M., Pask, H., Piper, J., & McKay, T. (2004). Efficient, all-solid-state, Raman laser in the yellow, orange and red. *Opt. Express*, 12(5), 785–790.

[123] Mildren, R. & Rabeau, J. (2013). *Optical Engineering of Diamond*. Wiley.

[124] Mildren, R. P., Butler, J. E., & Rabeau, J. R. (2008). CVD-diamond external cavity Raman laser at 573 nm. *Opt. Express*, 16(23), 18950–18955.

[125] Mildren, R. P. & Sabella, A. (2009). Highly efficient diamond Raman laser. *Opt. Lett.*, 34(18), 2811–2813.

[126] Mitchell, M., Hryciw, A. C., & Barclay, P. E. (2014). Cavity optomechanics in gallium phosphide microdisks. *Applied Physics Letters*, 104(14), 141104.

[127] Oh, D., Sell, D., Lee, H., & Yang, K. (2014). Supercontinuum Generation in an On-Chip Silica Waveguide. *opticsinfobase.org*, 39(4), 1046–1048.

[128] Okawachi, Y., Yu, M., Venkataraman, V., Latawiec, P. M., Griffith, A. G., Lipson, M., Loncar, M., & Gaeta, A. L. (2017). Competition between Raman and Kerr effects in microresonator comb generation. *Opt. Lett.*, 42(14), 2786–2789.

[129] Oliver, D. W. & Slack, G. A. (1966). Ultrasonic attenuation in insulators at room temperature. *Journal of Applied Physics*, 37(4), 1542–1548.

[130] Osswald, S., Yushin, G., Mochalin, V., Kucheyev, S. O., & Gogotsi, Y. (2006). Control of sp2/sp3 carbon ratio and surface chemistry of nanodiamond powders by selective oxidation in air. *J. Am. Chem. Soc.*, 128(35), 11635–11642.

[131] Parrotta, D. C., Kemp, A. J., Dawson, M. D., Hastie, J. E., & Member, S. (2013). Raman Laser With Intracavity Frequency-Doubling to the Visible Region. *IEEE J. Sel. Top. in Quantum Elect.*, 19(4).

[132] Paschotta, R. (2017). Q Factor.

[133] Passaro, V. M. N. & de Leonardis, F. (2009). Investigation of SOI raman lasers for mid-infrared gas sensing. *Sensors*, 9(10), 7814–7836.

[134] Patel, R. N., Schr, T., Wan, N., Li, L., Mouradian, S. L., Chen, E. H., & Englund, D. R. (2015). Efficient Photon Coupling from a Diamond Nitrogen Vacancy Centre by Integration with Silica Fibre. *arXiv*, (pp. 1–11).

[135] Piper, J. a. & Pask, H. M. (2007). Crystalline Raman lasers. *IEEE J. Sel. Top. Quantum Elect.*, 13(3), 692–704.

[136] Pollnau, M. (2015). Rare-earth-ion-doped channel waveguide lasers on silicon. *IEEE Journal on Selected Topics in Quantum Electronics*, 21(1).

[137] Rackauckas, C. & Nie, Q. (2017). DifferentialEquations . jl – A Performant and Feature-Rich Ecosystem for Solving Differential Equations in Julia. *Journal of Open Research Software*, 5(15).

[138] Ramachandran, G. N. (1947). Photoelastic constants of diamond. *Proceedings of the Indian Academy of Sciences - Section A*, 25(2), 208–219.

[139] Ramelow, S., Farsi, A., Clemmen, S., Levy, J. S., Johnson, A. R., Okawachi, Y., Lamont, M. R. E., Lipson, M., & Gaeta, A. L. (2014). Strong polarization mode coupling in microresonators. *Optics Letters*, 39(17), 5134.

[140] Ranka, J. K., Windeler, R. S., & Stentz, a. J. (2000). Visible continuum generation in air-silica microstructure optical fibers with anomalous dispersion at 800 nm. *Optics letters*, 25(1), 25–7.

[141] Rath, P., Ummethala, S., Nebel, C., & Pernice, W. H. P. (2015). Diamond as a material for monolithically integrated optical and optomechanical devices. *Physica Status Solidi (A) Applications and Materials Science*, 15, 1–15.

[142] Richter, I., Sun, P. C., Xu, F., & Fainman, Y. (1995). Design considerations of form birefringent microstructures. *Applied Optics*, 34(14), 2421–9.

[143] Rivière, R., Arcizet, O., Schliesser, a., & Kippenberg, T. J. (2013). Evanescent straight tapered-fiber coupling of ultra-high Q optomechanical micro-resonators in a low-vibration helium-4 exchange-gas cryostat. *Review of Scientific Instruments*, 84(2013), 1–8.

[144] Rogers, L. J., Jahnke, K. D., Metsch, M. H., Sipahigil, A., Binder, J. M., Teraji, T., Sumiya, H., Isoya, J., Lukin, M. D., Hemmer, P., & Jelezko, F. (2014). All-Optical Initialization , Readout , and Coherent Preparation of Single Silicon-Vacancy Spins in Diamond. *Phys. Rev. Lett.*, 113(December), 263602.

[145] Rong, H., Jones, R., Liu, A., Cohen, O., Hak, D., Fang, A., & Paniccia, M. (2005). A continuous-wave Raman silicon laser. *Nature*, 433(7027), 725–728.

[146] Rong, H., Kuo, Y.-H., Xu, S., Liu, A., Jones, R., Paniccia, M., Cohen, O., & Raday, O. (2006). Monolithic integrated Raman silicon laser. *Opt. Express*, 14(15), 6705–6712.

[147] Rong, H., Xu, S., Cohen, O., Raday, O., Lee, M., Sih, V., & Paniccia, M. (2008). A cascaded silicon Raman laser. *Nat. Photonics*, 2(3), 170–174.

[148] Ryu, J.-H., Cho, B.-O., Hwang, S.-W., Moon, S. H., & Kim, C.-K. (2003). Trajectories of ions inside a Faraday cage located in a high density plasma etcher. *Korean J. Chem. Eng.*, 20(2), 407–413.

[149] Sabella, A., Piper, J. A., & Mildren, R. P. (2010). 1240 nm diamond Raman laser operating near the quantum limit. *Opt. Lett.*, 35(23), 3874–6.

[150] Sabella, A., Piper, J. A., & Mildren, R. P. (2011). Efficient conversion of a 1064 μm Nd:YAG laser to the eye-safe region using a diamond Raman laser. *Opt. Express*, 19(23), 23554.

[151] Sabella, A., Piper, J. a., & Mildren, R. P. (2014). Diamond Raman laser with continuously tunable output from 3.38 to 3.80 μm. *Opt. Lett.*, 39(13), 4037–40.

[152] Saifullah, M. S. M., Kurihara, K., & Humphreys, C. J. (2000). Comparative study of sputtered and spin-coatable aluminum oxide electron beam resists. *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, 18(6), 2737.

[153] Saifullah, M. S. M., Subramanian, K. R. V., Tapley, E., Kang, D. J., Welland, M. E., & Butler, M. (2003). Sub-10 nm Electron Beam Nanolithography Using Spin-Coatable TiO 2 Resists. *Nano Letters*, 3(11), 1587–1591.

[154] Sarabalis, C. J., Hill, J. T., & Safavi-Naeini, A. H. (2016). Guided acoustic and optical waves in silicon-on-insulator for Brillouin scattering and optomechanics. *ArXiv*.

[155] Shin, H., Qiu, W., Jarecki, R., & Cox, J. (2013). Tailorable stimulated Brillouin scattering in nanoscale silicon waveguides. *Nature ...*, 4, 1944.

[156] Shul, R. & Pearton, S. (2000). *Handbook of Advanced Plasma Processing Techniques*. Springer.

[157] Singh, N., Hudson, D. D., Yu, Y., Grillet, C., Jackson, S. D., Casas-Bedoya, A., Read, A., Atanackovic, P., Duval, S. G., Palomba, S., Luther-Davies, B., Madden, S., Moss, D. J., & Eggleton, B. J. (2015). Midinfrared supercontinuum generation from 2 to 6 μm in a silicon nanowire. *Optica*, 2(9), 797.

[158] Sipahigil, A., Evans, R. E., Sukachev, D. D., Burek, M. J., Borregaard, J., Bhaskar, M. K., Nguyen, C. T., Pacheco, J. L., Atikian, H. A., Meuwly, C., Camacho, R. M., Jelezko, F., Bielejec, E., Park, H., Lončar, M., & Lukin, M. D. (2016). Single-Photon Switching and Entanglement of Solid-State Qubits in an Integrated Nanophotonic System. *ArXiv*, (pp. 1–15).

[159] Sohn, Y.-I. & Lončar, M. (2015). Free-standing Nanostructures in Single-crystal Quartz. In *60th Int. Conf. EIPBN*.

[160] Sohn, Y.-I., Meesala, S., Pingault, B., Atikian, H. A., Holzgrafe, J., Gundogan, M., Stavrakas, C., Stanley, M. J., Sipahigil, A., Choi, J., Zhang, M., Pacheco, J. L., Abraham, J., Bielejec, E., Lukin, M. D., Atature, M., & Loncar, M. (2017). Engineering a diamond spin-qubit with a nano-electro-mechanical system. (pp. 1–10).

[161] Soref, R. (2010). Mid-infrared photonics in silicon and germanium. *Nat. Photonics*, 4(8), 495–497.

[162] Soref, R. (2015). Group IV photonics: Enabling 2 μm communications. *Nat. Photonics*, 9(6), 358–359.

[163] Spear, K. & Dismukes, J. P., Eds. (1994). *Synthetic Diamond: Emerging CVD Science and Technology*. Wiley.

[164] Spillane, S. M., Kippenberg, T. J., Painter, O. J., & Vahala, K. J. (2003). Ideality in a fiber-taper-coupled microresonator system for application to cavity quantum electrodynamics. *Physical review letters*, 91(4), 043902.

[165] Spillane, S. M., Kippenberg, T. J., & Vahala, K. J. (2002). Ultralow-threshold Raman laser using a spherical dielectric microcavity. *Nature*, 415(6872), 621–623.

[166] Sque, S. J., Jones, R., & Briddon, P. R. (2006). Structure, electronics, and interaction of hydrogen and oxygen on diamond surfaces. *Physical Review B - Condensed Matter and Materials Physics*, 73(8), 1–15.

[167] Srinivasan, K., Barclay, P. E., Borselli, M., & Painter, O. (2004). Optical-fiber-based measurement of an ultrasmall volume high- Q photonic crystal microcavity. *Phys. Rev. B*, 70(8), 81306.

[168] Srinivasan, K. & Painter, O. (2007). Linear and nonlinear optical spectroscopy of a strongly coupled microdisk-quantum dot system. *Nature*, 450(7171), 862–5.

[169] Taillaert, D., Bogaerts, W., Bienstman, P., Krauss, T. F., Daele, P. V., Moerman, I., Verstuyft, S., & Mesel, K. D. (2002). An Out-of-Plane Grating Coupler for Efficient Butt-Coupling Between Compact Planar. *IEEE Journal of Quantum Electronics*, 38(7), 949–955.

[170] Takahashi, S., Suzuki, K., Okano, M., Imada, M., Nakamori, T., Ota, Y., Ishizaki, K., & Noda, S. (2009). Direct creation of three-dimensional photonic crystals by a top-down approach. *Nat. Mater.*, 8(9), 721–725.

[171] Takahashi, Y., Inui, Y., Chihara, M., Asano, T., Terawaki, R., & Noda, S. (2013). A micrometre-scale Raman silicon laser with a microwatt threshold. *Nature*, 498(7455), 470–4.

[172] Thompson, J. D., Tiecke, T. G., de Leon, N. P., Feist, J., Akimov, A. V., Gullans, M., Zibrov, A. S., Vuletić, V., & Lukin, M. D. (2013). Coupling a single trapped atom to a nanoscale optical cavity. *Science*, 340(6137), 1202–5.

[173] Tiecke, T. G., Nayak, K. P., Thompson, J. D., Peyronel, T., de Leon, N. P., Vuletić, V., & Lukin, M. D. (2015). Efficient fiber-optical interface for nanophotonic devices. *Optica*, 2(2), 70.

[174] Tzoar, N. & Pascone, R. (1981). Radiation loss in tapered waveguides. *Journal of Op*, 71(9), 1107–1114.

[175] Vázquez-Córdova, S. a., Dijkstra, M., Bernhardi, E. H., Ay, F., Wörhoff, K., Herek, J. L., García-Blanco, S. M., & Pollnau, M. (2014). Erbium-doped spiral amplifiers with 20 dB of net gain on silicon. *Optics express*, 22(21), 25993–6004.

[176] Walavalkar, S. S., Homyk, A. P., Henry, M. D., & Scherer, A. (2013). Three-dimensional etching of silicon for the fabrication of low-dimensional and suspended devices. *Nanoscale*, 5(3), 927–31.

[177] Wang, C., Burek, M. J., Lin, Z., Atikian, H. A., Venkataraman, V., Huang, I.-C., Stark, P., & Lončar, M. (2014). Integrated high quality factor lithium niobate microdisk resonators. *Opt. Express*, 22(25), 30924.

[178] Wang, H., Liu, X., Wang, L., & Zhang, Z. (2013). Anisotropic optical properties of silicon nanowire arrays based on the effective medium approximation. *Proc. SPIE*, 8495, 62–69.

[179] Wang, T.-J., Chu, C.-H., & Lin, C.-Y. (2007). Electro-optically tunable microring resonators on lithium niobate. *Opt. Lett.*, 32(19), 2777–2779.

[180] West, P. R., Stewart, J. L., Kildishev, A. V., Shalaev, V. M., Shkunov, V. V., Strohkendl, F., Zakharenkov, Y. a., Robert, K., & Byren, R. (2014). All-dielectric subwavelength metasurface focusing lens. *Optics express*, 22(21), 1593–1595.

[181] Williams, O. A. & Jackman, R. B. (2003). Surface conductivity on hydrogen terminated diamond. *Semiconductor Science and Technology*, 18(3).

[182] Williams, R. J., Kitzler, O., McKay, A., & Mildren, R. P. (2014). Investigating diamond Raman lasers at the 100 W level using quasi-continuous-wave pumping. *Opt. Lett.*, 39(14), 4152.

[183] Xiong, C., Pernice, W. H. P., Ryu, K. K., Schuck, C., Fong, K. Y., Palacios, T., & Tang, H. X. (2011). Integrated GaN photonic circuits on silicon (100) for second harmonic generation. *Opt. Express*, 19(11), 10462–10470.

[184] Xiong, C., Pernice, W. H. P., & Tang, H. X. (2012). Low-loss, silicon integrated, aluminum nitride photonic circuits and their use for electro-optic signal processing. *Nano Letters*, 12(7), 3562–3568.

[185] Xu, Q. & Lipson, M. (2006). Carrier-induced optical bistability in silicon ring resonators. *Optics Letters*, 31(3), 341–343.

[186] Yang, K. Y., Beha, K., Cole, D. C., Yi, X., Del'Haye, P., Lee, H., Li, J., Oh, D. Y., Diddams, S. A., Papp, S. B., & Vahala, K. J. (2016). Broadband dispersion-engineered microresonator on a chip. *Nature Photonics*, 10(5), 316–320.

[187] Yang, L., Carmon, T., Min, B., Spillane, S. M., & Vahala, K. J. (2005). Erbium-doped and Raman microlasers on a silicon chip fabricated by the sol-gel process. *Applied Physics Letters*, 86(9), 1–3.

[188] Yang, L., Saavedra, S. S., Armstrong, N. R., & Hayes, J. (1994). Fabrication and characterization of low-loss, sol-gel planar waveguides. *Analytical chemistry*, 66(8), 1254–63.

[189] Yariv, A. (1973). Coupled-Mode Theory for Guided-Wave Optics. *IEEE Journal of Quantum Electronics*, QE-9(9), 919–933.

[190] Yariv, A. (2002). Critical coupling and its control in optical waveguide-ring resonator systems. *IEEE Photonics Technology Letters*, 14(4), 483–485.

[191] Yi, X., Yang, Q. F., Zhang, X., Yang, K. Y., Li, X., & Vahala, K. (2017). Single-mode dispersive waves and soliton microcomb dynamics. *Nature Communications*, 8, 1–9.

[192] Yoon Oh, D., Yang, K. Y., Fredrick, C., Ycas, G., Diddams, S. A., & Vahala, K. J. (2017). Coherent ultra-violet to near-infrared generation in silica ridge waveguides. *Nature Communications*, 8, 13922.

[193] Yu, M., Okawachi, Y., Griffith, A. G., Lipson, M., & Gaeta, A. L. (2016a). Modelocked mid-infrared frequency combs in a silicon microresonator. *ArXiv*, (pp. 1–8).

[194] Yu, M., Okawachi, Y., Griffith, A. G., Picqué, N., Lipson, M., & Gaeta, A. L. (2016b). Silicon-chip-based mid-infrared dual-comb spectroscopy. *ArXiv*, (pp. 1–5).

[195] Zakharov, V. E. & Manakov, S. V. (1974). On the complete integrability of a nonlinear Schrodinger equation. *Theoretical and Mathematical Physics*, 19(3), 551–559.

[196] Zalalutdinov, M. K., Ray, M. P., Photiadis, D. M., Robinson, J. T., Baldwin, J. W., Butler, J. E., Feygelson, T. I., Pate, B. B., & Houston, B. H. (2011). Ultrathin single crystal diamond nanomechanical dome resonators. *Nano Lett.*, 11(10), 4304–8.

[197] Zhong, T., Kindem, J. M., Miyazono, E., & Faraon, A. (2015). Nanophotonic coherent light–matter interfaces based on rare-earth-doped crystals. *Nat. Comm.*, 6(May), 8206.

This thesis was typeset using LaTeX, originally developed by Leslie Lamport and based on Donald Knuth's TeX. The body text is set in 11 point Egenolff-Berner Garamond, a revival of Claude Garamont's humanist typeface. The above illustration, "Science Experiment 02", was created by Ben Schlitter and released under CC BY-NC-ND 3.0. A template that can be used to format a PhD thesis with this look and feel has been released under the permissive MIT (X11) license, and can be found online at github.com/suchow/Dissertate or from its author, Jordan Suchow, at suchow@post.harvard.edu.