Preface

The intent of this preface is to present an informal history of my research as a graduate student at Caltech. It is my hope that this introduction will provide a readable overview of my research, which will be described in much greater technical detail in the bulk of the thesis. In addition, I hope that this section will help the reader understand the logic behind our choices in research topics. On a personal note, the work described in this thesis has consumed an enormous amount of my time and energy, and I would like to take this opportunity to document some aspects of our research that are not included in journal papers. This includes tasks such as infrastructure building, which take place in any graduate career, but are particularly prevalent when you work for a newly minted faculty member. Finally, I will use this preface as a means to discuss how our work fits in with the work of others, and will attempt to provide the reader with some sense of the progress that has been made subsequent to our contributions in a given area.

Starting graduate school

Although I officially started graduate school in September, 2001, my introduction to optical microcavities and photonic crystals began a couple of months earlier, in discussions with Dr. Oskar Painter, who had not yet joined the faculty at Caltech, but whom I already knew through our work at XPonent Photonics (a start-up company in nearby Monrovia, CA). Oskar gave me the opportunity to work with him on a couple of papers he was writing. The first was a review article on how one can tailor the properties of photonic crystal microcavities [1], while the second was a short letter on the polarization properties of a pair of modes within a single defect, hexagonal lattice photonic crystal microcavity [2]. Along with a careful reading of his thesis [3] and Joannopoulos's and Sakoda's books on photonic crystals [4, 5], working on these papers gave me a thorough introduction to photonic crystal (PC) microcavities. During this time, Oskar and I had several discussions as to what type of research he would pursue upon starting at Caltech (as a Visiting Associate in the fall of 2001, and as an Assistant Professor in January, 2002). I committed to joining his group, confident that I would have the opportunity to engage in the electromagnetic design, fabrication, and optical characterization of microphotonic devices.

The two projects that Oskar and I discussed in greatest detail were the development of high quality factor (Q) photonic crystal microcavities, and efficient evanescent coupling between optical fiber tapers and photonic crystal waveguides (PCWGs). We agreed that I would pursue the former topic, while my officemate Paul Barclay, who joined Oskar's group a few weeks after I did, would focus on the latter. From my experiences at XPonent Photonics, I was already very familiar with the finite-difference time-domain (FDTD) method for electromagnetic simulations [6], which is well suited for studying the loss properties of wavelength-scale devices. As a result, there were very few barriers to plunging into research during my first year of graduate school.

Designing photonic crystal microcavities

One of the most interesting and useful aspects of photonic crystal microcavities is their ability to confine light to an extremely small mode volume (V_{eff}), approaching a cubic half-wavelength in the material [7, 8]. These ultrasmall volumes correspond to extremely large per photon electric field strengths, which translate to strong light-matter interactions. This is at the heart of a number of potential applications in quantum optics, nonlinear optics, laser physics, and sensing/detection. The advantages of an ultrasmall volume are most potent when combined with a high cavity quality factor (Q), which corresponds to a long photon lifetime within the cavity. One specific application of interest to us is strong coupling in cavity quantum electrodynamics (cavity QED, or cQED), which examines the coherent interaction of a quantized electromagnetic field with a two-level system, such as an atom or a semiconductor quantum dot (which, at a first level of approximation, is a two-level system). Strong coupling places very strict requirements on the cavity used: the atom-photon coupling rate g, which scales as $1/\sqrt{V_{\text{eff}}}$, must exceed the atom decay rate $\gamma_{\perp}$ and the cavity decay rate κ, which scales as $1/Q$. The strong coupling regime (fig. 1) is one in which coherent interactions between a single atom and a single photon can take place, and is at the core of many applications of cavity QED in quantum information processing and computing.

Within the optical domain, cavity QED experiments have typically taken place in a system consisting of a single atom coupled to a mode of a high-finesse Fabry-Perot cavity [9]. Over the past several years, there has been a real push to extend this work to semiconductor-based systems, where the Fabry-Perot cavity would be replaced by some semiconductor microcavity, and the atom would be replaced by a semiconductor quantum dot [10]. There are a number of motivating factors behind this drive, as semiconductor-based systems offer significant experimental simplification in comparison to the atom-Fabry-Perot systems. In particular, the semiconductor microcavities are

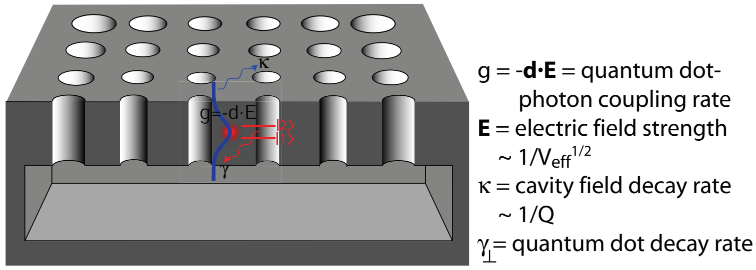

Figure 1: Schematic of quantum-dot-photon coupling in a photonic crystal cavity.

monolithic elements that should not require active stabilization, and a semiconductor quantum dot can be naturally integrated in a microcavity during the initial growth of the material comprising the device. Furthermore, the quantum dot is physically trapped within the device, eliminating the infrastructure required to trap atoms. Finally, these devices are potentially scalable and easy to integrate, as a result of the planar processing techniques by which they are created.

As I mentioned above, one of the key properties of semiconductor microcavities is the very small volumes to which they confine light. These small volumes lead to a very strong coupling rate to even a single quantum dot, so that $g/2\pi \sim 10$ GHz for $V_{\mathrm{eff}} \sim (\lambda/n)^3$. For comparison, typical coupling rates in an atom-Fabry-Perot system are on the order of 10–100 MHz. Nevertheless, even for coupling strengths of several GHz, a cavity $Q \sim 10^4$ is needed to be within the strong coupling regime. As of a couple of years ago, these Qs were not accessible to ultrasmall volume semiconductor microcavities.

Indeed, when I first started this work, the performance of photonic crystal (PC) microcavities, in terms of Q, was significantly worse than it is today. The highest predicted Qs from simulations were on the order of 10^4 [7, 11], and the highest experimentally demonstrated Qs were an order of magnitude less at 2,800 [12]. Nevertheless, PC microcavities still seemed to be a preferred architecture, as the experimentally demonstrated Qs in semiconductor microdisk cavities ($\sim 12,000$ in ref. [13] and $\sim 20,000$ in ref. [14]) were not sufficiently high to compensate for their increased modal volumes (over 10 times larger than those of the PC microcavities). Micropost cavities [15, 16] had experimentally demonstrated Qs that were even lower than those of the PC cavities, for slightly larger mode volumes. Table 1 presents, to the best of my knowledge, the state of the art for microcavities in early 2002 (in this table, I have only considered experimentally demonstrated results). The cavities are generally classified into two types; large mode volume, ultra-high Q structures such as Fabry-Perot cavities and microspheres, and wavelength-scale, relatively low-Q

Table 1: Q and Veff for several experimentally demonstrated microcavity structures, as of 2002

Fabry-Perot	Microsphere	Micropost	Microdisk	Photonic Crystal
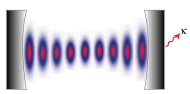	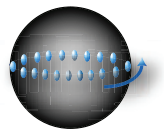	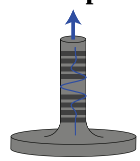	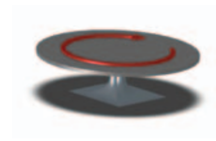	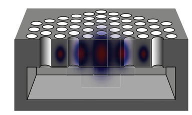

Geometry	Material	Q	$V_{\mathbf{eff}}((\lambda/n)^3)$	Reference
Fabry-Perot	Air (w/dielectric mirrors)	10^9	10^5	[17, 18]
microsphere	SiO_2	10^9	10^4	[19, 20]
micropost	GaAs/AlGaAs	2×10^3	5	[15]
microdisk	GaAs/AlGaAs	1.2×10^4	6	[13]
photonic crystal	InP/InGaAsP	500	0.3	[8]
	GaAs/AlGaAs	2.8×10^3	0.4	[12]

structures such as semiconductor microdisks, microposts, and photonic crystals. As our interests were focused on cavity QED with semiconductor quantum dots, we were most interested in the development of semiconductor microcavities. The PC cavities appeared to be the most promising candidate, as their ultrasmall volumes implied that the requisite Qs for achieving strong coupling would be less than what would be necessary in a micropost or microdisk.

In approaching the design of PC microcavities, Oskar and I set the following goal: a theoretical $Q \gtrsim 5\times10^4$ for $V_{\text{eff}} \sim (\lambda/n)^3$. These numbers were chosen because they would be sufficient, in the ideal case, to allow for strong coupling experiments in cQED. More important than these numbers, we wanted to understand and elucidate the causes for radiation loss in these devices, and to develop a fairly general framework that could be used to design high-Q PC cavities. Of particular interest to us was the development of designs that would be robust to perturbations such as those that one might encounter in the fabrication of the cavities. This seemed to us to be a crucial point, as relatively small discrepancies between a fabricated device and the intended structure would, in many designs, cause the maximum achievable Q to degrade by almost an order of magnitude, such as in ref. [12].

As I describe in chapter 2, the natural way to study losses in these devices is in momentum space. In particular, by considering the Fourier components of the cavity mode, one can understand the sources of loss; small momentum components lead to vertical radiation losses, while in-plane losses occur when the cavity mode contains components in regions of momentum space for which the photonic lattice is no longer reflective. Taking the spatial Fourier transform of a generated cavity mode serves as a diagnostic for understanding the source of loss in a given design. In addition to our paper on this topic [21], which was published in July, 2002, Vučković and co-workers published a paper during the same month that also noted the relationship between vertical radiation loss and the presence of small momentum components within the spatial Fourier transform of the cavity mode [22].

Rather than just having a diagnostic tool for understanding the performance of a given design, we were also interested in the development of the designs. Knowing which parts of momentum space are to be avoided is the first step toward this. As a second step, we chose to consider modes of a specific odd symmetry, as this leads to zero DC Fourier components and an automatic reduction in vertical radiation loss. We combined this with a methodology for tailoring the defect regions in these cavities in such a way so as to avoid problematic regions in Fourier space, and used FDTD simulations to quantitatively calculate the properties of the designs we generated. We were able to meet our goals, and develop designs in both square and hexagonal lattice photonic crystals with Qs as high as 10^5 and $V_{\text{eff}} \sim (\lambda/n)^3$ [21, 23]. These designs used a grading of the hole radius in the cavity, both to help achieve high-Qs and to make the cavity robust to perturbations to the design. We did not attempt to optimize our results, as Qs of 10^5 were sufficient for the initial experiments of interest to us, and more importantly, at the time, the experimentally demonstrated Qs for PC cavities were significantly lower than what had been predicted theoretically, so it was unclear whether further design optimization would be beneficial.

In addition to using FDTD, we were able to develop a number of semianalytic tools to aid in the cavity design (chapter 1). One of the most useful tools was a simple symmetry-based analysis that used the basic tools of point group theory to classify the symmetries of PC defect cavity modes and determine their dominant Fourier components [1, 3]. The results of this analysis meshed perfectly with the cavity design principles described above, and we relied upon it heavily in our high-Q designs [21, 23]. In particular, this analysis identified the modes within defect cavities in square and hexagonal lattice photonic crystals that satisfied our symmetry criterion for reducing vertical radiation loss. Oskar and I wrote a separate article that fully details the development of this symmetry-based analysis [24], which was published in July, 2003. In the same month, we published another article [25], which details the development of a Wannier-like equation that yields the envelope for resonant modes within PC defect cavities. These two methods are very complementary. The group theory analysis identifies and classifies the symmetries of defect modes within PC cavities, but has nothing to say about the localization of the modes. The Wannier analysis, on the other hand, describes the localizing effect of the PC lattice on the defect modes. When combined with the symmetry analysis, this Wannier method qualitatively matches the results of the FDTD simulations.

After these initial publications [21, 23, 24, 25], we did not return to the topic of PC microcavity design, but several research groups across the world continued to work on this topic, and have succeeded in developing designs with predicted Q factors in excess of 10^6 [26, 27, 28, 29, 30]. Momentum space design principles remain the basis for most of these works, albeit in combination with new geometries such as modified photonic crystal waveguides [28, 30], and the use of inverse design and optimization methods [31, 29] to help reduce the dependence on the trial-and-error approach that dominated early cavity design work. From my perspective, the Wannier and symmetry-based methods remain attractive starting points for any cavity design (whether or not the focus is on high-Q structures), as they provide a wealth of information with very little computational expense, and can be used to build physical insight on how to appropriately tailor the cavity geometry for the application of interest. Ultimately, FDTD or some other numerical simulation method will be used to calculate the cavity properties in detail, and it is likely that some further amount of tweaking of the cavity design will be required. During this step, the aforementioned optimization and inverse design methods will be of great use.

Building the cleanroom and characterization lab

While developing the PC cavity design techniques, I also spent a great deal of time helping Oskar plan the setup of our labs. Oskar, Paul, and I were able to attend meetings with the project manager in charge of building our labs, and we carefully designed a cleanroom facility and a characterization lab that would be able to support the experiments that we had planned for the first few years of the research group. The centerpiece of the cleanroom would be an Oxford Instruments Plasma Technology (OIPT) inductively coupled plasma reactive ion etch (ICP-RIE) tool clustered to a plasma-enhanced chemical vapor deposition (PECVD) tool. This tool, along with a Hitachi cold cathode, field emission scanning electron microscope (SEM), would allow us to do all of the fabrication steps necessary for creation of optical devices such as the photonic crystal waveguides and cavities that Paul and I were studying. The ICP-RIE/PECVD and SEM would be located in a Class 1000 portion of the cleanroom, and would be adjacent to service chaises that would house the support equipment for the machines (such as vacuum pumps, chillers, gas lines, and bottles). In addition, the cleanroom would contain a Class 100 section that would house fume hoods, a spinner, and a mask aligner for doing optical lithography. Our characterization lab would initially contain setups that would allow for the fabrication of tapered optical fibers, optical probing of microfabricated devices in the telecommunications band ($\lambda \sim$1.5 μm) using these tapers, and microphotoluminescence measurements of light-emitting devices. Within the period of September 2001 – June 2002, we ordered the bulk of all of the equipment to be housed in the cleanroom and characterization labs, and helped install basic equipment such as the optical tables. By June 2002, I had finished the bulk of our work on the design of high-Q PC microcavities as well as coursework for the year, so that I could focus exclusively on lab work. That summer, we installed the SEM and ICP-RIE, the latter of which required a considerable amount of setup time. Oskar, Paul, Tom Johnson (another graduate student of Oskar's who had just joined the group), and I did all of the stainless steel and PVC plumbing required to service the machines and their accompanying vacuum pumps. By August, we were ready to begin using the tools that we had spent the last couple of months installing.

In August 2002, Oskar and I flew to Murray Hill, New Jersey, to meet with members of the quantum cascade laser group at Bell Laboratories, headed by Dr. Federico Capasso. Our primary collaborator at Bell Laboratories was Dr. Raffaele Colombelli, and he and Dr. Mariano Troccoli met with us and discussed a project that they and Oskar had initiated earlier, which was the incorporation of photonic crystal microcavities within a quantum cascade heterostructure to create electrically injected microcavity laser arrays operating at mid-infrared wavelengths. Initially, our role was primarily to provide photonic crystal design expertise. However, after visiting Raffaele and Mariano, we learned that they were having great difficulty etching the quantum cascade heterostructures using their plasma etching system. Oskar and I agreed to attempt this part of the fabrication at Caltech, while Bell Labs would remain responsible for material growth, all other device fabrication steps, and subsequent device characterization.

Developing fabrication processes

After returning from Bell Labs, I began work in the cleanroom in earnest. Having had some experience in the development of plasma etching processes, I took the lead on developing the ICP-RIE etch recipes, while Paul was in charge of getting the electron beam lithography to work. For Paul, this meant starting with a code that had originally been written by a Caltech undergraduate, Oliver Dial, to control the scan coils and stage of a Hitachi SEM. Paul needed to update the code for our SEM (a different model than that originally used), data acquisition card, and operating system. For me, the ICP-RIE etching would involve building upon my previous experience in plasma processing to develop recipes for etching silicon and InP-based materials. We had decided on silicon-on-insulator (SOI) as a platform for near-infrared (near-IR) PCs, due to its low optical losses in this wavelength range and the ready availability of high quality wafers from commercial vendors such as SOITEC. InP-based materials were of potential interest for future experiments with PC microcavity lasers in quantum-well-based materials, including the quantum cascade photonic crystal laser project we had agreed to work on with Bell Laboratories.

The development of fabrication processes for creating micro-optical devices is really at the heart of the work that I have done in graduate school. Very simply put, the physical phenomena that we have been interested in require superior device performance. For example, as described above, to reach the strong coupling regime in cQED, wavelength-scale PC microcavities with Qs in excess of 10^4 are required. Thus, all fabrication-induced losses (such as roughness-induced scattering and etch-induced absorption) had to be minimized as much as possible.

The process flow for fabricating a device such as a photonic crystal microcavity typically consists of 1) deposition of a hard mask layer, 2) coating of the sample in electron beam resist and subsequent electron beam lithography, 3) plasma etching (also known as dry etching) of the mask layer, and 4) plasma etching of the primary material layer (typically a semiconductor layer in the applications we consider). For some devices, such as the passive PC resonators and optically pumped lasers described in this thesis, these steps are followed by a wet etch step to undercut the devices. For more complicated structures, such as electrically contacted devices, a number of additional fabrication steps are required. In appendix C, I provide a qualitative overview of some of the considerations that must be taken into account when developing dry etch processes for microphotonic structures.

The first material system that we had significant success with was the quantum cascade material, a heterostructure containing dozens of InP-based layers. The requirements on this etch were relatively strict; we needed to etch $4-5$ μm deep holes (with a radius of ~ 1 μm) that were as vertical as possible. In this case, we had the fortune of starting with a very good etch mask, consisting of a 500 nm thick SiO_2 PC mask that had been fabricated by our colleagues at Bell Labs. The process we developed for etching the semiconductor layers, which will be described briefly in appendixes A and C, was able to meet our requirements, and gave us the first indication that we would be able to fabricate world-class structures with our system.

At the same time, we also worked towards fabricating near-IR PCs in silicon and InP. Our group was exclusively responsible for the fabrication of these devices, so that we needed to develop the electron beam lithography and all subsequent etches. By September 2002, Paul had succeeded in adapting the control software to allow us to do the lithography with our SEM, and we began by using polymethylmethacrylate (PMMA) as our electron beam resist. We were able to do reasonably good lithography with this resist, but its performance under dry etching was very poor. There was no question that we would need to use a dielectric mask before transferring the PC pattern into the semiconductor layer (either Si or InP), but the performance of the PMMA under dry etching was so bad that it seemed unlikely that we would be able to suitably etch our dielectric mask (SiO_2) with it. For about four months, we struggled with the fabrication, and attempted to use a two-level mask consisting of a Au metal layer on top of SiO_2, with the hope that the PMMA would be more robust to the etch that would transfer the PC pattern to Au, at which point the Au would serve as a strong etch mask for the SiO_2 layer. Although we made some progress, none of our results were up to our expectations, and we did not attempt to optically characterize any of the devices created using these processes.

In January 2003, we switched electron beam resists to ZEP-520A, manufactured by Zeon Chemicals. Oskar had heard of this resist while at a conference, and in looking at its specifications, we realized that it was developed to be significantly more dry etch resistant than PMMA. Paul and I were quickly able to recalibrate our lithography for this new resist, and the resulting PC patterns looked very good. For the time being, we decided to exclusively focus on Si fabrication, and furthermore, we scrapped the original process we were working on (an Ar/Cl_2 etch) and started from scratch, using a SF_6/C_4F_8 chemistry. As briefly described in chapter 4 and in more detail in appendix C, this chemistry etches silicon beautifully, particularly for our purposes, where deep etches were not required (in such instances, a cryogenic temperature etch or the Bosch process etch are typically used). We were quickly able to develop an etch that produced smooth, vertical sidewalls. Equally important, this was a low power etch, so that the ZEP electron beam mask suffered minimal etch damage. As a result, we tried using this electron beam resist as the only etch mask for the silicon layer, foregoing the dielectric mask that we had previously attempted to use. This worked extremely well, so that creation of testable PC devices was imminent. Paul and I calibrated the electron beam lithography to create patterns that matched our designs (for waveguides and cavities, respectively), and the final processing step, undercutting of the sample, was relatively simply accomplished through a hydrofluoric acid wet etch to remove the underlying SiO_2 layer of the SOI material. By the end of February 2003, we were ready to test our devices.

Our plan was to use optical fiber taper waveguides as a method for coupling light into and out of these PC devices, and to study the characteristics of this coupling to understand the device properties. Such coupling had previously been demonstrated to *silica* devices such as microsphere cavities [32, 20], and these works served as an inspiration for our research. However, our devices were significantly different from what had been used in previous demonstrations; we were working with high-index semiconductor microphotonic elements in which the modal fields, particularly in the PC microcavities, were significantly more spatially localized. Overcoming the refractive index mismatch between the fiber taper and the PCWG was at the heart of Paul's work, as they had been designed [33] to have a phase velocity (and hence effective refractive index) equal to that of light propagating through the optical fiber. As a result, it was clear that phase matched coupling should be achievable. When combined with the significant field overlap between the taper and PCWG modes, this would enable efficient power transfer to take place over the length of the waveguide. For the PC cavities, we did not employ any specific technique to overcome the refractive index mismatch. Instead, we were essentially hoping that the field overlap between the fiber taper mode and the cavity mode would be strong enough to enable some amount of coupling between the two devices; efficient power transfer was not our initial goal. Rather, we simply hoped to have enough coupling to be able to use the taper as a means to learn about the modes of these PC cavities.

Quantum cascade photonic crystal microcavity lasers

As I touched on above, while working on the high-*Q* PC cavity project, I also had the opportunity to collaborate with Bell Labs on a project involving the integration of PC cavities with quantum cascade heterostructures. As this project lies somewhat outside the focus of the rest of this thesis, I have chosen to include a summary of the technical details within appendix A. Nevertheless, this work presented an exciting opportunity for us, both in terms of collaborating with a very accomplished group of scientists at Bell Labs and also in terms of the technical potential of the project.

The basic goal of the project was to create quantum cascade, photonic crystal surface-emitting lasers, or QC-PCSELs. These devices were of interest to us for a number of reasons. For quantum cascade lasers, they offered the potential for direct surface emission, a non-trivial property due to the transverse magnetic (TM) polarization of the intersubband transitions in QC lasers, but a highly desirable one for applications such as trace gas sensing (one of the most important applications of QC lasers). In addition, PC cavities offered the promise of device miniaturization and integration, allowing for the realization of multi-wavelength laser arrays on a single chip, again of great potential for sensing applications. For PC microcavities, success would represent the first demonstration of an electrically injected PC microcavity laser, and an important milestone for PC devices.

The PC microcavity design we initially decided to employ was a very simple geometry within the hexagonal lattice, as our initial goal was simply to demonstrate lasing from a QC-PCSEL, without particular regard for how high the Q of the cavity was (of course, the Q needed to be high enough for gain to overcome loss in the cavity). By the time I started working on the project in earnest in August 2002, the main obstacle was not in design but in fabrication of the devices, and in particular, the dry etch of the quantum cascade heterostructure. As mentioned above, by the end of 2002, we had developed a suitable dry etch, and Raffaele and Mariano had worked out all of the other processing steps. Toward the end of 2002 and in early 2003, they began measurements on fabricated devices, and lasing action was observed. After carefully examining the data, it became clear that the cavity modes that we observed in electroluminescence measurements were not due to defect modes, but were rather due to band-edge states (low group velocity modes) in the photonic lattice. From the perspective of what we had set out to demonstrate, this distinction was relatively unimportant as the band-edge modes were still confined within the microcavity structure and had reasonable far-field distributions. However, for future applications, the increased localization and higher Qs of the defect modes are of considerable interest.

To verify our understanding of these devices, the initial electroluminescence data, which showed the tuning of the cavity resonances as a function of the cavity geometry (lattice constant and hole size) as well as lasing for certain modes that were well aligned with the QC material gain, was supplemented by additional measurements of the laser mode's polarization and far-field distribution. These experimental measurements were complemented by FDTD simulations and simple symmetry-based arguments, and by May 2003, we had a good understanding of how the devices were functioning. Our principal results were presented in ref. [34], and subsequent publications presented detailed discussions of device fabrication [35] and lasing mode identification [36].

Around the time of completion of this first round of work, there was a redistribution of the Bell Labs group to different universities (Raffaele went to Université Paris-Sud as a research faculty member, Mariano became a postdoc in Federico Capasso's newly formed research group at Harvard, and Claire Gmachl became a faculty member at Princeton). On our end at Caltech, I continued to work on the cavity design some [37], but otherwise stopped working on the project to focus all of my efforts on the near-IR high-Q microcavities. Raviv Perahia has taken over the project at Caltech, and has been working on developing the capability to do a larger chunk of the fabrication at Caltech (including the electron beam lithography and dielectric mask etching), and on new etch recipes to handle different QC heterostructures. Initial efforts will be focused on achieving room temperature operation of the devices (the first devices operated at $\sim$10 K), through improvement of a number of device characteristics, including the QC material quality, the cavity design and fabrication, and the efficiency of current injection.

First attempts at taper coupling and near-IR photonic crystal microcavity lasers

While I worked with Paul and Oskar on developing the fabrication processes to implement our photonic crystal cavity and waveguide designs, Matt Borselli took the lead on building a station to fabricate optical fiber tapers and a second station for probing microphotonic structures with them. The former essentially consists of a pair of motorized stages that pull on the fiber (in opposite directions) while it is being heated by the flame from a hydrogen torch. The characterization setup consists of a motorized xy stage upon which the sample sits, and a motorized z stage upon which the optical fiber taper is attached. A scanning tunable laser is used as the input to the fiber taper, and the transmission past the microphotonic element is detected by a simple InGaAs photodetector. By February 2003, these two setups were ready to be used in conjunction with our newly fabricated silicon devices.

My initial attempts at probing our PC cavities with fiber tapers were unsuccessful and somewhat discouraging; we were not able to see any spectral characteristics of our cavity modes within the taper's transmission spectrum. After spending about a week testing a number of different devices, we became concerned that the technique was simply ill suited for our application; the coupling between the waveguide mode of the fiber taper and our ultrasmall volume PC cavity mode was just too small for appreciable power transfer to take place. While I contemplated this possibility, Paul took over the testing setup and was almost immediately successful in seeing coupling between his PCWGs and the fiber tapers. His initial results in this area were sent out for publication in late March, 2003 [38], and demonstrated some of the basic principles of this technique. Over the next few months, Paul continued to study taper coupling to PCWGs, and was able to not only demonstrate highly efficient power transfer to the PCWG mode of interest [39], but also the utility of fiber tapers as a probe for the spatial and dispersive properties of PCWGs [40].

Soon after our initial failure at taper coupling to PC cavities, I decided that it was worthwhile to pursue another route to an experimental demonstration of our high-Q designs. One method that had been commonly used for studying the Qs of semiconductor microcavities was the use of an active material to create a light-emitting structure whose emission properties could be studied [41, 13, 8, 14, 12, 42]. I was aware that this method had some limitations, such as the need to pump the structure to material transparency in order to achieve a true estimate of the bare cavity Q, but it was nevertheless an already established technique that would, at the least, give us a strong indication as to whether we were able to create PC cavities with a $Q > 10^4$. Furthermore, active devices such as microcavity lasers were an area of research that we had planned on investigating at some point, and we had in fact already ordered and set up many of the components necessary to do microphotoluminescence measurements. Of course, one drawback was that active devices would require the development of new etch processes. More importantly, as of January 2003, we did not have access to any material from which such devices could be fabricated.

Luckily, in February 2003, Claire Gmachl put us in touch with Dr. Jianxin Chen, a Bell Labs scientist with expertise in the growth of InP-based quantum well structures. Jianxin had in fact already grown 1.3 μm InAsP/InGaAsP multi-quantum-well laser material that was very close in design to what was of interest to us. We agreed to collaborate on this project, and by mid-April, Jianxin had provided us with laser material from which we could make our PC microcavity devices.

Starting in March, I began focusing on InP-based PC microcavity lasers in earnest (chapter 3). The fabrication promised to be challenging, as InP is typically considered to be a much more difficult material to dry etch than Si. Although there are other possible chemistries that can be used, we were committed to an Ar/Cl_2 etch due to its relative cleanliness, which is an absolute necessity for us as we use a single ICP-RIE to etch a number of different material systems with a number of different gas chemistries. While clean, one problem with the Ar/Cl_2 chemistry is that the $InCl_x$ etch byproducts are not volatile at room temperature, so that the material will not etch cleanly unless steps are taken to increase the volatility. The most common way to do this is through heating; in our QC work, this heating was accomplished by the plasma itself. Although the material that we etched in the QC project was very similar to the material we wanted to etch for these near-IR PC microcavity lasers, it quickly became clear that we could not use the same etch. The QC etch is very nonlinear in time, resulting from the finite amount of time (at least a couple of minutes) required for the etch to heat the sample to some sufficiently high temperature. For the QC lasers, this was not a problem, as we had a 500 nm thick SiO_2 etch mask that could withstand the etch for several minutes. Such a thick etch mask was possible due to the mid-IR operating wavelength ($\lambda \sim 8$ μm) of these devices, which allows the use of a chemically amplified photoresist (which is relatively dry etch resistant) as an electron beam resist. For the near-IR ($\lambda \sim 1.3$ μm) PC microcavity lasers, we had to use an electron beam resist with much better resolution (such as ZEP 520A), due to the much smaller feature sizes. These resists are much less dry etch resistant, limiting the thickness of dielectric mask that could be used.

Fortunately, we had another option to heat the sample and increase the $InCl_x$ volatility, which was to directly increase the temperature of the ICP-RIE lower electrode using a resistive heating element that the tool came equipped with. We quickly added the necessary plumbing to allow us to run the ICP-RIE in this heated wafer table mode, which would let us achieve temperatures in excess of 250 °C. Initial tests indicated that temperatures greater than $\sim 150°$C would be sufficient for achieving the volatility necessary to etch the InP cleanly.

Of course, such elevated temperatures precluded any possibility of directly etching into the InP using only the electron beam resist as a mask, as we did with silicon. We decided to use SiO_2 as an etch mask, and for the next several weeks, I spent a considerable amount of time and energy in trying to develop etch recipes for both the SiO_2 mask and for the InP-based heterostructure layers. I was eventually able to develop a reasonable SiO_2 etch, although it was somewhat fickle in its performance, and has remained so to this day. One difficulty in etching so many different materials in one chamber is that it is hard to be certain that you are starting your etch from some well-known initial condition. As a result, it is very important for us to develop relatively robust processes. Nevertheless, the SiO_2 etch seemed to be at least adequate, and the subsequent InP etch did not take long to optimize once we committed to using elevated temperatures. By early May 2003, all of the steps, including the subsequent wet etching steps to release the PC membranes (which are somewhat more complicated for these structures than the SOI ones), had been worked out to make the lasers. Chapter 3 discusses the development of these fabrication procedures in detail.

Oskar and I tested the initial rounds of devices together, and we were immediately rewarded with lasing action in photoluminescence measurements of the first devices we tested. Over the next couple of weeks, we tested many devices, and were able to measure cavity Qs of 1.3×10^4 at material transparency. This principal result was very exciting, and represented about a 5 times increase in Q over what had been demonstrated for a PC microcavity to date [12]. In addition, we were encouraged to see that our measured Qs were essentially limited by the resolution of the spectrometer, and furthermore, were very reproducible from device to device. These results, along with studies of the lasing mode polarization, light-in-light-out characteristic, and approximate mode localization, were submitted for publication in late May, and were published in September [43]. In addition, we wrote a detailed article describing the fabrication of these devices [44], which, along with our new cavity designs, was clearly instrumental in their superior performance. This point was particularly apparent to us as Oskar's own graduate research involved fabrication of PC microcavity lasers in InP-based multi-quantum-wells, and the discrepancy between our new fabrication processes and the fabrication processes he had used (which employed a chemically-assisted ion beam etch) was quite pronounced.

We initially had many plans for continuing work on these InP multi-quantum-well lasers, including studying laser performance as a function of the number of quantum wells in the device, developing surface passivation methods to reduce non-radiative recombination, and starting a program to investigate some interesting phenomena in microcavity physics (such as enhanced radiative rates due to the Purcell effect or measurements of the spontaneous emission coupling factor). These plans were never followed up on, however. Initially, we were more interested in giving the taper coupling approach another try, and in the long term, quantum-dot-based materials seemed to be a more suitable choice for some of the microcavity physics experiments. As a result, we stopped efforts of InP-based near-IR PC lasers after this initial round of success. Nevertheless, the Qs that we demonstrated remain the highest values that have been demonstrated for an InP-based PC microcavity (to the best of my knowledge), and quantum-well-based lasers do have potential for various applications, such as fluid-based sensing [45, 46, 47] where the additional quantum confinement provided by quantum dots is not needed. Furthermore, the fabrication processes developed can certainly be used in combination with structures incorporating InP-based quantum dots in an InP-based matrix [48].

Fiber taper probing of silicon photonic crystal microcavities

After completing work on the high-Q PC microcavity lasers in June 2003, we returned the idea of using the fiber tapers to investigate the cavities. We felt that there were many good technical reasons for doing so. In particular, Q measurements involving emission from a light-emitting material are complicated by the need to pump the material to transparency; in practice, the pump level at which this condition is achieved can be hard to experimentally identify. In addition, these emission-based measurements are limited by the resolution of the spectrometer used; for us, this was ~ 0.1 nm, but even for instruments with a longer path length, achieving resolutions better than 0.01 nm can be difficult. Other resolving instruments, such as Fabry-Perot optical filters could possibly be used, but would rely upon collecting a significant amount of emitted power from the PC cavities, a non-trivial feat. A strictly passive measurement, which we proposed, would instead probe the cavity transmission (or reflection) through the fiber taper as a function of the wavelength of the input probe light. The wavelength resolution in such a measurement would then be limited by the probe laser's linewidth, which could be less than 10 MHz ($<$ 0.1 pm). Passive measurements of PC microcavities had certainly been done in the past [49], in experiments where an in-plane waveguide had been fabricated to couple light into and out of the cavity. However, such experiments had their own difficulties; coupling light into a PCWG is itself non-trivial, and furthermore, somewhat inflexible in that a waveguide must be fabricated for each cavity on the chip, and the waveguide position, and therefore the waveguide-cavity coupling, is fixed and cannot be adjusted. On the other hand, an optical fiber taper waveguide can be fabricated with very little loss ($<$ 10%), and a single device could be used to probe all of the PC cavities on a chip, in a rapid and flexible manner. In that sense, we envisioned the taper acting as an optical probe for chip-based resonant microphotonic elements.

Furthermore, Paul's experiences in coupling between tapers and PCWGs convinced us that our initial failures in coupling to cavities were not fundamental, but were rather due to simple technical difficulties that could easily be remedied. In particular, he discovered some problems with how our detectors were set up, which would have prevented us from measuring modest changes in the transmitted signal. As the taper-PC cavity coupling was not expected to be particularly strong (unlike the taper-PCWG coupling), this was a vitally important discovery. At the time of our previous attempt at taper-PC cavity coupling, we also had not properly calibrated our SEM, so that our measurements of the dimensions of our cavity (such as the hole radius and lattice spacing) were overestimated by around 10%. As the scanning range of the laser source used in the taper probing measurements was around 5% of the center wavelength (λ=1595 nm), it was certainly important to know the cavity geometry to a reasonably high degree of accuracy.

Thus, when we began our second attempt at taper coupling to the PC cavities (chapter 4), we were much better prepared. We were immediately able to see some amount of coupling between a taper and the first set of new SOI PC cavities that we fabricated, and were able to confirm that the modes were localized defect states. In particular, we were able to demonstrate coupling to the high-Q mode of interest from our design work (chapter 2), and by July 2003, we had used the taper probing technique to measure cold-cavity Qs as high as 25,000. This measurement technique worked extremely well; we could easily probe all of the devices on a chip with a single taper, and the cold-cavity Q could be accurately determined by measuring the linewidth of the cavity resonance (within the taper's transmission spectrum) when the taper was positioned several hundred nanometers away from the cavity (to reduce taper loading effects).

In addition to using the taper to determine the resonant wavelength and Q of the cavity modes, we also began to consider its use as a tool to study the spatial properties of the modes. In particular, by varying the taper's position with respect to the cavity, the position-dependent coupling could be ascertained. This coupling was clearly a function of the overlap between the taper and cavity fields, and would thus, at some level, describe the spatial localization of the cavity mode. We fabricated new devices that could be probed along both axes of the cavity, and took measurements of this position-dependent coupling as the taper was moved along each of these axes. This data was then compared to the results of a simple coupled mode theory [50] that took into account the analytically determined taper field and FDTD-calculated cavity field, and the experimental data and numerical results matched very well. In early September 2003, we submitted a paper detailing the taper-based measurements of a Si PC microcavity with an experimentally measured $Q \sim 40,000$ and spatial localization consistent with $V_{\text{eff}} \sim 0.9(\lambda/n)^3$ to the journal *Nature*. Although the Qs that we measured were only a factor of 3 times larger than what we had demonstrated in the multi-quantum-well laser cavities, we felt that the taper probing technique, which confirmed the simultaneous demonstration of high-Q and ultrasmall V_{eff} in our PC microcavities, added a significant amount of content to the work. The demonstrated Q and V_{eff} values were particularly exciting from the perspective of cQED, where they would satisfy the strong coupling requirements for both neutral alkali atoms and self-assembled semiconductor quantum dots, with the potential for coupling rates on the order of tens of GHz, which would exceed both the cavity and atom (quantum dot) decay rates.

About three weeks later, we learned that *Nature* had just decided to accept another paper, by Susumu Noda's group at Kyoto University, Japan, detailing the experimental demonstration of high Qs in PC microcavities, and therefore would not consider publication of our work. The Kyoto group work [51], which was published at the end of October 2003, showed measurements of Qs as high as 45,000 in an SOI PC microcavity, using the in-plane waveguide coupling approach. Our rejected manuscript was submitted to *Physical Review* in September 2003, and was eventually published as a *Rapid Communication* in *Physical Review B* in August 2004 [52]

On the surface, the Kyoto work and our work seem to be very similar, but I have always felt that within the context of PCs and optical microcavities, there are several important differences. As I have already mentioned, our work was partly focused on the application of fiber taper probing to wavelength-scale semiconductor microcavities, and in demonstrating that the fiber taper could serve as a useful tool for studying both the spectral and spatial properties of the cavity modes (in particular, measuring both Q *and* V_{eff}). In addition, although both designs employed Fourier space methods to reduce vertical radiation loss from the cavities, there were some important differences. Upon glancing upon the geometry of the cavities, the most glaring difference would seem to be in the complexity of the design; the Kyoto group design simply consists of three missing holes and two shifted holes within a hexagonal lattice PC. In contrast, our design incorporated a graded square lattice in which the hole radius continuously changed as a function of position within the lattice. While we did use this to achieve our highest *Q*s, it was not completely essential for obtaining reasonably high *Q*s. We were previously able to show, in our simulation work, that *Q*s of ∼40,000 could be achieved in simple square lattice cavities in which the defect consisted of the reduction in size of two air holes in the lattice. In addition to increasing this maximum achievable *Q* to $\sim 10^5$, part of the reason for incorporating the grade in hole radius was to make the design robust to fabrication imperfections, so that *Q*s in excess of 10^4 could be achieved even when the cavity significantly deviated from the prescribed geometry. On the other hand, the Kyoto group design (as well as other designs involving the fine tuning of a small number of holes in the PC lattice) required very precise control of the cavity geometry to achieve high *Q*s; small variations in the geometry could easily reduce the *Q* by a factor of 10.

We spent the next couple of months exploring this idea more quantitatively, fabricating and testing different devices and comparing the results to simulations. We amassed the data from several devices, and were able to conclude that our design was indeed robust to fabrication imperfections, from both an experimental and a theoretical standpoint. In particular, we were able to show that the cavity *Q* could remain above 20,000 even in the presence of variations in the cavity geometry that caused the resonant frequency to vary by ∼ 10% of its nominal value. These geometrical variations included perturbations that broke the desired symmetry of the structure, as well as variations in the hole radius that were significantly different from the intended grade in hole radius. As described in chapter 4, there are relatively simple physical explanations for this robustness, which I believe can be incorporated as principles in future designs of high-*Q* PC microcavities. We published some of these ideas and numerical and experimental data on the robustness of our cavities in an article in *Optics Express* in April 2004 [53]. Although robustness has not necessarily gained widespread recognition as an important element of PC cavity designs, I think that eventually it will, particularly as optimization and inverse design techniques [31, 29] become more commonly used, as robustness can possibly be entered as a condition within such algorithms. Certainly, from the perspective of an experimentalist, I can attest to the practical utility of having every device on the chip support a high-*Q* mode, rather than an isolated number of devices that precisely match the nominal design.

At this point, although we had been able to show the versatility of the fiber taper as a probe for the spectral and spatial properties of PC microcavities, as described in detail in a review article [54], one important thing that we had not yet demonstrated was highly efficient coupling to the cavities. Our results thus far had been limited to maximum coupling depths of around 10%-20%. On the one hand, such levels of coupling were fairly high because the taper is such a low-loss device; coupling losses to an in-plane waveguide using standard end-fire coupling approaches can easily be in excess of 3 dB per coupling junction (for a total of 6 dB, or 75 %). On the other hand, phase matched coupling, such as that which had been achieved for silica microsphere and microtoroid cavities [20, 55, 56], and by Paul for PCWGs [38, 39, 40], could achieve coupling efficiencies approaching 100 %. There was thus significant room for improvement.

The strategy we adopted, in work led by Paul, was to first couple light from the taper to the PCWG, which we had shown could be very efficient, and then couple light from the PCWG to the PC cavity. This had been our plan from the outset, as the PCWGs had been specifically designed to mode match to our PC cavities [33]. In addition to potentially allowing for very high coupling efficiencies, this technique had the advantage of more optimally loading the cavity than direct taper probing does, so that the loaded Qs can typically be much higher. Paul's results [57] showed a loaded Q of 3.8×10^4 for $\sim$ 44% coupling efficiency to a cavity with an unloaded Q of 4.7×10^4. In comparison, for direct taper coupling, the loaded Q was $\sim 2.0\times10^4$ for a coupling depth of about 10 % to a cavity with an unloaded Q of 4.0×10^4. Having achieved these impressive results by May 2004, Paul then worked on an initial demonstration of what could be done with these fiber-coupled PC cavities by studying the nonlinear response of these devices. He was able to show evidence of optical bistability at very low input powers of a couple hundred microWatts, corresponding to femtoJoule internal cavity energies. These results were published in February 2005 [57], and were soon followed by several additional studies of nonlinear optics in silicon PC microcavities (see ref. [58] and its follow-up articles, for example).

Since the initial publications on high-Q PC cavities from our group [43, 52, 53] and the Kyoto group [51], there has continued to be impressive progress in the experimental demonstration of higher and higher Q PC cavities. The Kyoto group was able to fine tune their initial design to achieve Qs of $\sim 10^5$ [59], and more recently, have employed an entirely different design to achieve $Q \sim 6.0\times10^5$. This very interesting and impressive result relies upon the slight modification (through adjustment of the lattice constant) of a PCWG mode to form a cavity that has very little vertical radiation loss, and still has $V_{\text{eff}} \sim$1-2$(\lambda/n)^3$ [28]. Even more recently, the group of Masaya Notomi at NTT Laboratories has used a similar design (this time adjusting the size and position of some of the PCWG holes) to achieve $Q \sim 9.0\times10^5$ [30]. One key to these designs is that the PCWG modes, in principle, have no vertical radiation loss, so that they serve as a good starting point from which the cavity design can be built. Another key is that strong localization in these PC lattices can be achieved through very weak perturbations to the lattice, which allows for the creation of defect modes that do not contain lossy Fourier components (such as those within the cladding light cone). In some ways, these results mirror, and of course amplify, what we have been able to see in our graded lattice PC designs, particularly in regards to the idea that a seemingly weak modification to the PC lattice can still confine light to an ultrasmall volume, which was perhaps not appreciated in the earliest PC cavity work. This can be seen in the evolution of PC cavity designs, starting from Oskar's original work [7, 8] employing a single missing air hole in a hexagonal PC lattice, to some of our designs that graded the lattice hole radius, to these latest designs that utilize small modifications to a PCWG.

Switching to AlGaAs: Even more fabrication

As soon as we achieved our first results on passive probing of high-Q Si PC microcavities in June 2003, we began to consider their potential use in cavity QED experiments. There were two specific experiments of interest to us, involving trapped neutral atoms and self-assembled quantum dots, respectively. Of course, our group does not have expertise in either atomic physics or material growth, so both projects would require a collaboration with other groups. At Caltech, we are fortunate to have two research groups in the physics department, those of Professors Hideo Mabuchi and H. Jeff Kimble, that are world leaders in atomic physics and its application to cQED. Hideo's group, in particular, had already began to consider the potential for incorporating PC cavities with magnetostatic atom chips to create integrated devices in a scalable architecture [60]. Hideo's student, Benjamin Lev, had already developed the ability to trap and guide neutral Cs atoms on microfabricated atom chips [61]. However, despite some initial design work [31], their group had not yet experimentally demonstrated high-Qs in their PC microcavities, and perhaps more importantly, had not been able to effectively address the issue of coupling light into and out of the devices.

In early July 2003, Paul, Oskar, and I met with Ben and Hideo to discuss the potential for a collaboration, and we agreed that it was something of interest to all parties. Initially, there was not that much that we could do on our end, as Paul and I still had much work to do in establishing the efficacy of the taper probing for PC microcavities. Nevertheless, it was useful for us to meet with Ben on a periodic basis, to learn about his atom trapping setup, and to understand what some of the biggest obstacles in this experiment would be. The basic proposal and some simple simulations for the experiment that we planned to do, optical detection of single atom transits through a PC microcavity, were put together in article that we coauthored with Ben and was published in July 2004 [62].

It was immediately apparent that the one issue that was most pressing on our end was the demonstration of a high-Q, ultrasmall V_{eff} at the near-visible wavelengths at which atomic cavity QED experiments are conducted. The Si microcavities we had developed would not be option, as Si is opaque at wavelengths below 1 μm. InP-based structures were a possibility, but we instead decided to go with GaAs-based devices, as that would be the most likely host material for any future experiments with semiconductor quantum dots (QDs). For the atomic physics experiments, pure GaAs was not a possibility, as it too was opaque at the Cs transition wavelength of interest (the D2 transition at $\lambda \sim 852.3$ nm). Instead, we would use an AlGaAs structure (with an Al percentage of $\sim 30\%$) which in principle, would be transparent at 852 nm. As the refractive index of the AlGaAs structure would be very close to that of Si, the PC cavity design would not change, with the exception of the proper scaling for operating at shorter wavelengths. The fabrication, however, would be completely different. Although optimization of device fabrication in yet another material system was not a very attractive project for me, it was certainly a necessity for us to be able to move forward with cavity QED experiments. Indeed, we placed more importance on this than on any efforts to improve on the performance of our Si PC microcavities.

After taking my candidacy examination in February 2004, I began calibration of etch processes for AlGaAs-based microcavities. By this point, we had begun a collaboration with Professors Andreas Stintz and Sanjay Krishna at the Center for High Technology Materials at the University of New Mexico. Andreas and Sanjay are experts at material growth and in particular, have a good deal of experience in growing self-assembled QDs for applications such as lasers and hyperspectral detectors. They had agreed to provide us with both pure $Al_{0.3}Ga_{0.7}As$ waveguides for the Cs cavity QED experiments and QD-containing waveguides for the creation of microcavity lasers and eventual studies of single photon, single QD interactions.

In a very qualitative sense, AlGaAs tends to etch somewhat less easily than Si but more easily than InP. As a result, it was not clear as to whether we would have to use a dielectric mask in our processing, or if we could transfer the PC pattern directly from the electron beam resist as we had done in the Si devices. Other groups had used the direct transfer approach, but the air holes used in their devices were significantly larger than what ours would be, so that achieving a vertical etch for our devices would be somewhat more difficult. Furthermore, the demonstrated Qs for AlGaAs PC microcavities had been relatively modest [12, 63] (a few thousand at best), and in order to improve upon this, we wanted to optimize the fabrication processes as much as possible, even if this meant resorting to the added steps required to use a dielectric etch mask.

As I discuss in appendix C, we eventually settled on a process that made use of a SiN_x dielectric etch mask for the subsequent dry etch of the AlGaAs layer. The key here is that the SiN_x can be etched using a recipe similar to that we used for etching Si. This low power etch minimally damages the electron beam resist, allowing for the initial cavity pattern to be faithfully reproduced in the dielectric layer. By August 2004, we were in a position to start testing devices.

Over the previous several months, Matt had begun research on Si microdisk cavities. This was a natural follow-up to our Si PC cavity work in many ways. The etch processes needed to fabricate the devices would be identical to what we had already developed, and the devices could be tested using the fiber taper waveguides. At that point the highest demonstrated Qs for small mode volume devices were on the order of 10^4 [13, 14].

Matt's results, which were published in October 2004 [64], were truly impressive; he was able to measure Qs as high as $\sim 5\times10^5$ for a mode volume of $\sim 6(\lambda/n)^3$. These results, in terms of the metrics Q/V_{eff} and $Q/\sqrt{V_{\text{eff}}}$, which are relevant to processes such as enhanced spontaneous emission and strong atom-photon coupling, were actually better than what we (or any other group) had achieved in PC cavities. For applications requiring the introduction of another material to couple to the optical field, such as cavity QED with neutral atoms or colloidal quantum dots, the microdisk geometry was not optimal, as its peak electric field lies buried within the semiconductor layer. However, for experiments involving embedded materials, such as semiconductor quantum dots (which was to be my primary focus), they are completely suitable. Furthermore, one particularly attractive aspect of Matt's results was that coupling depths of $\sim 50\%$ could be achieved while still maintaining a loaded $Q > 10^5$. Thus, not only could the fiber taper be used to probe the Q and spatial localization of the microdisk modes, it could serve as an efficient coupler as well. Although Paul had already established the method for efficiently coupling to our high-Q PC microcavities by this point, the simplicity of being able to directly couple to the microdisk while still maintaining a high-Q was certainly appealing.

In my estimation, the microdisks seemed to be a good starting point for future experiments. For much of the work we intended to do, including microcavity-quantum-dot lasers and chip-based cavity QED involving single quantum dot, single photon interactions, the cavity geometry is by itself not necessarily important; all of the important device properties are encapsulated by the cavity's Q, V_{eff} (more precisely, the peak electric field strength at the location of the emitter), and η_0, a parameter that defines the collection efficiency of photons from the cavity. Matt's results had clearly shown that the microdisks were very competitive in these regards. In addition, from a purely practical standpoint, high quality microdisks are in many ways much easier to fabricate than PC cavities (in terms of the requirements on sidewall angle, as one example). Thus, when starting new experiments, it seemed like an appropriate strategy was to first use microdisk cavities as a way to observe the basic physical phenomena of interest. Ultimately, PC cavities have a number of interesting properties that are somewhat unique to them (for example, very high spontaneous emission coupling factors), and they are certainly interesting from the perspective of future experiments involving integration of multiple devices on a chip, where the truly planar geometry of the PCs is a significant advantage. We thus planned to begin experiments by using microdisk cavities, and would switch to PC cavities when many of the experimental details had been worked out, or when we were interested in studying phenomena that would be specific to them.

Quantum dot microdisk lasers

Fabrication of the AlGaAs microdisk cavities was accomplished through largely the same fabrication steps as what we had developed for the PC cavities. The primary differences were slight adjustments in the etch chemistries for the SiN_x and AlGaAs dry etches, to make the disk sidewalls as smooth as possible, even at the expense of sidewall verticality. The reason for this is that for most microdisk geometries, sidewall angle is not nearly as important an influence on Q as it is for PC cavities, but sidewall roughness is. Another important element of the microdisk fabrication is a resist reflow process that Tom had developed for the Si microdisk cavities. This process largely eliminates radial variations in the mask, thereby improving the circularity of the disk and reducing surface scattering losses. Matt and Tom had recently used this technique to demonstrate Qs as high as 5×10^6 in a 60 μm radius disk, and they also made devices that exhibited $Q \sim 1.25\times10^6$ for $V_{\text{eff}} \sim 14(\lambda/n)^3$ [65].

The first devices I tested were microdisks fabricated in a passive $Al_{0.3}Ga_{0.7}As$ waveguide (no quantum dots), for potential use in the neutral Cs atom cavity QED experiments. At the desired operation wavelength of $\lambda \sim 852$ nm, the devices did not perform as well as hoped, and we did not observe Qs greater than $\sim 2\times10^4$. On the other hand, the devices performed much better at longer wavelengths, with $Q \sim 1.5\times10^5$ exhibited in the 1500 nm band. Over the next several months, this behavior was observed repeatedly, both by me and by Paul, who took over our group's portion of the neutral atom cavity QED project upon completion of his work on nonlinear optics in Si PC cavities [57]. Our working hypothesis is that deep impurity states (caused by O_2 incorporation, for example) in the AlGaAs are causing absorption at wavelengths above the semiconductor bandgap (around 700 nm for $Al_{0.3}Ga_{0.7}As$) but below 1 μm. There is some support for this theory in literature, through measurements of waveguide loss [66]. We are currently working with Kevin Hennessy in the group of Professor Evelyn Hu at UC Santa Barbara to study this phenomenon a bit more closely, with Chris Michael taking the lead on this project from our end. If true, this will have a significant impact on QD-microcavity devices, as much of the current work is being done at wavelengths below 1 μm.

Our measurements within the 1500 nm band were quite promising as far as our QD-based devices were concerned, as they were grown by Andreas and Sanjay to operate in the 1.2-1.3 μm band. The most immediate application for these devices would be room-temperature lasers. We began this investigation, described in chapter 5, by fabricating disks in an epitaxy containing three stacked layers of quantum dots. Such an epitaxy (often times containing even more layers of dots) is often used in QD laser research as a result of the increased amount of gain relative to a single layer of QDs. We were quickly able to demonstrate optically pumped lasing from such devices, and the thresholds were quite low, close to the predicted transparency level for the material. These low thresholds were consistent with the cavities sustaining very high Qs. Indeed for a cavity $Q \gtrsim 3{\times}10^4$, the modal gain from a single layer of dots was predicted to be sufficient to achieve lasing. This would be a significant demonstration, as at that point, no other microcavity-based device had been able to demonstrate room temperature lasing from a single layer of QDs (due to the relatively high Qs required).

We were soon able to demonstrate optically pumped lasing from a single layer of QDs, at room temperature, for both pulsed and continuous wave pumping. The threshold pump values were again quite close to the predicted transparency levels, indicating that the cavity Qs were most likely much larger than the minimum value ($Q \sim 3{\times}10^4$) needed for lasing. As this value is already higher than what could be resolved with our spectrometer (which has a resolution limit of around 0.1 nm), this would typically be all that could be quantitatively said about the cavity Q using standard measurement techniques. There is the possibility of fitting the lasing data with a rate equation model with the Q as a free parameter, but such fits depend upon a number of different parameters, such as the QD radiative lifetime, internal efficiency, and collection efficiency, which will not necessarily be precisely known, and the estimated Qs will have a relatively large uncertainty.

Using the fiber taper probing technique in conjunction with knowledge about the microdisk mode structure allowed us to learn a great deal more, however. In particular, we used the taper to probe the cavity modes within the 1.4 μm wavelength band, which was far enough red detuned from the emission wavelength that we did not expect to see any effects related to the absorption by the QD gain material, thus giving us an accurate estimate of the cold-cavity Qs of the modes. This would give us the desired information about the quality of our fabrication, and would let us know whether our cavities were good enough for future experiments, such as those in cQED. It would also give us additional information about the lasing modes we had studied through photoluminescence. Of course, a mode that we study in the 1.4 μm band is not the same mode that is lasing. However, by carefully studying the cavity's spectrum and comparing it to simulation results, we can study a mode within the 1.4 μm band that has the same polarization and vertical and radial order as the lasing mode, and only differs in its azimuthal mode number (i.e., the number of lobes in the azimuthal direction). As long as the radiation loss that occurs when the mode tunnels around the disk perimeter is not the dominant loss mechanism, modes whose azimuthal orders are only slight different will behave very similarly; in fact, modes at the longer wavelength will have lower radiation-limited Qs. As losses due to mechanisms such as surface scattering, surface absorption, and bulk absorption are not expected to be drastically different between the 1.2 μm and 1.4 μm bands, the Q measured in the 1.4 μm band can give a good estimate of the Q expected for the lasing mode in the 1.2 μm band.

Using this taper coupling technique, we were able to demonstrate Qs as high as 3.6×10^5. These Qs were over a factor of 10 higher than anything that had been demonstrated in an AlGaAs microcavity. The high-Q resonances we saw were actually doublets, corresponding to standing wave modes that form when the traditional clockwise and counterclockwise modes of the disk are coupled and split by surface scattering. Such doublet modes were also seen by Matt in his investigation of Si disks [64], and are indicative of a cavity loss rate that is low enough that coherent coupling between the propagating modes of the disk can occur. For our disk geometries (255 nm thick and 4.5 μm in diameter), $V_{\text{eff}} \sim 6(\lambda/n)^3$. For optimal coupling to a single quantum dot, this corresponded to a QD-cavity coupling rate $g/2\pi \sim 11$ GHz, which would greatly exceed (by over a factor of 10) the demonstrated cavity decay rate $\kappa/2\pi \sim 0.4$ GHz and typical QD decay rates of $\gamma_{\perp}/2\pi \sim 1$ GHz. These results indicated the potential for strong coupling in this system.

In addition to the V_{eff} and cold-cavity Q of these cavities, we were also able to demonstrate relatively efficient coupling to them, with power transfers very similar to what Matt had demonstrated in the Si microdisks. We measured coupling depths as great as 60 % for a loaded $Q \sim 10^5$. This was a key result; it indicated that we could still achieve strong coupling (with $g/\kappa \sim 10$) while also obtaining efficient coupling into and out of a coupled QD-microdisk system. For future experiments, this will be of great importance. Experiments in quantum optics will involve very low light levels, and efficient light collection is necessary from a detection standpoint, particularly at wavelengths greater than 1 μm, where InGaAs detectors, which suffer from significantly poorer performance characteristics than the Si detectors used at shorter wavelengths, are used. From the standpoint of future work in quantum information processing, applications in linear optics quantum computing [67, 68] require a near-unity collection efficiency of single photon pulses. Although not currently at that level, our results were a significant step towards these goals, and the path to more efficient coupling through tailoring of the disk geometry was apparent. Indeed, Matt and Tom have since demonstrated critical coupling in the Si microdisks.

This first set of results on optical loss and lasing characteristics in AlGaAs microdisk cavities with embedded quantum dots was submitted for publication in December 2004, and was published in April 2005 [69]. At this point, we began to seriously consider the possibility of doing cavity QED experiments in these resonators. By this point, vacuum Rabi splitting due to the interaction of a single QD with a microcavity had just been demonstrated by three groups [70, 71, 72], through measurements of the spontaneous emission from the devices. Although our devices had the potential for exhibiting Rabi splitting with a greater ratio of g/κ, it was clear to us that the most important contribution that we could make was to develop cavity QED experiments using the fiber taper as an input-output channel, as it appeared to be the key to opening up a number of future experiments and device applications. This would require us to set up a number of different pieces of equipment in the lab, such as a spectrometer with a very sensitive detector for studying the emission from single QDs, a laser source for doing near-resonant pumping and spectroscopy of the coupled QD-microcavity system, and the integration of the fiber-coupled devices in a liquid He cryostat. The last point was particulary important and challenging, as it had not yet been demonstrated and was essentially at the heart of our proposed experiments.

While waiting for some of the experimental equipment to arrive, we decided to further investigate the performance of our microdisk-QD lasers at room temperature (chapter 6). In particular, we were interested in explicitly demonstrating the utility of the fiber taper for future measurements of QD-containing microdisks. To this point, we had used the taper as means to probe the microdisk Qs, but there obviously was significant potential for using the taper as means to collect the emission from optically pumped devices, and furthermore, as a way to pump the devices as well (to this point, our pumping and collection was done through standard free-space optics). To most explicitly compare the efficiency of fiber taper collection with free-space collection, we revamped our photoluminescence measurement setup to allow for the microdisks to simultaneously be probed with optical fiber tapers while still allowing for standard free-space pumping and collection. This allowed us to continue to pump the devices with a free-space beam while comparing the amount of collected power through the free-space optics with that collected through the fiber taper. The difference was quite dramatic; whereas before we were only able to collect hundreds of picoWatts from our lasers, we were now able to collect as much as a couple hundred nanoWatts from the devices. To most precisely quantify our results, we compared the differential efficiency ξ of the laser with and without the fiber taper collection, and saw that the taper collection improved ξ by almost two orders of magnitude. By collecting a significant fraction of the laser emission, we put ourselves in the position to more quantitatively study some aspects of the laser's behavior, particularly subthreshold, where the poor collection typically obtained through free-space optics makes careful investigations of device behavior difficult.

Not only is the collected power significantly greater when using the fiber tapers, but the number and variety of modes that are observed is strikingly different. This is not too surprising; for microdisks, free-space collection essentially relies upon imperfections in the disk to scatter light vertically into our collection optics. This is a relatively inefficient process, and becomes increasingly more inefficient as device fabrication improves and higher and higher Q devices are realized. The fiber taper, on the other hand, directly evanescently couples light out of the disk WGMs, and can be a very efficient process. Although loading due to the taper does degrade the cavity Q, the key point is that this loading is efficient in that the added loss is primarily due to coupling into an observable channel (the fiber taper); the lost photons are collected into the fiber and are free to interact with some other part of the system, such as another cavity or photon counting detectors. This is not the same as loss due to absorption, for example, where the lost photons are not detectable and are of no benefit. This point is somewhat obvious in that it is true for all cavities; in a simple Fabry-Perot cavity, for example, one typically degrades the reflectivity of one of the mirrors to allow for light to be coupled out of the cavity. Despite the simplicity of this concept, it is nevertheless often misunderstood in the literature.

In the early part of 2005, we gained access to two new laser sources, which allowed us to do two additional measurements. The first laser source was a 980 nm external cavity laser that had been built by some members of our lab; we were able to use this as a source for optically pumping our microdisks directly through the fiber taper. We were thus able to make purely fiber-coupled devices that no longer relied upon any free-space optics. From a technological standpoint, the device simply consists of a fiber input, to be hooked up to an optical pump source, and a fiber output, through which the emitted laser light would propagate. Such purely fiber-coupled devices had been previously demonstrated by the Vahala group at Caltech in the context of erbium-doped glass microsphere and microtoroid cavities [73, 74]; our work extended this to the regime of wavelength-scale, semiconductor-based devices, for which interesting gain materials such as quantum well and quantum dot layers can be incorporated. The second laser source we acquired was a 1250 nm scanning tunable laser. The primary purpose of this laser was for future experiments employing near-resonant pumping of QDs at cryogenic temperatures (Andreas and Sanjay had recently grown new samples for us that emit at ~ 1.32 μm at room temperature and $\sim$1.25 μm at 4 K). For the time being, we were able to use it in conjunction with our current devices (which emitted at $\sim$1.2 μm) to measure the cavity Q at a wavelength that was much closer to the peak of the QD emission spectrum. Because the wavelength dependence of surface state and material absorption in AlGaAs is not particularly well known, and in addition, we had already seen strong differences in the cavity Qs between the 850-1000 nm band and the 1400 nm band, it was important for us to be sure that the Qs would still be high at wavelengths relevant to future cavity QED experiments. We were indeed able to measure $Q > 2\times10^5$ at $\sim$1250 nm, with the discrepancy in Q between the 1400 nm band and the 1250 nm band likely being due to absorption from the tail end of the inhomogeneous QD spectrum.

We wrote a relatively detailed paper on the various results concerning photoluminescence measurements using optical fiber tapers and submitted it for publication in June 2005; it was published a few months later in November [75]. We subsequently continued our study of these QD-microdisk devices by considering how device performance (cavity Q, V_{eff}, and lasing threshold) scales in going to smaller sized disks. In particular, we were interested in seeing how these microdisks could compare to devices such as photonic crystal microcavities, both in theory and in practice. On the theoretical end, we employed a finite element eigenfrequency solver method to calculate the radiation-limited quality factor Q_{rad} and mode volume V_{eff} of the first order radial modes of D=1.5-3 μm diameter microdisks in the 1200-1400 nm band. These simulations had been developed using the FEMLAB software by Matt, who had in turn received some assistance from Sean Spillane (a graduate student in Professor Vahala's group who had previously used FEMLAB to study microtoroid resonators). From the simulations, we were able to see that standing wave modes in the microdisks could have $Q_{rad} > 10^5$ for $V_{\text{eff}} \sim 2(\lambda/n)^3$, corresponding to a disk diameter of $\sim$1.5 μm. Furthermore, Q_{rad} sharply increases as the diameter increases, so that $Q_{rad} > 10^8$ can be achieved for $V_{\text{eff}} \sim 3.5(\lambda/n)^3$. Such values are quite comparable to the highest values predicted for 2D photonic crystal microcavities [28, 30]. From these simulation results, we could then calculate the predicted coupling and decay rates for interactions with a single QD. One significant difference between QD-based cavity QED and atomic-based cavity QED is in the emitter decay rates; a typical QD might have a radiative decay rate $\gamma_{sp}/2\pi \sim$0.2 GHz, but due to non-radiative dephasing, the total decay rate $\gamma_{\perp}/2\pi$ will often be closer to 1 GHz (a neutral Cs atom, on the other hand, has a transverse decay rate $\gamma_{\perp}/2\pi \sim 2.6$ MHz). An equivalent cavity decay rate would correspond to $Q \sim 10^5$; this indicates that there is some optimum diameter at which the ratio of the QD-microcavity coupling rate g to the maximum decay rate in the system (either $\gamma_{\perp}$ or the cavity decay rate κ) is maximum. From our simulations, this diameter is $\sim$1.5-2 μm.

We then fabricated 2 μm diameter microdisks and studied their Qs and lasing properties using the taper coupling techniques we had already established. These devices were fabricated in the latest material that Andreas and Sanjay had grown, which had been optimized for room temperature emission at $\sim$1.32 μm and low temperature (4K) emission at $\sim$1.25 μm. By taper testing the devices in the 1.4 μm band, we were able to demonstrate Qs as high as 1.2×10^5 for $V_{\text{eff}} \sim 2.2(\lambda/n)^3$. We also observed room-temperature, continuous-wave lasers with threshold pump powers as low as 1 μW of absorbed power, and used the fiber taper to create devices with differential collection efficiencies as high as 16 %. From a pure device performance standpoint, these devices were outstanding; the demonstrated lasing thresholds were orders of magnitude smaller than what had been demonstrated by other groups for microcavity-QD lasers. This is a result of the ultrasmall volume of these cavities and the correspondingly small number of QDs that need to be inverted (all the more so because we are only using a single layer of QDs); the ability to achieve lasing from such a small amount of gain is made possible by having a high cavity Q. Our study of these small diameter microdisks with embedded quantum dots was submitted for publication in November 2005, and published in February 2006 [76].

Towards cavity QED experiments

The most recent devices we had fabricated showed significant promise for cavity QED experiments, and over the past few months, we have begun to equip our lab to do such work. This includes setting up a continuous flow liquid He cryostat, which would allow us to reduce our sample temperatures to $\sim$4 K, a necessity in order to reduce non-radiative dephasing in the QDs, as well as a 0.5 m spectrometer system with a liquid-nitrogen-cooled InGaAs detector. Even with such a detector, single QD detection at wavelengths above 1 μm is very difficult, due to the already-mentioned poorer performance of InGaAs detectors relative to Si detectors.

With this equipment in place, we are in principle in the position to observe vacuum Rabi splitting in the spontaneous emission from a single QD in a semiconductor cavity, as has been recently demonstrated by three groups at the end of 2004 [70, 71, 72]. This is not to say that such a demonstration will be trivial; we certainly have to work hard to optimize our photoluminescence setup to be sure that we can observe single QD emission, but the techniques for doing so have been established by many research groups, although fewer have done so at the longer wavelengths at which our devices operate. We are most interested in performing a slightly different experiment, however, where Rabi splitting will be observed in the transmission (or reflection) past the cavity. This will be a unique experiment as a result of our capability to efficiently couple light into and out of the cavity, which can allow us measure the cavity response as opposed to spontaneous emission. In this sense, it will make our semiconductor-based system quite analogous to the neutral atom cavity QED experiments, where Rabi splitting is typically observed in the cavity's transmission spectrum [77, 78]. Most important, however, is the potential of this fiber-coupled system for observations beyond vacuum Rabi splitting, such as photon blockade [79] or few photon nonlinear optics. Chapter 8 presents a further discussion of some of these topics.

As our goal was to make the fiber taper an integral part of our cavity QED experiments, it was necessary to figure out how to best integrate it within our liquid He cryostat. In some ways, our problem paralleled the problem that Paul and Ben were facing in trying to integrate fiber-coupled microcavities into a UHV chamber for neutral atom cavity QED. They had decided that the appropriate way to tackle this latter problem was to try to robustly mount the taper to our microcavity chip with a UV-curable epoxy. This would result in a fiber-pigtailed device that, ideally, would be portable and easy to integrate in existing setups. Paul spent many weeks tackling this challenging, yet tedious task. There was little room for error, as a taper movement on the order of a micron would eliminate coupling to the cavity. A number of different methods and strategies were tried but were at best partially successful. Towards the end of the summer of 2005, however, he developed a technique that affixed the taper to on-chip support structures using microdroplets of UV-curable epoxy. This technique worked amazingly well, to the point that it was easier to break the support structures free from the rest of the chip than it was to remove the taper was from the support structures. Paul submitted these results, along with some impressive measurements of $Q \sim 4\times10^6$ for a SiN_x microdisk with $V_{eff} \sim 15(\lambda/n)^3$, for publication in April 2006 [80]. These SiN_x microdisks are transparent at $\lambda \sim 852$ nm, making them suitable for the Cs atom cavity QED experiments.

In the summer of 2005, Christopher Michael began working full-time in our group, and began working with me on the QD-microcavity project. In particular, he took the lead on trying to adapt the mounting techniques that Paul had developed for our low temperature setup. Of course, the primary concern here was the behavior of the UV epoxy at low temperatures. Chris was quickly able to reproduce Paul's mounting results for our AlGaAs microdisks at room temperature, and we had no significant difficulties in physically integrating the taper-coupled device into our cryostat. During the cooldown, however, the taper came free at around 200 K. Chris and I made multiple additional attempts at cooling down mounted samples over the next few weeks, and tried a number of heat-curable epoxies that in principle had good low temperature performance. None of these attempts succeeded, however. One important point is that the microdroplets of epoxy that we were using were significantly less than what any of these epoxies had been tested at during manufacturing. By mid-November 2005, we began to re-assess the situation.

Because Paul and Ben were tackling a similar problem, it seemed to make sense to try to leverage their recent success. In doing so, we overlooked the important point that Paul and Ben were doing taper mounting, rather than active positioning with mechanical stages, in part out of necessity, as they needed to make a device that was of low enough profile to avoid impeding any of the cooling and trapping laser beams within their vacuum chamber. This was not an all an issue for our QD work. Although the experiments using taper-mounted devices would potentially be relatively simple and elegant (and inexpensive), it is certainly not necessary, particularly in early experiments where the primary focus is on the physical phenomena we hope to see. We thus began to explore the possibility of incorporating a positioning setup within the cryostat. By the end of December 2005, I had worked out a design that would significantly modify our cryostat and let us integrate piezo-actuated stages into the chamber. The cryostat manufactures (Janis research) began construction of the modified parts in early 2006, and we ordered a set of piezo-actuated stages that were specifically engineered for operation at low temperatures and high vacuums.

While waiting for the new equipment to arrive, I began to more quantitatively investigate some of the phenomena we might expect to see in our experiments. This gave me the opportunity to learn some of the quantum master equation and quantum trajectory formalisms common to quantum optics research. Using the Quantum Optics Toolbox for Matlab developed by Sze Tan [81, 82], we began to simulate a coupled QD-microdisk system using the Q and V_{eff} values we would expect for our devices. One question of particular interest to us was how, if at all, the QD-microcavity interaction would differ in our devices, which had Qs that were high enough that surface scattering would couple the propagating modes of the disk and create standing wave modes. One immediate consequence of having standing wave modes is that their mode volumes are approximately half that of the traveling wave modes; the peak electric field strength for a standing wave mode is therefore $\sqrt{2}$ times larger than that for a traveling wave mode. Of course, another consequence is that the

Table 2: Experimentally demonstrated optical microcavities as of early 2006. The scanning electron microscope (SEM) images are for devices fabricated in the Painter research group (Si μdisk image courtesy of M. Borselli and SiN_x μdisk image courtesy of P.E. Barclay).

InP PC	**Si PC**	**Si μdisk**	**AlGaAs μdisk**	**SiN_x μdisk**
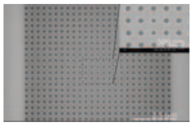	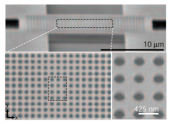		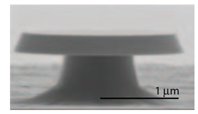	
Geometry	**Material**	Q	$V_{\mathbf{eff}}((\lambda/n)^3)$	**Reference**
photonic crystal	InP/InGaAsP	1.3×10^4	1.2	[43]
	Si	9.0×10^5	1.7	[30]
	Si	6.0×10^5	1.2	[28]
	GaAs/AlGaAs	2.0×10^4	1.0	[71]
microdisk	Si	1.3×10^6	14	[64]
		5.0×10^6	188	[65]
	GaAs/AlGaAs	3.6×10^5	6	[69]
		1.2×10^5	2.2	[76]
	SiN_x	3.6×10^6	15	[80]
microtoroid	SiO_2	1.0×10^8	500	[56]
		4.0×10^8	160	[83]
micropost	GaAs/AlGaAs	1.3×10^4	16	[84]
		2.8×10^4	79	[84]
	$AlGaAs/AlO_x$	4.8×10^4	51	[85]

new standing wave field would have nodes for which the QD-cavity interaction would be zero. We wanted to understand a bit more quantitatively how the system's spectral response would behave within different parameter regimes, where the quantities of interest are the QD-photon coupling rate g, the cavity decay rate κ, the QD decay rate $\gamma_\perp$, and the amplitude and phase of the backscattering parameter β. Finally, the per photon electric field strengths within these cavities are so large that nonlinear optics at the single photon level becomes a possibility. Our initial consideration of these topics are the focus of chapter 8.

In table 2, I have updated table 1 to report overall progress in the field that has occurred during the time of my graduate research. This table, which contains experimentally demonstrated values for Q and V_{eff} for a number of different microcavity geometries in different materials, is a good indication of the large amount of effort that has been placed on developing microphotonic structures with low optical losses (other types of devices, such as PC waveguides, have also exhibited markedly reduced losses over the last few years, although I have not included these structures within the table).[1] In addition to the semiconductor microcavities that I have focused on in this thesis, there has been significant progress in the development of chip-based dielectric cavities, such as the SiO_2 microtoroid cavities first demonstrated by Deniz Armani in Professor Kerry Vahala's group at Caltech [56], which exhibit Qs in excess of 100 million, and the SiN_x cavities demonstrated by Paul in our group [80], which have Qs of a few million and mode volumes that are significantly reduced in comparison to the microtoroid geometries. As these structures are fabricated in materials that are transparent throughout the visible and near-IR spectrum (unlike high-index semiconductors such as Si, GaAs, and InP), they have potential for integration with a number of systems, including colloidal quantum dots, impurity states in crystalline films, and alkali atoms, which all have optical transitions at less than 1 μm. In general, the results of table 2 indicate that the field has significantly and rapidly matured, through a combination of progress in fabrication technology, our understanding of how to design these devices, and in the development of tools with which these devices can be probed. For quantum optics applications, this means that we are finally in the position to take advantage of the intense electric fields that are supported by these ultrasmall volume microcavities, in order to study coherent light-matter interactions in solid-state systems.

The organization of this thesis is as follows. Part I details the design and experimental realization of ultrasmall volume, high quality factor photonic crystal microcavities. Measurements of InP-based multi-quantum-well lasers and passive Si resonators are presented. Part II is focused on microdisk resonators in the AlGaAs system. These microdisks contain an integral layer of InAs quantum dots, and measurements of the loss properties and lasing characteristics of these devices are presented. Simulations of the predicted behavior of these structures in the strong coupling regime of cavity QED are also considered. Finally, the appendixes contain relevant background material for the topics discussed within the body of the thesis.

---

[1] In addition to the cavities described in this table, there have been significant efforts in developing new cavity geometries. These include defect cavities in full three-dimensional photonic crystals [86] and other types of cavities employing the distributed Bragg reflection confinement mechanism, including circular Bragg resonators [87, 88] and Bragg "onion" resonators [89].

List of Publications

Listed below are publications that took place during my time in graduate school. They have been divided into two sections, to separate those that are most closely related to this thesis from those that are not.

Publications most closely related to this thesis

- K. Srinivasan and O. Painter, Momentum space design of high-Q photonic crystal optical cavities, *Optics Express*, **10**(15), pp. 670-684 Jul. 29, 2002.

- K. Srinivasan and O. Painter, Fourier space design of high-Q cavities in standard and compressed hexagonal lattice photonic crystals, *Optics Express*, **11**(6), pp. 579-593, Mar. 24, 2003.

- O. Painter and K. Srinivasan, Localized defect states in two-dimensional photonic crystal slab waveguides: a simple model based upon symmetry analysis, *Physical Review B*, **68**, 035110, Jul. 2003.

- O. Painter, K. Srinivasan, and P. E. Barclay, A Wannier- like equation for localized resonant cavity modes of locally perturbed photonic crystals, *Physical Review B*, **68**, 035214, Jul. 2003.

- K. Srinivasan, P. E. Barclay, O. Painter, J. Chen, A.Y. Cho, and C. Gmachl, Experimental demonstration of a high-Q photonic crystal microcavity, *Applied Physics Letters*, **83**(10), pp. 1915-1917, Sep. 8, 2003.

- R. Colombelli, K. Srinivasan, M. Troccoli, O. Painter, C. Gmachl, D. M. Tennant, A. M. Sergent, D. L. Sivco, A. Y. Cho, and F. Capasso, Quantum cascade surface-emitting photonic crystal laser, *Science*, **302**(5649), pp. 1374-1377, Nov. 21, 2003.

- K. Srinivasan, P. E. Barclay, and O. Painter, Fabrication-tolerant high quality factor photonic crystal microcavities, *Optics Express*, **12**(7), pp. 1458-1463, Apr. 5, 2004.

- K. Srinivasan, P. E. Barclay, O. Painter, J. Chen, and A. Y. Cho, Fabrication of high quality factor photonic crystal microcavities in InAsP/InGaAsP membranes, *Journal of Vacuum Science and Technology B*, **22**(3), pp. 875-879, May, 2004.

- K. Srinivasan, O. Painter, R. Colombelli, C. Gmachl, D. M. Tennant, A. M. Sergent, D. L. Sivco, A. Y. Cho, M. Troccoli, and F. Capasso, Lasing mode pattern of a quantum cascade photonic crystal surface-emitting microcavity laser, *Applied Physics Letters*, **84**(21), pp. 4164-4166, May 24, 2004.

- B. Lev, K. Srinivasan, P. E. Barclay, O. Painter, and H. Mabuchi, Feasibility of detecting single atoms using photonic bandgap cavities, *Nanotechnology*, **15**, S556-S561 2004.

- K. Srinivasan, P. E. Barclay, M. Borselli, and O. Painter, Optical-fiber-based measurement of an ultrasmall volume high-Q photonic crystal microcavity, *Physical Review B, Rapid Communications*, **70**, 081306(R), Aug. 25, 2004.

- M. Borselli, K. Srinivasan, P. E. Barclay, and O. Painter, Rayleigh scattering, mode coupling, and optical loss in silicon microdisks, *Applied Physics Letters*, **85**(17), pp. 3693-3695, Oct. 25, 2004.

- P. E. Barclay, K. Srinivasan, and O. Painter, Nonlinear response of silicon photonic crystal microresonators excited via an integrated waveguide and fiber taper, *Optics Express*, **13**(3), pp. 801-820, Feb. 7, 2005.

- K. Srinivasan, P.E. Barclay, and O. Painter, Photonic crystal microcavities for chip-based cavity QED, *Physica Status Solidi b*, **242**(6), pp. 1187-1191, March 24, 2005.

- K. Srinivasan, M. Borselli, T. J. Johnson, P. E. Barclay, O. Painter, A. Stintz, and S. Krishna, Optical loss and lasing characteristics of high-quality-factor AlGaAs microdisk resonators with embedded quantum dots, *Applied Physics Letters*, **86**, 151106, April 6, 2005.

- K. Srinivasan, P. E. Barclay, M. Borselli, and O. Painter, An optical fiber-based probe for photonic crystal microcavities, *IEEE Journal on Selected Areas in Communications (special issue on photonic crystals)*, **23**(7), pp. 1321-1329, July 2005.

- K. Srinivasan, A. Stintz, S. Krishna, and O. Painter, Photoluminescence measurements of quantum-dot-containing semiconductor microdisk resonators using optical fiber taper waveguides, *Physical Review B*, **72**, 205318, Nov. 10, 2005.

- K. Srinivasan, M. Borselli, O. Painter, A. Stintz, and S. Krishna, Cavity Q, mode volume, and lasing threshold in small diameter AlGaAs microdisks with embedded quantum dots, *Optics Express*, **14**(3), pp. 1094-1105, Feb. 6, 2006.

Other related publications

- O. Painter, K. Srinivasan, J. D. O'Brien, A. Scherer, and P. D. Dapkus, Tailoring of the resonant mode properties of optical nanocavities in two-dimensional photonic crystal slab waveguides, *Journal of Optics A-Pure and Applied Optics*, **3**(6), pp. S161-S170, Nov. 2001

- O. Painter and K. Srinivasan, Polarization properties of dipolelike defect modes in photonic crystal nanocavities, *Optics Letters*, **27**(5), pp. 339-341, Mar. 1 2002

- P. E. Barclay, K. Srinivasan, M. Borselli, and O. Painter, Experimental demonstration of evanescent coupling from optical fiber tapers to photonic crystal waveguides, *Electronics Letters*, **39**(11), pp. 842-844, May 29, 2003

- P. E. Barclay, K. Srinivasan, and O. Painter, Design of photonic crystal waveguides for evanescent coupling to optical fiber tapers and for integration with high-Q cavities, *Journal of the Optical Society of America B-Optical Physics*, **20**(11), pp. 2274-2284, Nov. 2003

- D. M. Tennant, R. Colombelli, K. Srinivasan, M. Troccoli, O. Painter, C. Gmachl, F. Capasso, A. M. Sergent, D. L. Sivco, and A. Y. Cho, Fabrication methods for a quantum cascade photonic crystal surface emitting laser, *Journal of Vacuum Science and Technology B*, **21**(6), pp. 2907-2911, Nov/Dec. 2003

- R. Colombelli, K. Srinivasan, M. Troccoli, O. Painter, C. Gmachl, D. M. Tennant, A. M. Sergent, D. L. Sivco, A. Y. Cho, and F. Capasso, Fabrication technologies for quantum cascade photonic-crystal microlasers, *Nanotechnology*, **15**, pp. 675-681, 2004

- P. E. Barclay, K. Srinivasan, M. Borselli, and O. Painter, Efficient input and output fiber coupling to a photonic crystal waveguide, *Optics Letters*, **29**(7), pp. 697-699, Apr. 1, 2004

- P. E. Barclay, K. Srinivasan, M. Borselli, and O. Painter, Probing the dispersive and spatial properties of planar photonic crystal waveguide modes via highly efficient coupling from optical fiber tapers, *Applied Physics Letters*, **85**(1), pp. 4-6, Jul. 5, 2004

- K. Srinivasan and O. Painter, Design of two-dimensional photonic crystal defect states for quantum cascade laser resonators, (http://arxiv.org/abs/physics/0410068), Oct. 11, 2004