ENHANCING MICROCAVITY POLARITONS FOR TECHNOLOGICAL APPLICATIONS

by

Mark D. Steger

B.S., Physics, Albright College, 2008

B.S., Chemistry, Albright College, 2008

M.S., Physics, University of Pittsburgh, 2012

Submitted to the Graduate Faculty of

the Kenneth P. Dietrich School of Arts and Sciences in partial

fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2016

UNIVERSITY OF PITTSBURGH

DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Mark D. Steger

It was defended on

December 9th 2015

and approved by

David W. Snoke, Professor, Physics and Astronomy

Robert P. Devaty, Associate Professor, Physics and Astronomy

Michael Wood-Vasey, Associate Professor, Physics and Astronomy

Robert D. Coalson, Professor, Chemistry, Physics and Astronomy

Kevin P. Chen, Associate Professor, Electrical and Computer Engineering

Dissertation Director: David W. Snoke, Professor, Physics and Astronomy

ABSTRACT

ENHANCING MICROCAVITY POLARITONS FOR TECHNOLOGICAL APPLICATIONS

Mark D. Steger, PhD

University of Pittsburgh, 2016

Microcavity exciton-polaritons, semiconductor quasiparticles that are a unique mixture of light and matter, are routinely used to study quantum many-body phenomena. Due to the light mass of the polariton, $\sim 10^{-4}$ times the bare electron mass, polaritons manifest noticeable quantum effects even at room temperature.

As solid state systems, microcavity polaritons are generally robust and compatible with current semiconductor technology. Microcavity chips could be integrated into electronic or optical circuits. I present a demonstration of microcavity polaritons as an all-optical transistor, where the strong nonlinearity of the system leads to a change in the reflectivity for a signal light-ray from high to low. I also discuss the promise of using strongly coupled microcavities as low-threshold polariton lasers, which could replace traditional lasers in some cases.

The last two decades have seen great strides in the material systems used in microcavities, even demonstrating strong coupling at room temperature. GaN, CdZnSe, organic semiconductors and more recently, MoS_2 have supported strong coupling at ambient conditions. This makes technological applications more promising. I present our current progress in this field. Also, the general quality of microcavities has advanced steadily over this time. I demonstrate that our long-lifetime polaritons persist for an order of magnitude longer than in similar samples. This opens up new regimes of study and technological application as these particles thermalize better and carry quantum coherence over macroscopic distances.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

PREFACE

Firstly, I want to thank my advisor, Dr. David Snoke, for mentoring me over the past seven years. He is an excellent teacher and selflessly devoted many hours to offering extracurricular lectures for graduate students, which have benefited me several times. Moreover, he has strengthened in me an eagerness for research and encouraged me to communicate with the broader scientific community, resulting in new contacts and new collaborations.

Also, I appreciate the input and support of my current and past committee members, Drs. Robert Devaty, Michael Wood-Vasey, Robert Coalson, Andrew Daley and Kevin Chen. Their attention and advice has been very helpful over the years.

Several collaborations have made this work possible. Loren Pfeiffer and Ken West at Princeton have produced phenomenal GaAs samples; Ken Burch and Marcel Hoek at Boston College have worked rapidly with us to develop new MoS_2 microcavities; and Angelo Mascarenhas and Kirstin Alberi have opened up a new material system by growing II-VIs.

My high school teachers and college professors were integral to my early interest in physics and research. Mr. Ken Alwine was a fantastic physics teacher who taught me that "physics is everywhere." His passing certainly left a hole in my hometown due to the impact he had on many students. My college professors, Drs. Brian Buerke and Devon Mason, pushed me to further apply myself to the subject. In particular, Dr. Buerke led me to further develop my burgeoning passion for optics and research in general, as he advised me on two undergraduate research projects as well as an intimate optics-based senior seminar.

Certainly, the past seven years would not have been the same without my labmates. I want to thank my seniors, Ryan Balili, Bryan Nelsen, Nick Sinclair, and Jeff Wuenschell for patiently guiding me and teaching me. I especially recognize Bryan for mentoring me so closely; I feel as though he taught me everything I know about the lab. I appreciate the

commiseration and cooperation of my peers, Gangqiang Liu and Chitra Gautham, as we have all assisted one another whenever possible. Lastly, I want to acknowledge David Myers and Jonathan Beaumariage for tolerating my harrying while trying to pass along all of the details of our projects and equipment to them. Otherwise, I have cherished the comradeship my labmates and other graduate students in the department, particularly Kevin Sapp.

The staff of the physics department was also vital to this dissertation. I owe great thanks to the administrative staff, machinists, and lab technicians.

Moreover, I want to thank friends and family for *not* helping me make progress on my dissertation. The greater Ultimate Frisbee community in Pittsburgh gave me countless opportunities to play high quality games on an informal level, both relieving stress and improving myself outside of the Physics Department.

I want to thank my family for always being supportive and accepting.

Lastly, but most important, I want to thank my wife, Soma Chatterji, for being strong and caring. You have kept me in touch with family and friends and kept our lives social when I otherwise might have lost touch. We have kept each other grounded and supported one another during all times, both terrific and trying.

I also want to thank my several sources of financial support. My work was supported under the National Science Foundation under grant nos. DMR-1104383, ECCS-010913, ECCS-1243778, and PHY-1205762. Also, I am especially honored to have been awarded the 2015 PQI Graduate Award and the 2010 Thomas Lain Essay Award.

1.0 INTRODUCTION

The exciton-polariton is a quasiparticle which is a quantum mixture of light and matter. One can envision it as a photon dressed by strong interactions with the electronic transitions in a medium. In this sense, the polariton is a very light mass particle (thanks to its photon constituent) with interparticle interactions contributed by the electronic matter of the system it interacts with.

Specifically, we work with semiconductor samples made from planar thin films as in Fig. 1.1 (a). The mirrors and cavity spacer create an allowed cavity photon state (see section 2.2). This changes the density of states of the electromagnetic field within the cavity and enhances the desired field mode. The quantum well is a thin layer of semiconductor material with a strong optical absorption into an exciton state (discussed in section 2.1.2). The absorption and reemission process couples the exciton and photon. In the case that the two states are distinguishable, we say that they are weakly coupled. However, if the coupling between the states is enhanced enough, a Rabi splitting appears and the system is said to be strongly coupled. The once degenerate exciton and photon become the non-degenerate upper and lower polaritons, as in Fig. 1.1 (b).

The effective mass of the polariton is $10^{-4}m_0$, where m_0 is the bare electron mass, as discussed in Section 2.3.3. This property allows these particles to exhibit quantum phenomena at readily accessible temperature and energy scales. For example, the de Broglie wavelength of a rubidium atom, a typical species used to study quantum mechanics in atomic, molecular and optical (AMO) physics, is on the order of 0.02 nm at room temperature, while that of a polariton at the same temperature is on the order of 800 nm. In cold atom condensation studies, the heavy atoms are routinely cooled to temperatures of 100's of nK to increase the de Broglie wavelength to a manageable scale, while in typical microcavity polariton

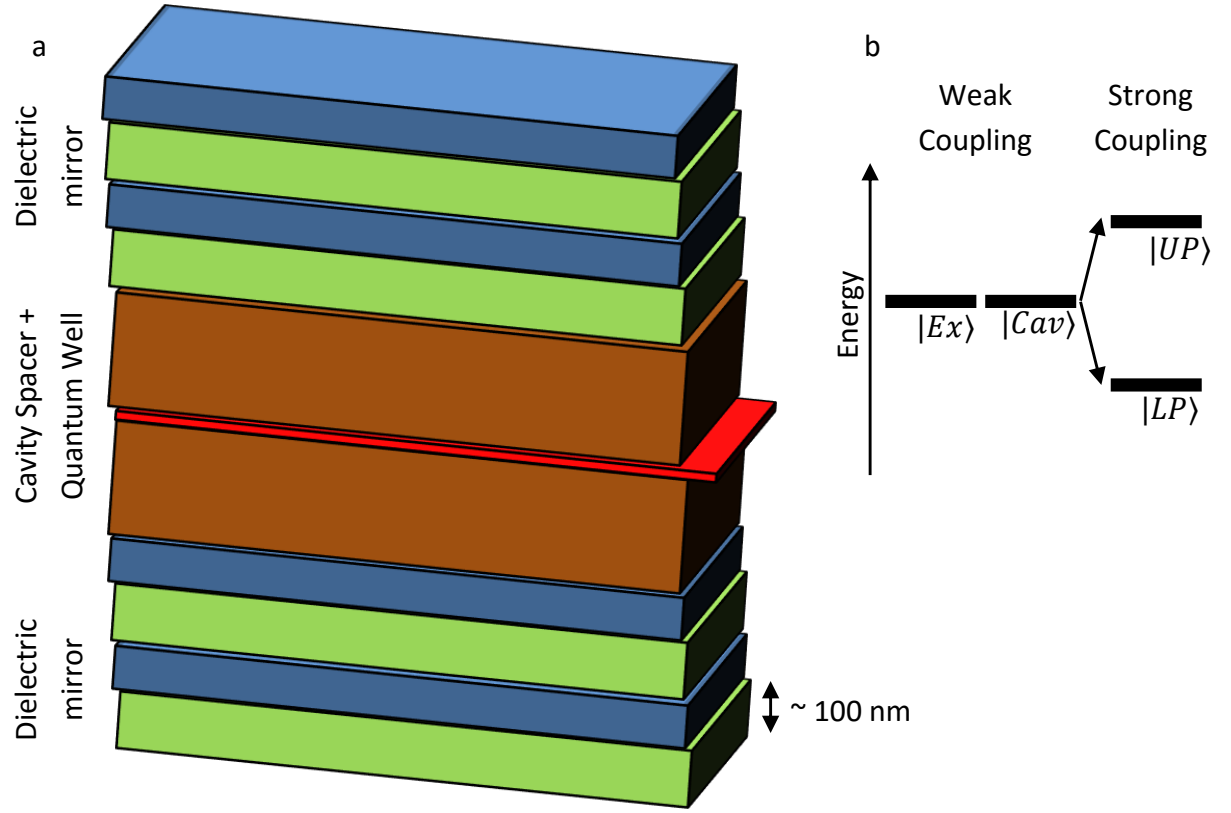

Figure 1.1: Frame (a): a semiconductor microcavity polariton consists of dielectric mirrors enclosing a cavity spacer. Embedded in the cavity spacer is a quantum well of semiconductor material. The cavity defines a confined photon state, which can be absorbed in the quantum well to form an exciton. Typical length scales of the optical layers are on the order of 50-100 nm, while the quantum well is on the order of 10 nm or less. Frame (b): in weak coupling, the exciton ($|Ex>$) and cavity photon ($|Cav>$) exist as degenerate states, but in strong coupling, the degeneracy is lifted to form the upper ($|UP>$) and lower ($|LP>$) polaritons.

experiments we work at 10-20 K.

Initially, polaritons were heavily studied to demonstrate Bose Einstein Condensation (BEC) and related effects. By 2006 and 2007, BEC was demonstrated in localized disorder states in microcavities [1] and in strain traps [2]. A further transition to lasing at higher density distinguishes this phase from traditional lasers [3] and suggests that polari-

ton condensation can be used to generate coherent light at lower density thresholds than is required for laser operation [4–6]. Since then, many more quantum condensation related effects have been demonstrated in polariton systems, including polariton lasing[1] [3, 4, 6–11], solitons [12–16], quantized vortices [17–20], and Jospehson junctions [21].

In addition to being a sandbox for studying quantum many-body effects, polaritons are increasingly being studied for technological applications. Since the polariton is a composite particle with properties governed by its light and matter constituents, typically in a micron-scale volume of structured material, it is relatively easy to engineer the properties of polaritons. One can tune the energies of the photon and exciton by changing thicknesses of particular layers. Even particle lifetimes can be engineered into a structure. Ultimately, the sample designer can change the particle interaction strength, mass, resonant energy, saturation threshold, and a slew of other parameters.

One application of microcavity polaritons is in the area of nonlinear optics. Exciton-polaritons exhibit huge optical nonlinearities, partly due to the transition from strong to weak coupling. This can, for example, be used to engineer all-optical switches in which one relatively weak gate pulse can be used to turn on or off absorption or reflection of a signal beam. Such optical transistors are not a new concept. In fact, over the past several decades a plethora of systems and geometries have been proposed for these devices, but none have proven themselves the way that electronic transitors have.

Most high quality, early results came from GaAs or CdTe systems studied at cryogenic temperatures on the order of 10 K. Physics in semiconductor systems tends to look 'cleaner' at cryogenic temperatures because the broadening due to phonon interactions is reduced. However, polaritons potentially permit us to study quantum phenomena at ambient temperatures. Additionally, any technological applications of polaritons are much less useful if samples must be maintained at liquid helium temperatures. For these reasons, a significant thrust of my work has been to develop high quality room temperature samples.

[1] Here I use the term polariton lasing to refer to the coherent quasi-condensate that forms in polariton systems that do not completely thermalize. Thus polariton lasing is less rigorous than polariton BEC, although it still refers to a coherent accumulation of particles in the ground state of the system.

1.1 DISSERTATION OVERVIEW

Traditional interest in microcavity polaritons has generally been analogous to AMO interest in cold atoms: researchers have seen these particles as a way to test our basic scientific knowledge by forming and manipulating quantum condensates. The past two decades have seen a substantial number of articles [1–3,22–28], reviews [7,29,30] and theses [31–33] focused on condensation of microcavity polaritons. In contrast, the focus of my work has been to develop technological applications for polaritons, and to transform the specific samples that we use to make them more applicable. As solid state systems, microcavity polaritons immediately lend themselves to real world applications, just like silicon-based transistor technology. One of my first projects was the demonstration of a microcavity polariton transistor switch, in which one light beam can switch the reflectivity for a second beam to be high or low. As novel as this application is, it isn't very useful with traditional microcavity samples that are only in strong coupling at cryogenic temperatures. Therefore, the long-term objective of my thesis work was to develop samples that exhibit strong coupling at room temperature. Such samples also bring the promise of other applications, such as the low-threshold coherent light emission from polariton lasing.

A second major improvement in microcavity polaritons over the past five years has been a drastic increase in the particle lifetime. This was immediately seen as a huge boon to condensation studies, since polariton condensates are closer to thermal equilibrium when the lifetime is longer than the thermalization time. Also, longer lifetime polaritons are useful for applications: I demonstrate propagation lengths on the order of millimeters, meaning that information can propagate across a macroscopic 'polariton chip,' should we ever make a polaritonic integrated circuit.

Table 1.1 in Section 1.2 gives a quick reference to most abbreviations used throughout this dissertation.

Chapter 2 covers the relevant physics of microcavity polaritons with great detail. I start by covering the electronic quasiparticles in solid state semiconductors in Section 2.1. Of specific importance is the derivation of excitons in Section 2.1.2 and how they couple to light in Section 2.1.3.

Section 2.2 covers many properties of dielectric mirrors and optical cavities. Here I derive the origin of the cavity photon mode in Section 2.2.2 and calculate many of its properties in Section 2.2.3. I demonstrate the relation between the physical structure of the cavity to its lineshape as a generic resonance in Sections 2.2.3.1 and 2.2.3.2.

In Section 2.3 I derive the basis of strong coupling between the exciton and photon modes assuming that they are coupled oscillators. Assuming that these two pure states mix, I develop many polariton properties in Sections 2.3.2 and 2.3.3. Section 2.3.4 shows how particle decay rates can cause strong coupling to collapse into the weak coupling regime, or how the broadened states can overlap to effectively spoil the anticrossing. Section 2.3.5 includes disorder in the exciton state, which can complicate the polariton system. These last two effects are very relevant to developing room temperature samples using new materials, as the disorder and decay times must be carefully considered in order to achieve strong coupling.

Chapter 3 discusses the numerous ways in which we interact with and make observations of the polaritons. Since polariton decay correlates directly to photon emission, our experimental techniques are nearly all optical.

Chapter 4 is an adaptation of my 2013 PRB publication [34]. In 2010, I modified previous designs attempting to produce long lifetime polaritons. The new samples showed qualitatively new condensation effects as published in [35], but the experimenters in the field wanted a clear measurement of the lifetime. In this paper, I develop multiple experiments and analyses to estimate the polariton lifetime in this system. Additionally, this paper directly demonstrates the long-range motion of these polaritons, and time resolves the ballistic motion according to the novel polariton dispersion.

In Chapter 5, based on my 2015 Optica publication [36], I make a more direct measurement of the polariton lifetime, improving on the results of Chapter 4. This experiment uses a carefully arranged geometry to resonantly inject polaritons against the inherent photon gradient in the sample, which causes them to turn around. I carefully aligned the gradient and the propagating jet with a streak camera time slit in order to carefully measure intensity vs position vs time. This yields the most direct measurement of a polariton lifetime ever presented in the literature, and clearly shows these samples to have a cavity mode lifetime

an order of magnitude higher than the next best samples. Since the data is taken when the polaritons are turning around, it is also a novel demonstration of photons (dressed as polaritons) decelerating as massive particles, turning around and traveling backwards.

Chapter 6 details the methods and results of my 2012 Applied Physics Letters publication [37] in which I demonstrate the operation of a microcavity reflectivity switch. This is one demonstration of using a microcavity as an all-optical transistor, where the state of the switch is given by the reflectivity: a gate beam locally modulates the reflectivity, which is probed by a signal beam. The years of 2011-2013 witnessed a race to demonstrate optical switching in microcavity polaritons, so I give an overview of alternative switching schemes in Section 6.1.

I present progress toward room temperature samples in Chapter 7. Several different systems have been pursued over the past decade and a half, and several groups have demonstrated strong coupling at room temperature. I give an overview of the history of room temperature samples, and also outline the specific systems that we have been working to improve.

1.2 GLOSSARY OF ABBREVIATIONS

The following abbreviations are commonly used throughout this dissertation. This section is intended to be used as a quick reference for clarification.

Table 1.1: Glossary of abbreviations

Abbreviation	Term
CW	Continuous Wave
DBR	Distributed Bragg Reflector
FSR	Free Spectral Range
FWHM	Full Width at Half Maximum
LP	Lower Polariton
MBE	Molecular Beam Epitaxy
NA	Numerical Aperture
PECVD	Plasma Enhanced Chemical Vapor Deposition
PL	PhotoLuminescence
QW	Quantum Well
TE	Transverse Electric
TM	Transverse Magnetic
TEM	Transverse ElectroMagnetic
UP	Upper Polariton
$\mathcal{Q}$	Q-factor
Ω	Vacuum Rabi splitting

2.0 MICROCAVITY POLARITON PHYSICS

In this chapter I develop the physics of strong coupling and the underlying systems necessary to understand microcavity polaritons. In Section 2.1, I briefly review the properties of electrons, holes and excitons in semiconductor systems, and discuss the coupling of these particles to light. Section 2.2 contains a discussion of the foundations of optical cavities and dielectric mirrors. A thorough understanding of the cavity mode is critical to designing successful microcavities, since the exciton must be overlapped spatially and spectrally to achieve strong coupling. I derive the origin of strong coupling between the exciton and photon in Section 2.3. This covers the boundary between strong and weak coupling, and details many important parameters of microcavity polaritons.

2.1 OVERVIEW OF SOLID STATE QUASIPARTICLES

In order to better understand the properties of relevant phenomena in condensed matter systems, we can renormalize our mental picture of the system and consider excitations out of the ground state to be genuine particles. For example, in a semiconductor at zero temperature, the electrons are collectively frozen in place so that the electrical conductivity is strictly zero. However, at finite temperature or under other excited conditions (such as optical illumination), some of those electrons in the valence band can be excited into the conduction band. If an electron is excited in this manner, that electron will be free to move through the crystal and carry a current thanks to its negative charge. Moreover, the promotion of that electron leaves a vacancy in the valence band. This means that the vast number of electrons in the valence band can move into the vacancy, much like one can move tiles in a

sliding block puzzle. However, it is again senseless to try and track all these electrons in the valence band. Instead, we track the *vacancy* and call this a particle of its own: we call it a *hole*. This hole is positively charged, being a lack of an electron, and can also carry current by moving through the crystal.

In this way, we can consider a crystal lattice and most electrons in it to be just a stage and these conduction band electrons and valence band holes to be the main actors. To differentiate these particles from fundamental particles (such as vacuum electrons, photons, or the whole cast of particles studied in particle physics), we call these quasiparticles. In fact, there are many more quasiparticles which become relevant depending on the system and the regime which is being investigated. However, just because these particles are not fundamental does not mean they are less significant. In fact, many quasiparticles exhibit unique phenomena which we could not understand in terms of the collective motion of the underlying fundamental particles. Some clear examples of this are as follows: on the scale of chemistry, molecules and atoms are the most useful particles, where molecules can be changed by exchanging atoms or groups of atoms; however, on the scale of fission and fusion, one must recognize that atomic nuclei are made up of protons and neutrons, which are further made up of (and can change form by exchanging) leptons and quarks. Just because the chemist may not consider the fine structure of the atomic nucleus does not make his predictions any less valid. On the contrary, it is much easier to understand molecular interactions on typical energy scales when we focus on atoms as the main actors.

For the purpose of this dissertation I will develop the framework necessary to understand the exciton-polariton as a unique quasiparticle, and highlight the origin of its parameters in terms of the underlying crystal lattice and electronic states.

2.1.1 ELECTRONS AT THE BAND EDGE

Here I will discuss the general methods for computing electronic effects in crystalline semiconductors. While the specifics of a solid state system can make calculations very difficult, we will utilize a series of methods that make problems tractable. For example, by recognizing the inherent symmetry of specific crystal lattices, we can impose the same symmetries on

the electronic quantum solutions.

The Schrödinger equation for a single electron inside a crystal lattice is:

$$\hat{H}\psi(\vec{x}) = \left[-\frac{\hbar^2}{2m_0}\nabla^2 + V(\vec{x})\right]\psi(\vec{x}) = E\psi(\vec{x}), \tag{2.1}$$

where we assume that the potential, $V(\vec{x})$ is a periodic potential in $\vec{x}$. Here we ignore spin-orbit coupling for simplicity. The reduced Planck's constant is given by $\hbar$, the bare electron mass by m_0, and the electron wavefunction in space by $\psi(\vec{x})$. We note that the functional form of the potential can be anything in general, but we will envision it being the Coulomb potential of the stationary atomic nuclei, possibly screened by the more tightly bound core electrons.

We can exploit the periodicity of the potential to make a generalization about the solution $\psi(\vec{x})$. Following the standard solid state procedure of applying Bloch's theorem (see e.g. Refs. [38–40]), we first note that observables related to the electronic states must share the periodicity of the underlying lattice. This means the electronic wave functions must be of the form:

$$\psi_{n\vec{k}}(\vec{x}) = u_{n\vec{k}}(\vec{x})e^{i\vec{k}\cdot\vec{x}}, \tag{2.2}$$

where $u_{n\vec{k}}(\vec{x})$ is a (sometimes unknown) periodic function. We have included an extra quantum number, n, to account for the possibility of multiple solutions u to the single unit cell potential, i.e. different bands. Here $\vec{k}$ indicates the crystal momentum, i.e. the motion of an electron wavepacket in the crystal lattice. The true momentum depends on the both $\vec{k}$ and $u_{n\vec{k}}(\vec{x})$. The maximum value of this wavevector is the reciprocal lattice constant $k_{max} = \frac{2\pi}{a_0}$, where a_0 is the lattice constant (e.g. for the case of a cubic crystal).

The electrons are inherently spin 1/2 particles. This classifies them as fermions, which means they can have at most one particle per state. Interestingly, as we will discuss below, when multiple fermions correlate with each other, they can act as composite bosons, such as the exciton and subsequently the exciton-polariton.

2.1.2 EXCITONS

Thus far we have discussed the states of free electrons in crystalline media. Even in an insulator, electrons can be excited into conducting states and leave behind positively charged holes. However, these particles do not strictly ignore each other after they form. Just as a positively charged proton (or positron) in free space will attract a free negatively charged electron, the valence band hole and conduction band electrons will feel a Coulomb attraction. In a direct analogy to the hydrogen atom, these quasiparticles will form bound Rydberg states.

Here I consider an isotropic toy-model of an exciton to highlight several features. Just like atomic energy levels of the electrons, there is a ladder of excited band states and a kinetic energy continuum of the center of motion:

$$E_x = -\frac{Ry}{n^2} + \frac{(\hbar k)^2}{2M}. \tag{2.3}$$

The second term in (2.3) accounts for the center-of-mass kinetic energy of the exciton where $M = m_e + m_h$ is the sum of the electron and hole masses. Here the binding energy scales with the effective Rydberg and has degeneracies identical to that of the hydrogen atom.

$$Ry = \frac{e^2}{8\pi\epsilon a_x}, \tag{2.4}$$

where e is the elementary charge, and a_x is the Bohr radius of the bound pair:

$$a_x = \frac{4\pi\hbar^2\epsilon}{e^2 m_r} \tag{2.5}$$

$$m_r = \frac{m_e m_h}{m_e + m_h}.$$

The Rydberg Ry reduces the total energy of an electron-hole pair and is thus a binding energy. If there are free electrons and holes in the crystal, they can cool down (typically by scattering with the lattice phonons) until they form these bound pairs. Once the electron and hole have significant spatial overlap, they eventually recombine, meaning that the conduction band electron falls into the corresponding hole in the valence band, and a photon is emitted to carry away the excess energy. It is not strictly necessary for the electron and hole to form

an exciton in order to recombine, but the probability of recombination is much greater due to the spatial overlap when bound together. Of course, there are other mechanisms by which electrons, holes and excitons can annihilate other than radiative recombination, but due to the effective overlap within the exciton, excitonic effects often dominate phenomena at the direct band edge.

By examining (2.3), (2.4), and (2.5), it is clear that the exciton Bohr radius (and therefore binding energy) will depend on several aspects of the crystal. In particular, the Rydberg is proportional to ϵ^{-2}, where ϵ is the dielectric constant (or tensor, for anisotropic media) of the crystal, which indicates the extent to which background electrons can screen the binding of the electron and hole. Also, the effective masses of the electron and hole enter into the Bohr radius, indicating that the band structure immediately at the gap can have a strong influence on the excitonic effects.

In general, a smaller Bohr radius corresponds to stronger excitonic effects and better light-matter coupling in microcavities. Again, this is due to the enhanced spatial overlap between the electron and hole. However, it can also be helpful because the excitons can persist at higher temperatures, since we should expect excitons to thermally ionize when $k_B T > Ry$, where k_B is the Boltzmann constant. For this reason, we will expect low-index (i.e. low dielectric constant) semiconductors to be better suited to excitonic applications.

In terms of semiconductor physics, we typically classify excitons into two categories according to the Bohr radius. Very strongly bound excitons with Bohr radius comparable to the lattice constant are called Frenkel excitons. These excitons appear commonly in organic molecules, where the exciton is just an excitation localized to a single molecule. The exciton can then hop from site to site, but slow dynamics limiting the hopping are a current problem in some fields, such as organic photovoltaics, where typically excitons must propagate from an absorption center to a charge separation interface.

Conversely, most moderately high-index semiconductors have much more weakly bound excitons, called Wannier excitons. These have a Bohr radius much larger than the crystal lattice constant, and therefore the constituent electron and hole can be interpreted as crystal excitations. These are the particles that we envisioned when deriving the Rydberg relations in (2.3), (2.4), and (2.5), which depend on the effective masses of these crystal excitations

and the background dielectric constant.

2.1.3 OPTICAL-ELECTRONIC TRANSITIONS

Fundamentally, optical absorption and emission arises from the electric dipole moment of electronic transitions in a crystal or molecule. The Hamiltonian for light-matter interaction is

$$\hat{H} = \frac{\hat{p}^2}{2m_0} + V(\hat{\vec{x}}) - \frac{e}{2m_0}(\hat{\vec{p}} \cdot \hat{\vec{A}} + \hat{\vec{A}} \cdot \hat{\vec{p}}) + \frac{e^2 \hat{A}^2}{2m_0}, \tag{2.6}$$

where $\hat{\vec{p}}$ is the momentum operator and $\hat{\vec{A}}$ is the vector potential. The first two terms account for the original Hamiltonian used in (2.1). The last term is negligible at typically achievable intensities in the optical regime, since $|eA| << |p|$. This means that most optical transitions are encoded in the $\hat{\vec{p}} \cdot \hat{\vec{A}}$ terms.

It is natural to work in the Coulomb gauge for the vector field–in which $\hat{\vec{A}}$ and $\hat{\vec{p}}$ commute–and express the field as a Fourier series of plane waves. In this case, $\hat{\vec{A}} = \frac{\vec{E}_0}{2\omega}[e^{i(\hat{\vec{k}} \cdot \hat{\vec{x}} - \omega t)} + \text{c.c.}]$, where $\vec{E}_0$ is the electric field amplitude (with polarization), ω is the optical angular frequency, $\hat{\vec{k}}$ is the wavevector operator, and c.c. is the complex conjugate. In this case the interaction terms in the Hamiltonian become

$$\hat{H}_I = -\frac{e\vec{E}_0}{2\omega m_0} e^{i(\hat{\vec{k}} \cdot \hat{\vec{x}} - \omega t)} \cdot \hat{\vec{p}} + \text{c.c.} \tag{2.7}$$

Primarily, we are concerned with the probability for an optical transition to couple states in the valence and conduction bands. When determining the matrix element for the optical transition, we assume that the optical field will have wavelength much larger than atomic orbitals or the typical lattice constants, so the $e^{i\hat{\vec{k}} \cdot \hat{\vec{x}}}$ becomes just a constant phase factor in the matrix element calculation[1]. Moreover, the time dependent exponential explicitly printed in (2.7), i.e. $e^{-i\omega t}$, will give rise to absorption, while the complex conjugate ($e^{i\omega t}$) leads to stimulated emission[2].

[1] It should be noted that the spatial dependence of the electric field does matter in a cavity, when the field amplitude is pinned in space. A radiator couples more strongly to the optical mode if it is at an antinode of the field

[2] This connection has roots in Fermi's golden rule, since non-energy conserving transitions will phase out [41].

Starting from the absorption matrix element,

$$H_{Icv} = \langle c|\hat{H}_I|v\rangle$$
$$H_{Icv} \propto -\frac{e\vec{E}_0}{2\omega m_0} \cdot \vec{p_{cv}}, \tag{2.8}$$

and noting the relation between $\hat{\vec{p}}$, $\hat{\vec{x}}$, and $\hat{H}_0$, the unperturbed Hamiltonian,

$$[\hat{\vec{x}}, \hat{H}_0] = \frac{i\hbar}{m_0}\hat{\vec{p}}, \tag{2.9}$$

this optical transition can be written in terms of the dipole moment:

$$H_{Icv} \propto -\frac{e(E_c - E_v)\vec{E}_0}{i\hbar\omega} \cdot \langle c|\vec{x}|v\rangle \approx i\vec{\mu_{cv}} \cdot \vec{E}_0, \tag{2.10}$$

where we have introduced $\vec{\mu_{cv}} = e\langle c|\hat{\vec{x}}|v\rangle$ as the dipole moment of the electronic transition. Here I have assumed that the broadening of the state is negligible and therefore the matrix element is only relevant if the energy difference between the levels $(E_c - E_v)$ is approximately equal to the energy of the photon being absorbed.

Embedded in the $\langle c|\hat{\vec{p}}|v\rangle$ matrix element is also a set of selection rules that will rule out certain possible transitions. For example, in GaAs at the band edge, the valence band is four-fold degenerate (after invoking spin-orbit coupling) while the conduction band is two-fold degenerate. In general, these transitions will have dipole moments (this is the spatial difference between the p-like valence band and the s-like conduction band), but the optical transition also requires the angular momentum of the photon to be conserved. Between all the possible transitions, some will change the electron spin by ± 1 and therefore be allowed. Other transitions that do not meet this requirement can not directly absorb or emit a single photon, and thus we call those transitions optically forbidden or dark states. States that are forbidden from emitting photons may have very long lifetimes. For example, the excitons that we use to make polaritons in GaAs samples have both bright and dark variants, which have total angular momentum J=1 or J=2, respectively. Since only the bright excitons will couple to the light and form polariton states, excitations will more readily decay via that channel. Any excitations that reach the dark state will first have to spin flip to form bright exciton-polaritons. Such effects are also ubiquitous in bulk and nanostructured excitonic materials, where the fastest decay channel is often radiative recombination of bright excitons [42, 43].

There is typically a continuum of states that will be optically excited in a semiconductor (or similarly in a molecular species), but excitonic contributions can be comparable or dominant over the nearby continuum transitions. Again, this is due to the large $\langle c|\hat{\vec{p}}|v\rangle$ matrix element. It is common to express the optical strength of a resonance in terms of the unitless oscillator strength, f:

$$f_{ba} = \frac{2|p_{ba}|^2}{m_0(E_b - E_a)}. \tag{2.11}$$

Here $E_b - E_a$ is the energy difference between the states.

2.1.4 CONFINED ELECTRONS, HOLES AND EXCITONS

It is common practice to engineer nanoscale structures into semiconductor devices to modify the electronic properties. Using growth methods such as molecular beam epitaxy (MBE), for example, allows us to accurately deposit only a few layers of a crystalline semiconductor material. This layer can then be clad with other materials, enabling us to effectively engineer the bandstructure on the nanometer scale.

The most basic such structure is a quantum well (QW), a 2D sheet of low-bandgap material clad on either side by higher bandgap material. In our GaAs-based structures, for example, we typically use 7 nm of pure GaAs embedded between pure AlAs layers. In the type-I quantum well, conduction-band electrons (and also holes) feel an effective potential landscape that confines them in the low-bandgap GaAs. Essentially, low energy electrons and holes will be confined in the growth direction, but can have any in-plane momentum.

Confinement in the QW increases the energy of the electron and hole (and therefore of the corresponding excitons and optical transitions). This means that we can tune the resonant energy of an exciton to some extent, but more importantly it gives us better control over the spatial location of excitons and usually enhances the binding energy and therefore oscillator strength of the exciton.

One can understand the enhancement of the oscillator strength by using the solutions of the 2D hydrogen or positronium atoms. In the case where particles are perfectly confined to a 2D plane, the exciton binding energy will be 4 times stronger than in the bulk case [44].

Importantly, this also enhances the oscillator strength, especially for light polarized parallel to the QW plane.

Ultimately, proper calculations of the properties of excitons in QWs is not a simple task. In principle, the electron and hole are trapped nominally in finite square well potentials, which already requires numerical solutions to the transcendental solutions to the Schrödinger equation. Moreover, the levels and offsets of the bandgaps in these regions must be properly accounted for (for example, see Appendix A in [44]). These offsets are typically predicted by the model-solid theory [45–47] and should be expected to have uncertainty up to the order of 0.1 eV. Furthermore, if the two materials are not perfectly lattice matched, one must account for the underlying strain, as it will deform the band structure. This is one reason why the GaAs/AlAs system is so well studied, as these materials are marvelously well lattice matched. Not only does this allow us to nominally ignore strain effects on the bandgap, but also relatively thick heterostructures of high quality can be grown before crystal order begins to degrade.

Phenomenologically, we can calculate the exciton ground state energy as the following:

$$E_x = E_{gap} + E_{e,conf} + E_{h,conf} - \text{B}_\text{E}, \tag{2.12}$$

where E_{gap}, $E_{e,conf}$, and $E_{h,conf}$ are the bandgap energy, electron confinement energy, and hole confinement energy, respectively. Typical confinement energies for electrons and holes in QWs in GaAs are on the order of 10-100 meV, as shown in table 2.1. B_E is the exciton binding energy (subtracted since it reduces the energy relative to unbound e-h pairs). Using the model-solid theory and the finite square well solutions, we can calculate a nominal solution to the first three terms; however, we note that the binding energy is not yet precisely calculated.

Exact solutions to the exciton binding energy are quite difficult to achieve. We can expect that the bound exciton will have B_E somewhere between the 3D and 2D cases, but we recognize that as we decrease the QW thickness, eventually the electron or hole may extend substantially into the barrier materials. This can drastically change the exciton binding. A straightforward (but costly) approach is using a variational principle technique to solve for the proper excitonic wavefunction [48], but these methods are computationally intense and

rely on making an acceptable guess as to *which* parameters to vary and *what* functionals to use. Others have proposed simpler methods to approximate the binding energy, such as so called dimensional interpolation [49].

Beyond 2D quantum structures, there are methods to further confine excitations. Researchers regularly create 1D wires and 0D dots to further change the physics.

2.1.5 EXCITONS: FURTHER CONSIDERATIONS

Quantum confinement and large crystal momentum can lead to other effects on electronic and excitonic states. In some cases, degeneracies are lifted, different states may be mixed, and in general the change of dimensionality results in a change of the density of states.

For example, in pure, bulk GaAs, the valence band is multiply degenerate. At lowest approximation, the p-like valence band is six-fold degenerate at the high symmetry Γ-point (including spin-degeneracy). Spin-orbit coupling lifts the degeneracy of one pair of states from the other four, and the so called spin-orbit (SO) split off band is shifted to lower energy. Shown in Fig. 2.1, what remains at the high-symmetry Γ point are the heavy-hole and light-hole valence bands (still both exhibiting 2-fold spin degeneracy). However, the bands are named for the different effective masses (curvature) of those bands at the $k = 0$ point. The light hole is almost an order of magnitude lighter than the heavy hole. This results in slightly different binding energies (order of 1 meV in bulk GaAs). This means that processes near the band edge of bulk GaAs must account for transitions from two degenerate valence bands and two nearly degenerate exciton states. In a QW, the light hole states are further split from the heavy hole due to the strong dependence of the confinement energy on the mass of a particle in a square well. Typical light-hole heavy-hole splittings are on the order of 20 meV in 7 nm QWs. This greatly simplifies the physics when considering optical processes at the band edge or coupling the exciton to the photon to achieve strongly coupled polaritons.

Table 2.1 contains some representative values for electron and hole confinement energies in a 7nm GaAs QW. The finite square well confinement is the more appropriate value to use, as the band offsets between the well and barrier may not be large compared to the corresponding infinite square well confinement energy. These band offsets are calculated as

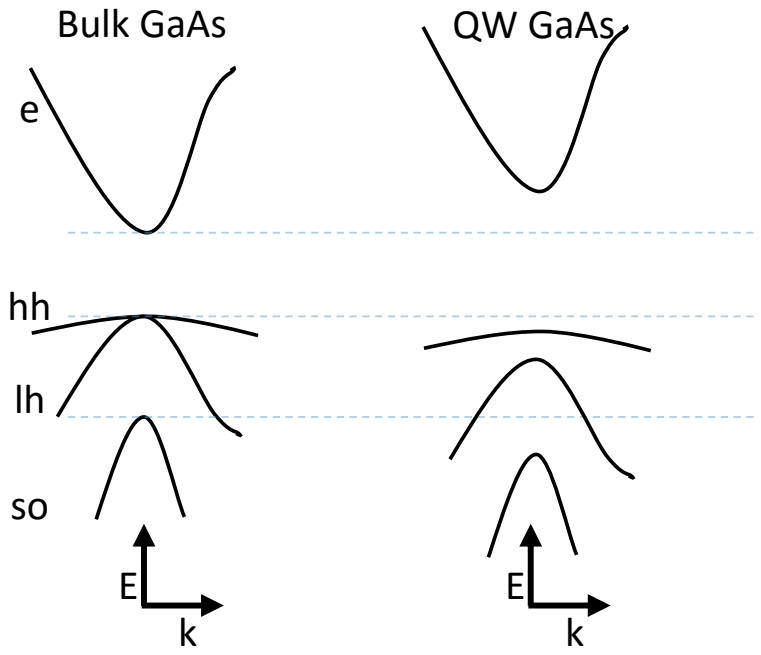

Figure 2.1: Band structure of GaAs at the Γ point. The conduction band (e) is singly degenerate (ignoring spin), while the valence band decomposes into the heavy hole (hh), light hole (lh) and split-off hole (so) band. QW confinement lifts the degeneracy of light and heavy holes. Energy scales and band curvature are not shown to scale.

described in Section 2.1.4. The large difference between the confinement energy of the light hole and heavy hole means that we can generally ignore the light-hole exciton since it is significantly higher in energy.

Exactly at the Γ point, we can easily calculate several properties of the excitons. We consider excitations as being purely light hole or heavy hole-like. However, at finite in-plane momentum, the Luttinger-Kohn (LK) Hamiltonian (based on the k·p method) introduces coupling between the valence bands [44]. Firstly, this means that the light hole and heavy hole (and in principle the SO hole) states are mixed at finite k. One implication is that this can change selection rules, such as for two photon-absorption [50]. Moreover, the anisotropy of the valence band masses results in interplay between the QW confined wavevector and the

Table 2.1: Confinement energies in a 7 nm GaAs/AlAs QW

Species	Infinite SQ confinement (meV)	Finite SQ confinement (meV)
Electron	115	66
Light hole	82	45
Heavy hole	23	18

in-plane wavevector. We often calculate the confinement energy at zero in-plane wavevector and assume it is fixed. To properly calculate the confined states, one should use the full LK Hamiltonian to generate self-consistent states [51].

In addition to the LK-Hamiltonian mixing otherwise independent states, other effects such as strain can mix states. The Pikus-Bir Hamiltonian accounts for this strain: the form of the mixing is very similar to that of the LK-Hamiltonian, except that the mixing terms are now generated by the deformation of the crystal, rather than the electron momenta [44]. This mixing is important when working with strained semiconductors (such as strained QWs), and when intentionally applying a strain field, such as for strain traps [51–53].

2.2 OPTICAL CAVITIES

So far, I have highlighted some fundamental electronic quasiparticles, namely the fermionic electron and hole, and the bosonic exciton. Photons, the fundamental bosons which are excitations in the electromagnetic field, can be trapped in a cavity between two mirrors. The mirrors can either be metallic (which are inherently lossy) or dielectric. By tuning the mirrors and cavity design, one can engineer the character of the photons that are permitted inside the cavity. Here I will highlight important aspects of mirror design and the dependence of a cavity photon on the structural design.

2.2.1 DIELECTRIC MIRRORS

The best mirrors for optical cavities tend to be thin films of dielectric or semiconductor materials. One can potentially also use thin or thick films of metal, in which the reflection coefficient is determined purely by the Fresnel reflection from a single interface. Any light that enters a thick region of metal will be absorbed on a length scale of the skin depth, so the only way to allow for any appreciable transmission is to use a layer of metal only a few nm thick.

For any real structure, it is useful to fully simulate the design in order to match it to the project requirements. Our lab employs a technique of solving Maxwell's Equations through an arbitrary 2D stack of materials referred to as the transfer matrix method, which is covered in detail in Appendix A. Essentially, one selects a wavelength and sets up a propagation matrix which propagates light through each layer and across each boundary of the stack, finally ending in a bulk region (the final 'layer' of the structure. This propagation matrix can be represented simply as a two-by-two complex valued matrix, so an n layer structure corresponds to just n, 2x2 matrix multiplications. The details of this method and the code used to implement it are given in this appendix.The simulations shown in Figs. 2.3 and 2.4 are generated using this method.

In order to produce higher quality cavities, we usually elect to use thin film interference effects from alternating layers of dielectric material with differing refractive indices. Such a stack of films is referred to as a distributed Bragg reflector (DBR). The standard design at wavelength λ_0 calls for each layer to be $\lambda/(4n)$ thick such that all reflected waves from the structure are perfectly in phase. Figure 2.2 demonstrates the accumulation of phase through the first several periods of a DBR.

Such a structure allows for low or high reflectivity mirrors, and the wavelength bandwidth of the reflectivity, called the stopband, can be tuned by material selection. Figure 2.3 shows simulated reflectivity from two DBRs designed to reflect 500 nm light. The first DBR simulation (heavy blue curve) uses index values of $n_L = 1.5$ and $n_H = 2$, which are the approximate values for SiO_2 and Si_3N_4, and uses 10.5 periods. The green curve replaces $n_H = 2.5$, but keeps the number of periods the same. The higher contrast DBR clearly has a

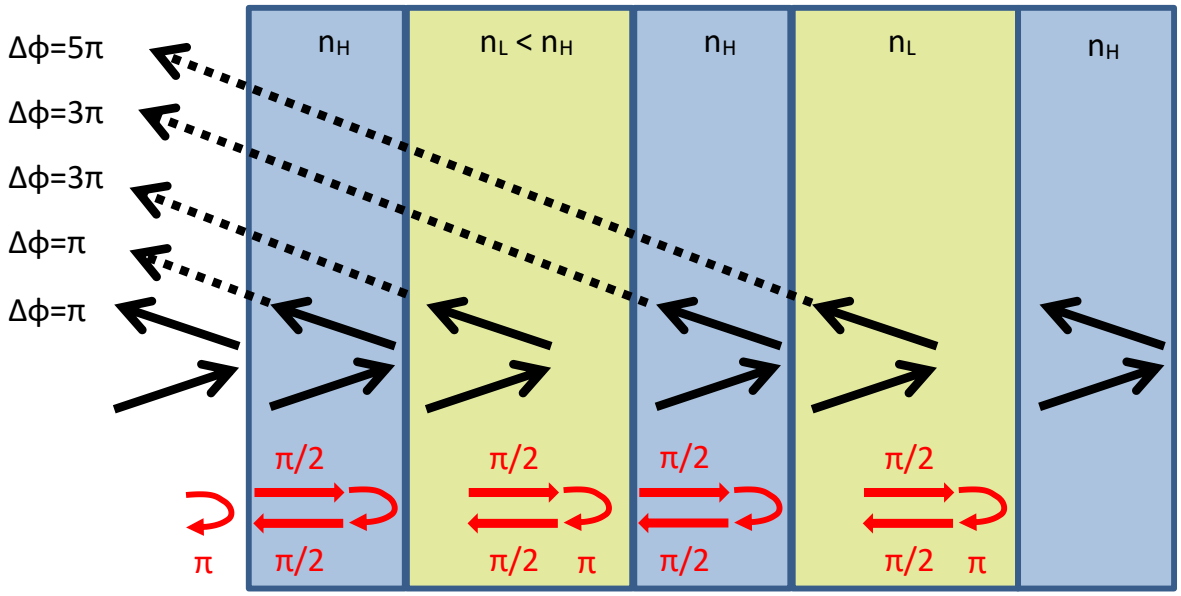

Figure 2.2: Phase delay from the first several reflections within a DBR. For clarity, the phase for each layer transit and each reflection is listed in red. Note that each reflection from low- to high-index results in an additional π phase shift. There are actually an infinite series of higher order reflections not shown in this figure.

broader stopband, but also has a slightly higher reflectivity. The red, dashed curve uses the same low contrast index values as the blue curve, but only has 5.5 periods in the DBR. This results in lower reflectivity, and washed out features. Thus, by selecting materials based on index contrast and engineering the number of periods, one can engineer the stopband width and maximum reflectivity.

Ultimately, it is useful to consider the complex phase induced upon reflection from a DBR in addition to the magnitude of the reflection. As discussed later in Equation 2.13, we must consider the phase induced from mirrors when designing a cavity thickness. Figure 2.4 demonstrates the variation of the induced phase of reflected light vs wavelength. The DBR is designed for 500 nm. This wavelength exhibits maximum reflectivity and also an induced phase of exactly π. However, away from the design wavelength we see a noticeable shift in

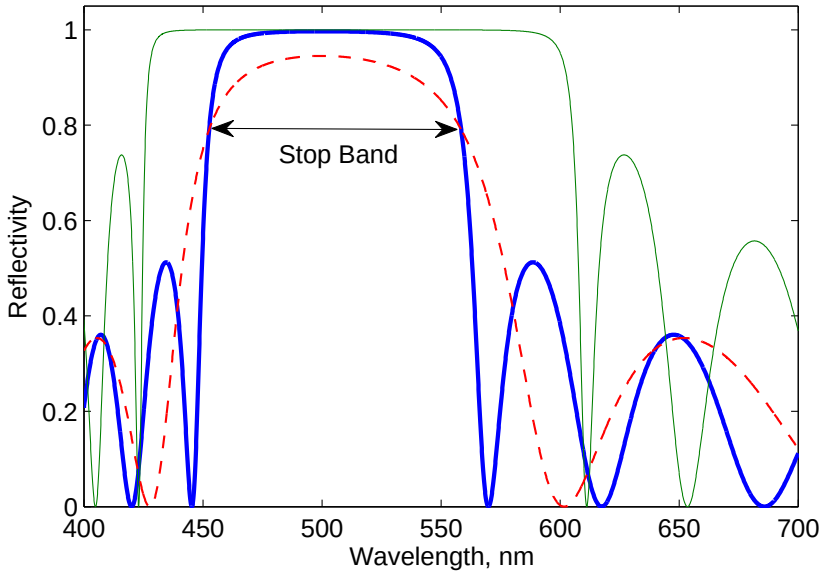

Figure 2.3: Reflectivity of DBR mirrors of different designs. Each is designed to be centered at 500nm. Keeping the same number of periods but increasing the index contrast increased the stopband width and maximum reflectivity of the light green curve ($R_{max} \approx 0.99998$) over the dark blue $R_{max} \approx 0.997$. The dashed red curve is a simulation with the same index contrast as the blue curve, but fewer periods.

the phase. Thus, if one elects to vary the cavity thickness, the Fabry-Pérot resonance may shift less (or more) than expected from just the cavity thickness change alone.

The second frame of Fig. 2.4 shows the index profile (blue line) of the simulated structure, which is a 10.5 period DBR surrounded by air. The sharp steps in the index profile indicate interfaces between materials. The green curve shows the normalized electric field magnitude at the simulated design wavelength. One can see that the intensity of 500 nm light decreases for each period of the DBR deeper into the structure from the illuminated side.

Such DBR structures can be produced by a large number of fabrication techniques. One technique employed regularly is molecular beam epitaxy (MBE), which is also used to grow high quality semiconductors in general. This technique is used by Loren Pfeiffer and Ken West at Princeton to produce the AlAs/GaAs samples which have been studied extensively

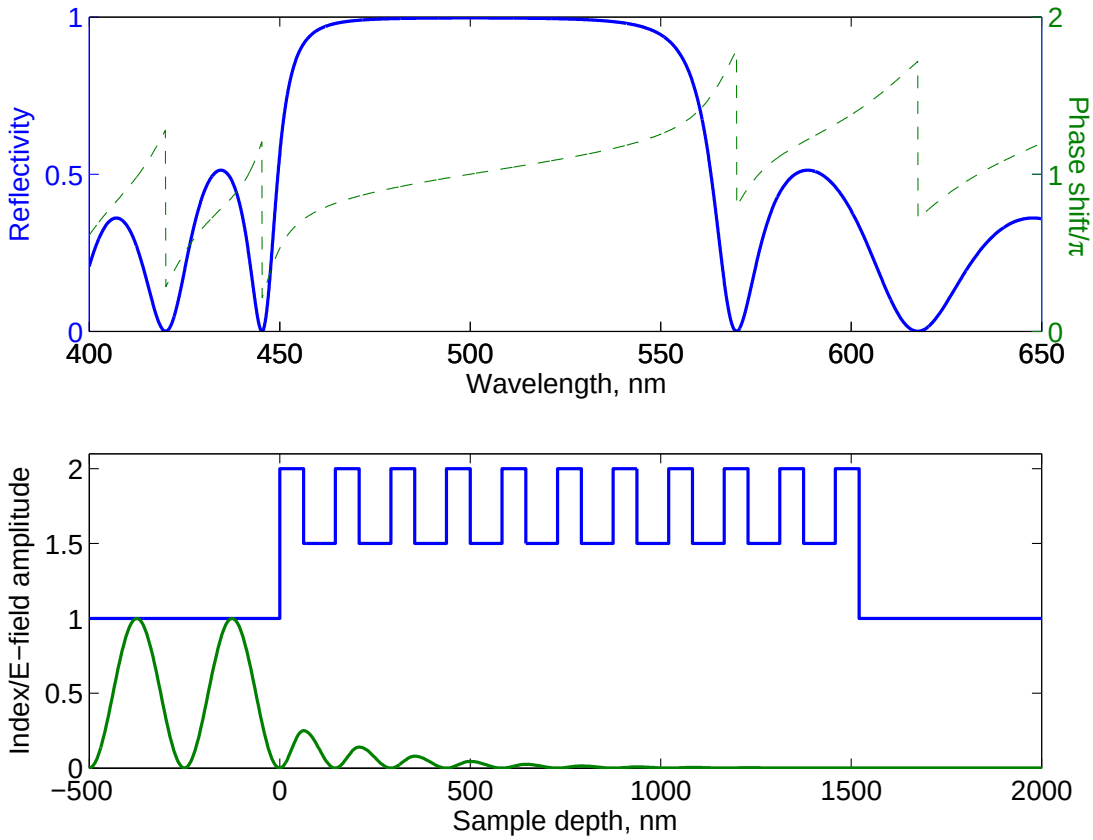

Figure 2.4: Top: Reflectivity and phase shift induced at a DBR. Bottom: index profile of the structure and electric field intensity.

in our group. However, this technique is costly, and requires materials to be lattice matched since crystalline films are deposited. MBE can be used to make monolithic AlGaAs DBRs and microcavities of very high quality, because those materials are almost perfectly lattice matched, but the low bandgap of GaAs limits its usefulness to the infrared range in optics. Other material systems, such as AlN/GaN and ZnSe and its material family, exhibit large lattice mismatch, making it difficult to grow high quality semiconductor visible range DBRs using MBE.

Thus we have begun to explore other deposition techniques, such as plasma enhanced

chemical vapor deposition (PECVD), to deposit amorphous dielectric films. These films can be deposited on almost any surface, so we can in principle embed any excitonic medium in a microcavity. Typically, this requires some creativity to structurally support the excitonic material in the submicron-scale cavity layer while removing any substrate to apply mirrors directly to the cavity. Such techniques are developed and employed in this dissertation for the development of room temperature samples.

2.2.2 THE FABRY-PÉROT INTERFEROMETER

The simplest case of a cavity (or Fabry-Pérot interferometer) is a slab of material between two mirrors. We will assume that the mirrors are partly transmitting with reflectivity R_m. Let us consider light of wavelength λ entering the cavity from the left. As the light propagates across the cavity it will pick up a phase factor according to the index of the cavity n and its thickness l. Strictly speaking, reflection from the mirror will impose a phase shift on the light, which may be wavelength dependent. Finally, the light propagates backwards and reflects from the first mirror. This is seen in Fig. 2.5. Thus the total phase accumulated over a round trip is:

$$\phi = 2\phi_M + \phi_F + \phi_B = 2\phi_M + 2\frac{2\pi nl}{\lambda}. \tag{2.13}$$

The light will continue to cycle through the cavity, picking up this phase each time and interfering with itself. In order for the optical mode to be allowed in the cavity, the total phase must be a multiple of 2π. For simplicity, we will assume that the phase shift induced at the mirror is precisely π. This is the case at the design wavelength of a DBR, and approximately the case for metallic mirrors. When we make this assumption, we can design a cavity thickness to support a particular wavelength by selecting

$$l = \frac{m\lambda}{2n} = \frac{m\lambda'}{2}, \tag{2.14}$$

where λ' is the wavelength of light in the material.

With each full cycle of the cavity, some light is transmitted through the second mirror with an advanced phase and amplitude reduced relative to the previous iteration due to two

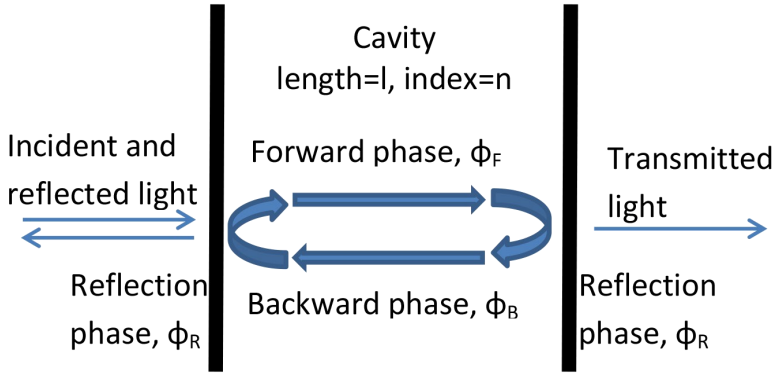

Figure 2.5: Accumulation of phase during round trip of light in a cavity. Solid black lines indicate mirrors. When the round trip phase accumulation is a multiple of 2π, that mode will constructively interfere and its transmission will be enhanced. Here ϕ_F and ϕ_B are the same, but I specify both to emphasize the fact that the round trip must be accounted for.

reflections. After summing an infinite series of these transmissions and accounting for the interference effects, one can compute the transmission as a function of the phase shift per cycle,

$$T = \frac{1}{1 + 4\frac{(R_M)}{(1-R_M)^2}(\sin^2(\phi/2))}. \tag{2.15}$$

What arises from this interference is a strongly wavelength-dependent transmission. At modes λ_m where a cavity mode is permitted, the transmission through the structure increases dramatically. Figure 2.6 shows the calculated transmission through two cavities of identical thickness but differing mirror reflectivities. The spacing between adjacent λ_m is called the free spectral range (FSR) and depends entirely on the optical thickness of the cavity. It follows from Eq. (2.14) that there will be a series of resonant wavelengths with spacing:

$$FSR = \lambda_{m+1} - \lambda_m = \frac{2nl}{m+1} - \frac{2nl}{m}. \tag{2.16}$$

Closer inspection of the state spacing shows that the series of resonant wavelengths is evenly spaced in energy. It is a cruel twist that history has elected to characterize FP modes

according to wavelength when frequency, energy or wavenumber would be a more natural scales. Thus in the future section on cavity photons, we will focus primarily on the energy, rather than the wavelength of the mode.

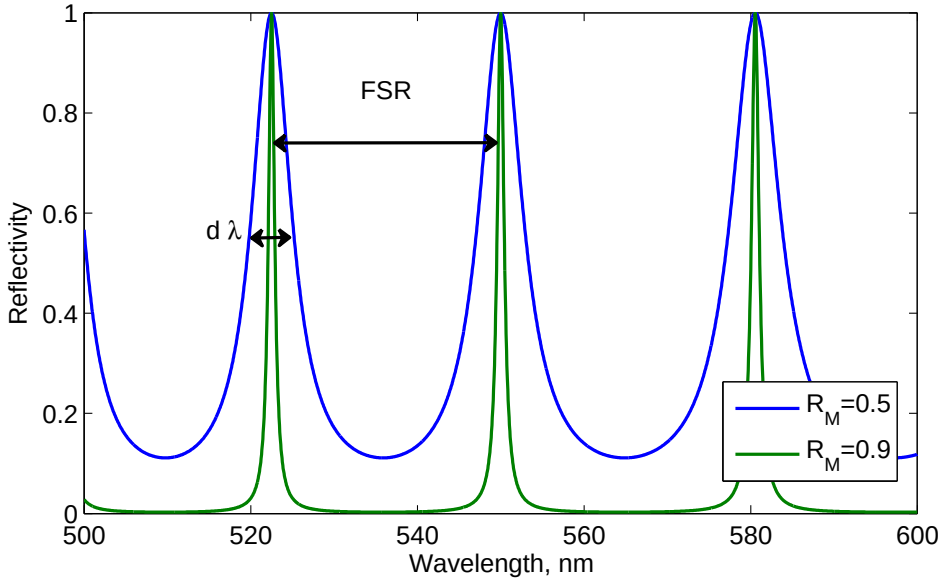

Figure 2.6: Transmission through a planar cavity of length L and given mirror reflectivity. FSR is the free spectral range, i.e. the separation between adjacent modes. $d\lambda$ is the FWHM of the mode.

Also apparent in Fig. 2.6 is dλ, the full width at half maximum (FWHM) of the individual cavity modes. This linewidth depends on not only the thickness of the cavity, but also the reflectivity of the mirrors. This is why the higher reflectivity cavity (green curve) exhibits much narrower linewidths. Often times cavities are quantified by the *finesse*:

$$\mathcal{F} = \frac{FSR}{d\lambda} \tag{2.17}$$

The constructive interference of light in the cavity greatly enhances the electric field amplitude of that cavity mode at particular positions. By placing excitonic materials, such as quantum wells, at the antinodes of this optical mode, the exciton couples more strongly to light than outside the structure. This can be described as the Purcell effect, where an enhancement (reduction) of the density of states of the vacuum optical mode can enhance (suppress) the rate of emission from a transition.

2.2.3 CAVITY PHOTON PROPERTIES

Once a cavity has been assembled from mirrors (either dielectric or metallic) and a cavity spacer (possibly an air gap or transparent material), we can characterize the emergent Fabry-Pérot resonance in terms of its linewidth and lifetime. Here we will refer to the single particle optical mode inside the cavity as a cavity photon. From the standpoint of a particle in the cavity, one can consider this as a well defined particle with an energy and a lifetime–as the photon leaks out of the cavity with an exponential decay. The photon can propagate in plane, with a mass dependent on its confinement energy. Here we will outline how these properties relate to the cavity geometry and confining mirrors.

The cavity mode resonance can be well understood as a damped harmonic oscillator. In the classical picture, the electric field can oscillate with a certain resonant frequency (and higher harmonics thereof), and the amplitude of that oscillation decreases exponentially. The damping is connected to the reflectivity of the cavity mirrors and geometry: higher reflectivity mirrors and increased cavity lengths will increase the photon lifetime.

2.2.3.1 Q-FACTOR AND LIFETIME

A good way to characterize a cavity is with the Q-factor, the ratio of the stored energy to the power dissipation:

$$\mathcal{Q} = 2\pi \frac{E_s}{\Delta E} = \omega_0 \frac{E_s}{P_d} \tag{2.18}$$

where E_s is the energy stored in the oscillator, ΔE is the energy dissipated per oscillation cycle, ω_0 is the angular frequency of the resonance, and P_d is the power dissipation. By considering the power dissipation of a cavity, we can assign a lifetime γ to the m^{th} cavity mode. Just as in the general case of damped oscillators, we expect energy to be dissipated exponentially:

$$E_s(t) = E_0 e^{-\gamma t} \tag{2.19}$$

such that the power dissipation is

$$P_d(t) = \gamma E_0 e^{-\gamma t} = \gamma E_s(t), \tag{2.20}$$

which can be used with Eq. (2.18) to show

$$\mathcal{Q} = \omega_m \frac{E_s(t)}{\gamma E_s(t)} = \frac{\omega_m}{\gamma}. \tag{2.21}$$

Furthermore, we can relate the Q-factor to the geometry of the cavity by considering the loss as due to transmission through a mirror. In a cartoon picture we can consider a pulse of light bouncing back and forth between the mirrors. Every time the light bounces, photons escape the cavity with probability of $1 - R_M$. We note that the energy density of the light is proportional to the number of photons $E_s = n\hbar\omega_m$, so each bounce reduces the stored energy $E_f = R_M E_i$. In this picture where the cavity is large, we average the stored energy and energy dissipated over the time it takes light to complete half a round trip. The average energy is

$$\overline{E_s} = \frac{E_i + E_f}{2} = \frac{(1 + R_M)}{2} E_i, \tag{2.22}$$

and the energy lost is

$$\Delta E_{trip} = (1 - R_M) E_i. \tag{2.23}$$

The time of flight for this half round trip is

$$\Delta t_{trip} = nl/c = m\lambda_m/2c. \tag{2.24}$$

The energy lost per *oscillation* cycle is therefore going to be the energy lost per half round trip times the ratio of oscillator to trip periods. The oscillation period is $\Delta t_{cycle} = \frac{2\pi}{\omega_{cycle}}$ such that

$$\Delta E_{cycle} = \Delta E_{trip} \frac{\Delta t_{cycle}}{\Delta t_{trip}} = \frac{2(1 - R_M)}{m} E_i, \tag{2.25}$$

and with the definition of $\mathcal{Q}$ in Eq. 2.18:

$$\begin{aligned}\mathcal{Q} &= 2\pi \frac{(1+R_M)E_i/2}{2(1-R_M)E_i/m} \\ &= \frac{m\pi}{2}\frac{1+R_M}{1-R_M}.\end{aligned} \tag{2.26}$$

Thus, we see that the Q-factor is directly determined by the cavity mode order and mirror quality. Comparing to Eq. (2.21), the inverse photon lifetime can be directly related to cavity design:

$$\begin{aligned}\gamma &= \frac{2\omega_m}{m\pi}\frac{1-R_M}{1+R_M} \\ &= \frac{2\omega_1}{\pi}\frac{1-R_M}{1+R_M}\end{aligned} \tag{2.27}$$

where ω_1 is the fundamental mode of the cavity. Thus all cavity photons have the same lifetime which is fixed by the cavity length and mirror quality. However, one can increase the lifetime of a desired wavelength by making the cavity longer such that the desired mode is a higher order. The tradeoff that one makes by doing so is the reduction of the free spectral range.

2.2.3.2 RESONANCES: LINEWIDTH AND LIFETIME

As mentioned previously, each cavity mode can be treated as a damped harmonic oscillator. In analogy to the position vector of a mechanical oscillator, the electric field will oscillate and damp out as it leaks out of the cavity. However, we are generally interested in the intensity of the light field to quantify a measurable decay, since the number of photons is proportional to the intensity. Thus we are interested in the energy response of the damped mechanical oscillator.

Starting from the classical definition of a damped, driven harmonic oscillator [54], we will relate the solution to both the frequency response and energy decay of the oscillations. The equation of motion of the oscillator is

$$\ddot{x} + \tilde{\gamma}\dot{x} + \omega_0^2 x = A\cos(\omega t + \theta_0). \tag{2.28}$$

Here ω_0 indicates the natural frequency of the undamped oscillator, ω is the frequency of the driving field, A is the amplitude of the driving field, and θ_0 is its absolute phase, included for generality. x indicates the position of the oscillator, which could be the position of a mass in a mechanical oscillator, or in our case the E-field. $\tilde{\gamma}$ defines the damping of the system.

The proper solution to the motion, $x(t)$, requires homogeneous and inhomogeneous parts. We must include two constants of integration $\mathcal{A}$ and $\mathcal{B}$ to allow for initial motion:

$$x(t) = \mathcal{A}e^{-\tilde{\gamma}t/2}\cos(\tilde{\omega}t + \mathcal{B}) + \frac{A\cos(\omega t - \phi)}{\sqrt{(\omega_0^2 - \omega^2)^2 + \tilde{\gamma}^2\omega^2}}. \tag{2.29}$$

This solutions include ϕ, the phase delay between the oscillator and the driving force. This is dependent on the frequency detuning between the drive and resonance, and is calculated as

$$\phi = \tan^{-1}\left(\frac{2\tilde{\gamma}\omega}{\omega_0^2 - \omega^2}\right) \tag{2.30}$$

so that exactly at resonance the phase delay is exactly 90^o. The first term in Eq. (2.29) is the homogeneous solution and accounts for the decay of transients in the system. The exponential decay of this equation is precisely the source of the photon lifetime in a cavity. Strictly speaking, transient oscillations will occur at frequency $\tilde{\omega} = \sqrt{\omega_0^2 - \tilde{\gamma}^2}$, but as we will focus on cases with low damping, this shift is negligible.

Since we are interested in the decay of photon number from a cavity, we will consider a populated cavity with no external driving field. Thus we can ignore the second term in Eq. (2.29). Since number is proportional to the energy stored, we work with the square of the amplitude response envelope:

$$\begin{aligned} n(t) &\propto (\mathcal{A}_0 e^{-\tilde{\gamma}t/2})^2 \\ n(t) &= n_0^2 e^{-\tilde{\gamma}t}. \end{aligned} \tag{2.31}$$

Here n_0 is the initial number of photons in the cavity. Thus we see that the damping parameter $\tilde{\gamma}$ is precisely the decay constant of a cavity photon, γ, and thus we will use the latter for the rest of these calculations.

Now, let us examine the frequency response of the oscillator. Experimentally, one could probe the response of an oscillator by optical absorption, reflectivity, or luminescence, for

example. Spectrally resolving light emitted from a simple resonance, for example, will reveal an intensity vs frequency profile that is precisely the square of the amplitude of the second term in Eq. (2.29).

$$I(\omega) = \frac{A^2}{(\omega_0^2 - \omega^2)^2 + \gamma^2\omega^2}, \tag{2.32}$$

which we can simplify for the case of moderately high $\mathcal{Q}$ oscillators, specifically since $\gamma/\omega_0 \to 0$ and since the envelope is strongly peaked around $|\omega_0 - \omega| \sim \mathcal{O}(\gamma)$. Immediately around the resonance at ω_0, Eq. (2.32) simplifies to

$$I(\omega) = \frac{(A/\omega_0)^2}{4(\omega_0 - \omega)^2 + \gamma^2}. \tag{2.33}$$

This is exactly the form for a Cauchy distribution (also known as a Lorentz distribution). Such a distribution is characterized by a central value, namely ω_0, and a FWHM of γ, which is exactly the lifetime of a photon in the cavity. Thus, it becomes a straightforward process to experimentally determine the lifetime of a simple particle, as long as one can accurately measure its frequency response. Of course, particles that are interacting amongst themselves or with external disorder can exhibit further broadening of luminescence or absorption spectra. Thermally excited particles will likewise result in luminescence spectra that are further broadened. Nevertheless, a simple FWHM measurement of a spectrum can at least imply an upper limit on the lifetime ($\tau = 1/\gamma$) of a particle.

2.2.3.3 PHOTON MASS

Perhaps a surprising, yet straightforward result of cavity optics is that the normally massless photon gains a mass. We can understand this directly in terms of the in-plane momentum of this relativistic particle. In general, the photon obeys the relativistic dispersion relation:

$$\begin{aligned} E &= \frac{\hbar c}{n}\sqrt{k_x^2 + k_y^2 + k_z^2} \\ E &= \frac{\hbar c}{n}\sqrt{k_\parallel^2 + k_z^2}, \end{aligned} \tag{2.34}$$

where E is the energy of a photon, and we have selected the z-direction as the confinement axis of the cavity, leaving $k_{\|}$ to be the in-plane wavevector. Here we will assume that k_z is strictly determined by the normal incidence mode. While in principle the mirror reflectivity and phase will change with angle of incidence, we will ignore that here, as we choose to consider small values of $k_{\|}$.

Perturbations about the $k_{\|} = 0$ point will result in parabolic curvature of the photon dispersion. Namely we can define the effective mass to be

$$m^* = \hbar^2 / \frac{\partial^2 E}{\partial^2 k_{\|}}\bigg|_{k_{\|}=0}$$
$$m^* = \frac{E_0}{c^2/n^2}. \tag{2.35}$$

For optical cavities with index values on the order of 1.5-3.5, typical values for the mass are on the order of $1 - 3 \times 10^{-5} m_0$.

It should be noted that calculations of oblique incidence transmission through a real structure is not entirely straightforward. Exact calculations (see appendix A) must take into account the polarization of light, since the continuity conditions on $\vec{E}$ and $\vec{B}$ require using different formulations of the Fresnel equations at boundaries. At large angle, there will be a slight splitting between the TE and TM (transverse electric and transverse magnetic) modes. For typical angles observed in polariton experiments, this splitting may be on the order of up to $100 \mu eV$.

2.2.3.4 E-FIELD DISTRIBUTION

Again, the principal objective with our cavities is to strongly couple an excitonic resonance to the enhanced electric field inside the cavity. Thus, we must take care when designing a microcavity polariton structure to properly place the excitonic structures at the antinodes of the cavity. The transfer matrix procedure outlined in Appendix A is particularly useful for designing real structures.

Based on the index contrast between the cavity material and the bulk cavity layer, we can expect the E-field mode magnitude to exhibit even- or odd-symmetry, just as a standing wave

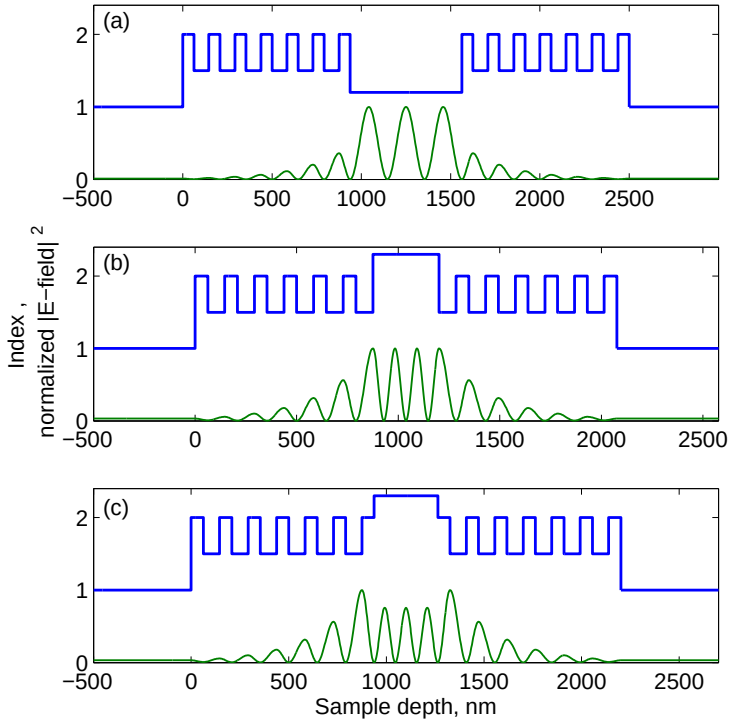

Figure 2.7: E-field mode profiles are plotted in three cavities designed to be $3\lambda/2$ at 500 nm. Bold, blue lines indicate the index structure of the sample, while the fine, green lines are the E-field mode intensity at 500nm. In each case, the DBRs are either 6 or 6.5 periods with layer refractive indices of 1.5 and 2 for the two materials. The cavity index is either lower or higher than both DBR materials. In frame (a), the cavity index is low, and so we elect to terminate the DBRs with the high index material to optimize index contrast. The cavity mode is even, with nodes located at the ends of the cavity. Frame (b) shows a cavity designed with a high index material. Here we terminate the DBRs with a low index layer to optimize contrast, but the boundary conditions force the cavity mode to be even, having antinodes at the boundaries. Frame (c) also uses a high index spacer, but also terminates the DBR with a high index layer. Because of the boundary conditions, the first high index layers of the DBR act like part of the cavity, and much of the E-field is lost into those layers.

on a string can be either even or odd according to the boundary conditions. Specifically, consider the three cavities in Fig. 2.7 (a)-(c). Here we have designed three very similar idealized cavities. Frame (a) shows an odd-symmetry cavity which has nodes at the mirror surfaces, while frame (b) shows an even-symmetry mode with antinodes at the cavity ends. This nuance in the design can be very relevant in placing the active materials within a cavity. For example, when using MBE grown QWs, it can be difficult to place QWs very close to the surface of the structure due to the need for barriers and buffer layers.

Furthermore, Fig. 2.7 (c) exemplifies that unexpected mode features can arise due to interference effects within a structure. In this case, we used a high index DBR layer adjacent to the high index cavity layer. At first glance, this appears to have shifted the cavity mode to an odd-type symmetry, but careful analysis shows that those final DBR layers are actually part of the cavity. The E-field penetrates significantly into those DBR layers, resulting in slightly lower than expected reflectivity and a longer than expected cavity. The loss of mode volume into regions not containing excitonic material is not ideal for strong coupling, as highlighted in Eq. (2.10), since the extension of the mode volume decreases $\vec{E}_0$ at the QWs.

Ultimately, any structure based on QW exciton resonance should be carefully simulated in order to place the thin QWs at the antinodes of the E-field mode. In many cases, multiple QWs can be placed at each antinode in order to increase the Rabi splitting and reduce the saturation density. This is particularly effective if the thickness of the well plus that of the barriers is less than the wavelength of light in the cavity material. That way multiple wells will fit into a region where the E-field profile is relatively uniform at its maximum value.

2.3 STRONG COUPLING AND POLARITON STATES

Finally, after fully developing optical cavities and excitonic states, we study the coupling between them that leads to the Rabi oscillation and the formation of two new states: the upper and lower polaritons. In this section, I derive the origin of strong coupling, and highlight several important results of the state mixing on the properties of polaritons.

In a semi-classical sense, one can envision the coupling between exciton and photon as

a series of absorptions and emissions on a very fast time scale. The exciton is an electron-hole pair with some finite lifetime to recombine and emit a photon. Inside the cavity, that radiative lifetime is made made much shorter due to the Purcell effect of the enhanced vacuum photon mode [55]. This is because the cavity readjusts the density of states (DOS) for the optical field, and the Fermi's golden rule indicates that this DOS of the final state affects the transition rate. Thus an exciton will only exist for a short time before emitting a photon.

Upon recombination (decay of the exciton) and emission of a cavity photon, that photon will rapidly cycle the cavity, reflecting from the high quality mirrors. As the photon traverses the excitonic material, it has a high probability to be absorbed as an exciton.

The process repeats many times on a timescale governed by the Rabi frequency, $\omega_R = 2g_0/\hbar$, where g_0 is the electric dipole coupling strength between the exciton and photon. An excitation in the system keeps changing: exciton→photon→exciton→photon. When this oscillation rate is much faster than decay rates or decoherence rates of the exciton and photon, this classical description of the oscillation is not valid. Instead, we say that the particle is in a superposition of an exciton and a photon. This superposition is exactly what we call the polariton states.

We typically identify strong coupling by the effect it has on the state energies. When there are two states close in energy and strongly coupled, they will repel each other to form the new polariton states. The most recognizable manifestation of this is an anticrossing or avoided crossing between two states. This is shown later in Figs. 2.8 and 2.9.

It should be noted that in the absence of state broadening and lifetime considerations, *any* coupling is strong coupling. Section 2.3.4 discusses the boundary between strong and weak coupling, and what can cause the anticrossing between coupled states to collapse. Also, coupling can become more complicated when there is disorder or multiple states coupling to the optical mode. In Section 2.3.5, I show that multiple exciton states can overlap with the photon to result in more complicated states than the upper and lower polaritons alone. This can be relevant in the case of excitons originating in different bands (such as the light and heavy hole excitons in GaAs) or it can arise in the case of disordered quantum wells.

2.3.1 COUPLED OSCILLATOR MODEL

Recognizing that the exciton will strongly interact with the photon mode of a cavity via the electronic dipole transition, $-\vec{\mu}\cdot\vec{E}$, as discussed in Section 2.1.3, we can express the Hamiltonian of the coupled system. Working in the second quantized formalism the interacting Hamiltonian is

$$\hat{H}_p = E_c(\hat{a}^\dagger\hat{a}) + E_x(\hat{b}^\dagger\hat{b}) + g_0(\hat{a}^\dagger\hat{b} + \hat{b}^\dagger\hat{a}), \tag{2.36}$$

where $\hat{a}^\dagger$ ($\hat{a}$) is the creation (annihilation) operator for a cavity photon, and with the corresponding $\hat{b}$ terms for the excitons. E_c and E_x are the energies of single cavity photons and excitons respectively. All of these terms are inherently dependent on in-plane wavevector, so there is an implied sum over all k-states. The coupling strength is quantified by g_0 and scales with the exciton oscillator strength and the optical density of states. The operators in the interaction term embody the physics of the cyclic absorption and reemission of the photon: $\hat{a}^\dagger\hat{b}$ will annihilate an exciton and create a photon, while $\hat{b}^\dagger\hat{a}$ annihilates the photon to create an exciton. This process conserves in-plane momentum (i.e. wavevector) in a bulk or QW type structure, so each $k_\parallel$ state of the exciton and photon couple separately.

We can diagonalize this Hamiltonian as demonstrated by J.J. Hopfield in 1958 [56]. The diagonalization transforms the independent exciton and photon states into two new superpositions: the upper and lower polaritons (UP and LP). These are the proper eigenstates of the coupled Hamiltonian, with the redesigned creation and annihilation operators for the lower ($\hat{c}_{LP}$) and upper ($\hat{c}_{UP}$) polaritons:

$$\hat{c}_{LP} = X\hat{b} + C\hat{a} \tag{2.37}$$

$$\hat{c}_{UP} = -C\hat{b} + X\hat{a} \tag{2.38}$$

$$\hat{H}_p = E_{LP}(\hat{c}^\dagger_{LP}\hat{c}_{LP}) + E_{UP}(\hat{c}^\dagger_{UP}\hat{c}_{UP}), \tag{2.39}$$

where X and C are the so-called Hopfield coefficients and indicate the character of a given polariton state. $|X|^2$ thus is the fraction of the lower polariton which is excitonic while $|C|^2$

measures the photon fraction. These assignments are reversed for the upper polariton. The energies of these states are

$$E_{LP}^{UP} = \frac{1}{2}(E_x + E_c) \pm \frac{1}{2}\sqrt{(E_c - E_x)^2 + 4g_0^2}, \tag{2.40}$$

while the Hopfield coefficients are calculated as

$$\begin{aligned} |X|^2 &= \frac{1}{2}(1 + \frac{(E_c - E_x)}{\sqrt{(E_c - E_x)^2 + 4g_0^2}}) \\ |C|^2 &= \frac{1}{2}(1 - \frac{(E_c - E_x)}{\sqrt{(E_c - E_x)^2 + 4g_0^2}}). \end{aligned} \tag{2.41}$$

2.3.2 POLARITON PROPERTIES

The strongest indicator of strong coupling is the avoided crossing observed in the spectrum of polariton systems. If there is a way to continuously vary the energy difference between the exciton and photon in the system, then one may observe two energy states appear to avoid each other. For example by sweeping the thickness of the cavity or changing the observation angle, we can effectively change the photon energy that we are observing, while the exciton energy will remain essentially constant. Real samples will exhibit thickness variations, either intentional or accidental, and those thickness variations will tend to impact the strongly confined and very light photon much more than the exciton. Likewise, by changing the angle of observation, we are using the drastic difference between the exciton and photon masses to observe different detunings.

Figures 2.8 and 2.9 show typical avoided crossings that we regularly observe in both real and k-space. Due to anisotropy of the MBE growth process, our samples often exhibit noticeable dependence of the normal incidence ($k_{\|} = 0$) cavity energy on position, which leads to a gradient in the photon energy. From Fig. 2.8, one can see that far away from $x = 0$, the states approach the pure exciton and photon energies. However, close to the resonant point ($x = 0$), the states appear to repel each other.

Similar to the avoided crossing observed in real space, an avoided crossing can be observed in momentum space when the $k_{\|} = 0$ cavity energy is lower in energy than the exciton, since at higher angle the photon will become resonant to the exciton. Figure 2.9 is calculated

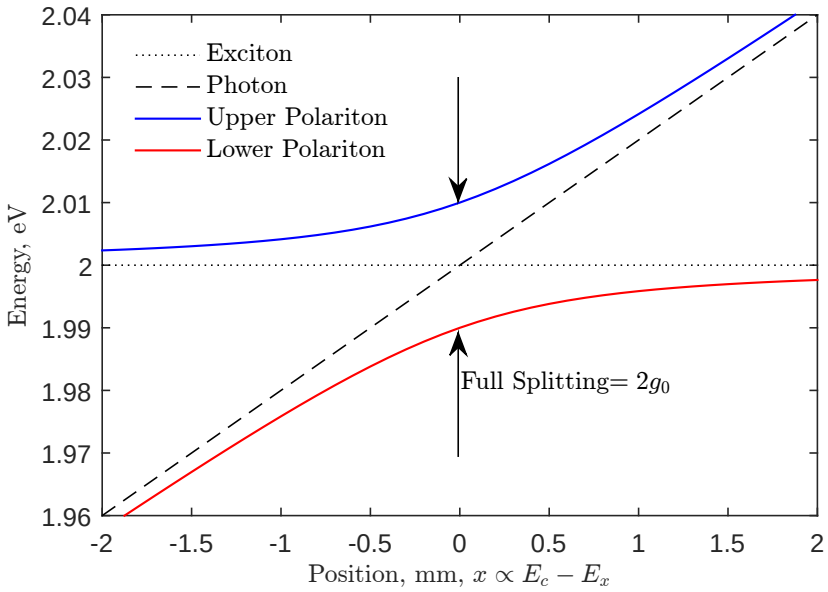

Figure 2.8: Polariton avoided crossing vs exciton-photon detuning for $k_{\|} = 0$. We often observe dependence of the photon energy on position due to variation in the cavity thickness. The polariton is often characterized by the minimum UP-LP splitting. This occurs at resonance and the full splitting is equal to twice the coupling constant, g_0, in Eq. (2.36).

for a case where the photon is slightly red-detuned from the exciton to highlight this. Near $k_{\|} = \pm 4\,\mu\text{m}^{-1}$, the bare exciton and photon energies cross, but we see the lower and upper polaritons avoid the crossing. It is a general result that the lower polariton is more excitonic at higher angles, and more photonic close to $k_{\|} = 0$.

The calculation in Fig. 2.9 would correspond to the k-resolved spectrum at a single spatial point, which would correspond to $x = -1.5\,\text{mm}$ in Fig. 2.8. Thus, if we generate a population of polaritons at this spatial point, the k-space spectrum shows the allowed states at that point. Since we have a wide range of k-states well above $|k_{\|}| = 4\,\mu\text{m}^{-1}$, we note that there are many exciton states available to be populated. Essentially, any very-high k lower polariton is effectively an exciton. These excitons can in principle serve as a particle reservoir and a population to scatter with. This exciton reservoir can be important either if it a thermally equilibrated component of a long lifetime system or if it is a dynamic

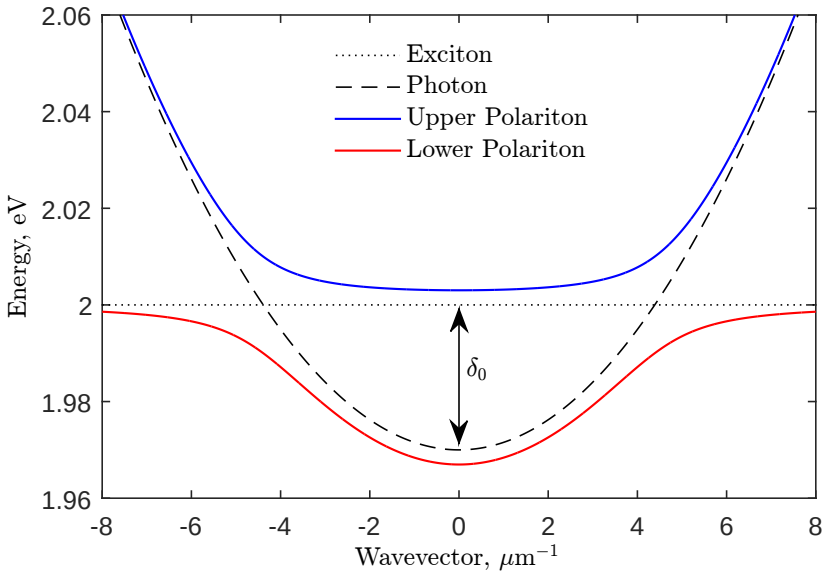

Figure 2.9: Polariton avoided crossing vs $k_\parallel$. The very light mass of the photon means that at a single point on the sample we can observe the mixing of a variety of angles. At this spatial point the polariton is negatively detuned with $\delta_0 = E_c(k_\parallel = 0) - E_x(k_\parallel = 0) = -3\,\text{meV}$

non-equilibrium population in transient or steady state conditions.

2.3.3 DETUNING DEPENDENT PROPERTIES

A benefit of the spatial variation in the photon energy in Fig. 2.8 is that it enables us to actively change the character of the polaritons as we study them. Typically, we quantify the character of a polariton under certain conditions by its detuning relative to the Rabi coupling of the system. Detuning is defined as the energy difference between the exciton and cavity photons, usually at normal incidence:

$$\delta_0 = E_c(k_\parallel = 0) - E_x(k_\parallel = 0). \tag{2.42}$$

Thus, negative detuning corresponds to conditions under which the lower polariton is more photon-like at $k_\parallel = 0$, while the upper polariton is more exciton-like. Exactly at $\delta_0 = 0$, the

uncoupled exciton and photon are degenerate and the upper and lower polariton are equal mixtures of exciton and photon.

The extent to which the polariton is exciton- or photon-like depends on the detuning relative to the coupling strength. These ratios are the Hopfield coefficients, already calculated in (2.41). As stated before, $|X|^2$ ($|C|^2$) represents the percent that the lower polariton is excitonic (photonic). (These assignments are reversed when considering the character of the upper polariton.) When considering the polariton as a superposition of the exciton and photon, these coefficients are the magnitudes-squared of those superposition amplitudes. These coefficients are strictly dependent on in-plane wavevector, as E_c and E_x are inherently k-dependent. Therefore, it is normal for the character of polaritons to be dependent on both position and momentum.

The implication of varying polariton character is that we can tune many parameters of the polariton, including mass, interaction strength and lifetime. Since the exciton is more massive, more strongly interacting, and usually longer lived than the cavity photon, we can expect these parameters to increase with increasing excitonic fraction.

Specifically, the effective mass of the polariton (valid over a small range of $k_\parallel \sim 0$) can be calculated directly from the Hopfield coefficients:

$$\begin{aligned} \frac{1}{m_{LP}} &= \frac{|X|^2}{m_x} + \frac{|C|^2}{m_c} \\ \frac{1}{m_{UP}} &= \frac{|C|^2}{m_x} + \frac{|X|^2}{m_c}. \end{aligned} \tag{2.43}$$

Here m_{LP}, m_{UP}, m_x, m_c are the masses of the lower polariton, upper polariton, exciton and photon respectively. Since the exciton is several orders of magnitude heavier than the photon, we can approximate the polariton mass (for $|C|^2 >> 0.01$):

$$\begin{aligned} m_{LP} &\approx \frac{m_c}{|C|^2} \\ m_{UP} &\approx \frac{m_c}{|X|^2}. \end{aligned} \tag{2.44}$$

From this, it is easy to see that at resonant detuning, the polariton mass is only twice that of the bare photon.

The lifetime of the polariton can be predicted from the intrinsic decay rates of the exciton and photon. If we recognize that particle decay rate can be included as an imaginary component of its energy, we can include exciton and photon lifetimes in (2.40). The resulting imaginary components of E_{LP} and E_{UP} are precisely the decay rates of those species. Again, since the decay rate of the photon (leakage through mirrors, $\tau \sim 1 - 10\,\text{ps}$) is so much faster than that of excitons (nonradiative decay, $\tau \sim 1\,\text{ns}$) we expect all but the most excitonic polaritons to have lifetimes fixed by the cavity quality. At resonance, the polariton will have a lifetime twice that of the photon.

Similarly, scattering rates and thermalization times will be dependent on the Hopfield coefficients. Since we assume that all interparticle interactions are mediated by the excitonic fraction of the polariton, these scattering rates should scale with $|X|^2$. Very excitonic polaritons should scatter rapidly and thermalize well, while very photonic particles will scatter more weakly and possibly not thermalize at all, just as photons themselves cannot thermalize.

2.3.4 WEAK COUPLING

Based on (2.40), we would always expect strong coupling between excitons and photons. Indeed, in an infinite crystal with no scattering or losses, we should expect this. However, we see experimentally that most systems require us to work very hard to achieve strong coupling. High quality mirrors are used to enhance the electric field and contain the cavity photon, and MBE growers painstakingly work to minimize disorder seen by QW excitons.

The reasons for these efforts and the collapse of strong coupling can be easily understood by including the finite lifetime of excitons and photons as an imaginary self-energy. Inclusion of the lifetime in the energy is equivalent to including broadening in the Lorentzian osciallator model in Section 2.2.3.2, and the lifetime of a quantum state puts a limit on its linewidth due to a time-energy uncertainty principle. The fact that an imaginary self energy will result in gain or decay of a quantum state can be seen in applying the time evolution operator: while quantum phase oscillates at the frequency of the energy, the amplitude will increase or decrease according to any imaginary part of the energy.

Therefore, using the explicit assignment of $E \to E + i\Gamma$ where $\Gamma = \hbar\gamma$, we can include

these lifetimes in (2.40). We note that this lifetime includes all processes that cause the exciton or photon to scatter. This includes emission through the cavity mirrors (decay of photon) or scattering from crystal defects or nonradiative recombination of the exciton. Radiative recombination of the exciton does not actually contribute to this lifetime, since that mechanism actually maintains the coherence of the two-state system. By including the imaginary self-energy of the bare exciton and photon, we can extract the imaginary self-energy of the polariton modes as

$$E_{LP}^{UP} + i\Gamma_{LP}^{UP} = \frac{1}{2}(E_x + E_c + i(\Gamma_x + \Gamma_c)) \pm \frac{1}{2}\sqrt{(E_c - E_x + i(\Gamma_c - \Gamma_x))^2 + 4g_0^2}. \tag{2.45}$$

Figure 2.10 demonstrates the collapse of strong coupling due to the lifetime mismatch between coupled states. For this calculation, we assume that one state has infinite lifetime, and increase the decay rate of the other state. As the mismatch between the decay rates increases, the effective Rabi splitting decreases. When the mismatch is equal to twice the coupling strength, the splitting collapses completely. This marks a clear boundary between the strong coupling regime (where the degenerate exciton and photon split to form the UP and LP) and the weak coupling regime (where the real parts of the energies are unchanged).

Here we observe that the distinctly broadened exciton and photon states become more similar in the strong coupling regime. In this case of zero-detuning, both the UP and LP have the same lifetime. This is not the complete story, since the UP has a separate decay channel: coupling between the upper and lower polariton branches presents an efficient mechanism for decay of upper polaritons into lower polaritons, which causes the UP to be intrinsically more broad than the LP.

It is a generic result of Hermitian coupling between states that the real self energy of the states is split while the imaginary energy (decay rate) is 'attracted.' As an aside, anti-Hermitian coupling can be observed, which causes a splitting of the imaginary energy while attracting the real self-energies (see [57]). Anti-Hermiticity relates to the sign of the coupling term between the states. In other words

$$g_0(\hat{a}^\dagger \hat{b} + \hat{b}^\dagger \hat{a}) \rightarrow g_0(\hat{a}^\dagger \hat{b} - \hat{b}^\dagger \hat{a}). \tag{2.46}$$

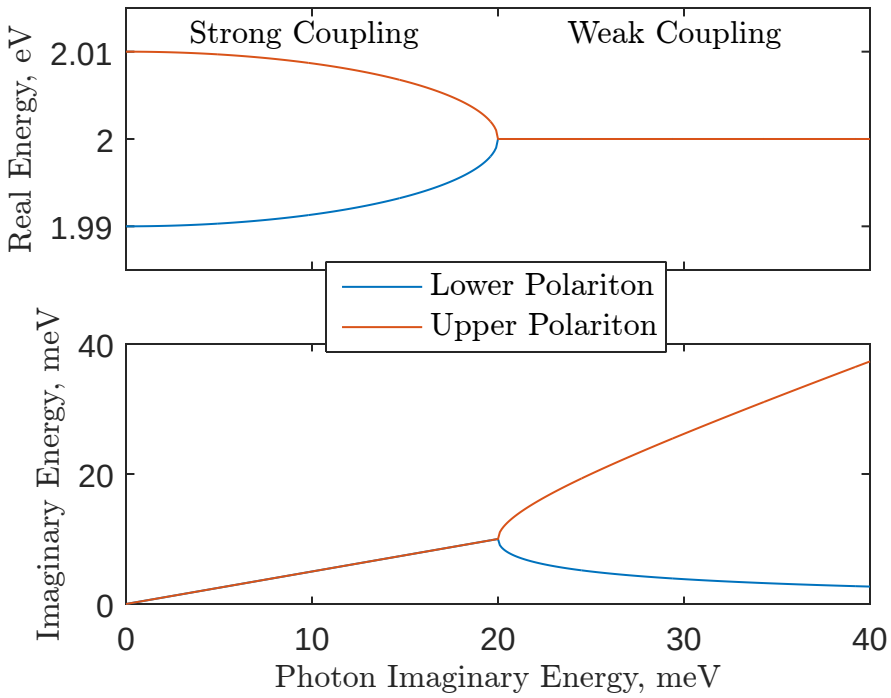

Figure 2.10: The effective splitting between states with mismatched lifetime/dephasing rates is reduced. Here we fix the imaginary component of one state (e.g. the exciton) strictly at zero, while varying the imaginary component of the other (i.e. the photon). The top frame shows the actual splitting in the real energies of these states. When the *difference* between the state decay rates is larger than the coupling $2g_0$ ($g_0 = 10$ meV in this calculation), we see the two energies collapse to form degenerate states. The bottom frame shows that the decay rates of the two states are identical in the strong coupling regime, while they split in the weak coupling regime.

It should be clearly noted that a *difference* in the decoherence rates between the exciton and photon give rise to a reduction of the Rabi coupling. In principle, even if the two states have large but similar Γ, (2.45) predicts that the energies will be split. However, we must also recognize that in order to resolve the UP and LP states, their individual broadenings must be less than the splitting between them. This gives us two requirements for optimally

strong coupling:

$$|\Gamma_x - \Gamma_c| << g_0 \tag{2.47}$$

$$\Gamma_{UP,\text{eff}}, \Gamma_{LP,\text{eff}} << g_0, \tag{2.48}$$

where Γ_{eff} is the broadening of the states including other sources of broadening, such as disorder. The first condition indicates that the effective splitting will be reduced by the lifetime mismatch between the states. The second condition is the requirement that the polariton states be distinguishable. For this condition, I specify the broadening of the polariton states rather than the bare exciton and photon, because the polaritons may interact with disorder and other particles (such as electrons and phonons) differently than bare excitons and holes.

Additionally, the disorder seen by the exciton and photon should be taken into account when assessing the quality of state coupling. While the previous discussions mention particle decay and dephasing as sources of the broadening, QW thickness fluctuations and interactions with background carriers or phonons can greatly broaden the excition state (see Section 2.3.5). Particularly in the case of inhomogeneous thickness fluctuation, one can imagine that a very narrow photon state will strongly couple to the variety of local excitons. Some excitons may be tuned out of resonance with that photon state. In this way, the narrowness of the photon can select out only a fraction of the exciton oscillator strength, effectively reducing the coupling over what one would expect from that many exciton states.

Finally, we point out that these considerations predict that a system can be tuned into and out of strong coupling. In particular, the exciton absorption will saturate at high density, due to the excitons being made up of fermionic electrons and holes, an effect often called phase space filling. Also, background carriers can screen the attraction felt by the excitons and reduce their oscillator strength. Both these mechanisms will decrease the exciton-photon coupling, g_0. The increased scattering between partices can simultaneously increase Γ_x. These conditions combine to work against the conditions in Eq. (2.46) to transition the system into weak coupling. This can drastically blueshift the polariton, and is much of the drive behind using polaritons as optical switches or other nonlinear devices. A series of works treat these saturation effects theoretically [58–60]

2.3.5 COUPLING BETWEEN DISORDERED STATES

The derivation of strong coupling and the polariton states in Section 2.3.1 using (2.36) inherently assumes that there are only two distinct oscillators. In Section 2.3.4 I added the lifetime broadening of these states to show that strong coupling is dependent on the lifetimes of the constituent states being longer than the coupling time. Here I discuss how disorder can impact polariton states.

In any real macroscopic system, we expect a certain amount of fluctuation and disorder. Figure 2.11 demonstrates two very real sources of excitonic disorder in quantum wells. A typical microcavity uses multiple QWs embedded within the cavity spacer to enhance the coupling strength. Traditionally, we use three clusters of QWs (one cluster at each optical antinode) with four QWs per cluster. Thus, Fig. 2.11 (a) shows that these separate QWs can potentially have different thicknesses. This is possible especially if the deposition rate of the MBE growth system changes during deposition. More fundamental to all QW disorder, Fig. 2.11 (b) shows that a single QW can have thickness variation in plane. While MBE growers tend to work very hard to control this sort of fluctuation, it can at least be expected to occur on the scale of single atoms.

Since the confinement potential of a QW is roughly a square well potential, the exciton energy increases approximately as the inverse square of the width. Therefore, similar width disorder in thin wells leads to larger exciton energy broadening than in wide wells.

Also, it should be noted that the observed disorder of excitons can depend on the lateral scale of the disorder in Fig. 2.11 (b). The exciton will average over the well widths on the order of the Bohr radius ($a_x \sim 10\,\text{nm}$ in GaAs QWs), but variation on larger scale can create a potential landscape or trapping for the exciton. Since the wavelength of light in the medium is much larger (~ 200 nm), we expect a photon to see a collection of exciton states. Therefore, if one observes luminescence directly from a QW even with diffraction limited resolution, the disorder will lead to broadening of the exciton emission.

Similarly, when the photon couples to the exciton to form a polariton, that photon will see a range of exciton states. Once the polariton state is formed, it will be spatially larger than a bare exciton. Such 'fluffy' quantum particles can exhibit less broadening than their

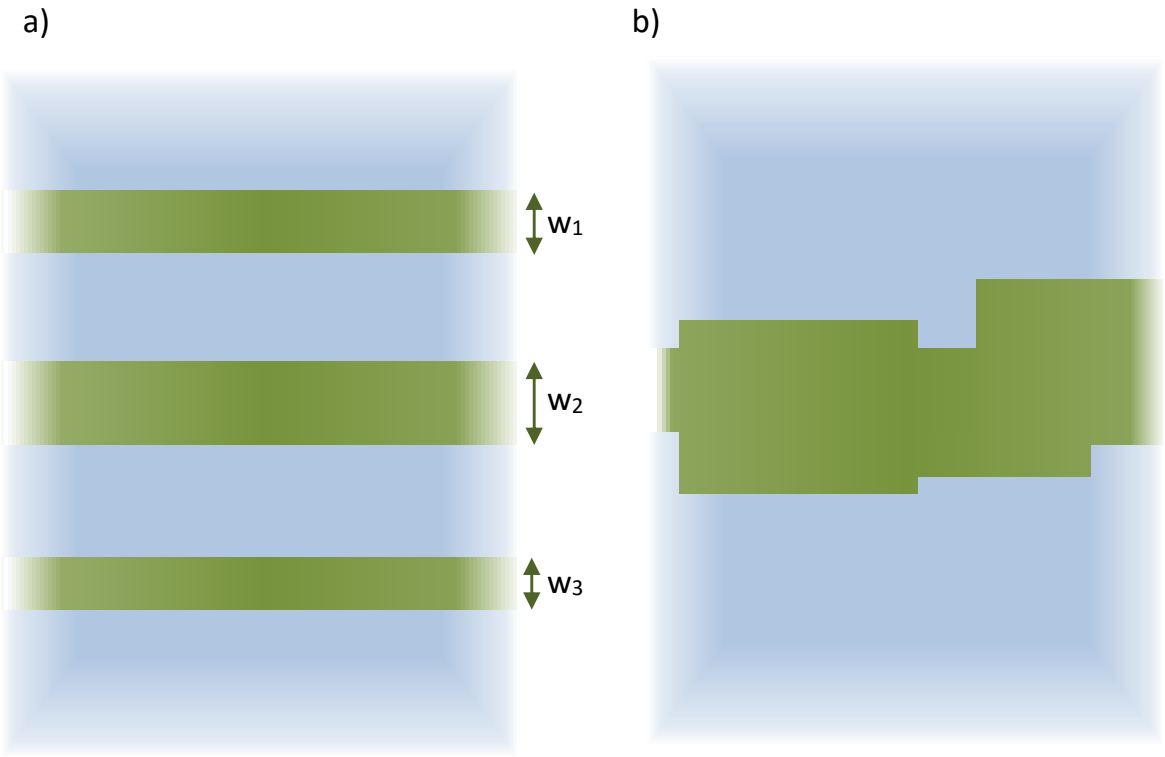

Figure 2.11: Two sources of exciton disorder in microcavity samples. Frame (a) shows that multiple QWs in a sample can have different thicknesses, while frame (b) shows that thickness of a single well can vary.

environment: if the wavefunction of the polariton is larger than the scale of the disorder, the energy disorder of the particle will be narrower than the disordered potential [61][3].

To better understand the impact of disordered exciton states on strong coupling, I will extend the Hamiltonian in Eq. (2.36) to allow for more states. Rather than solve this problem analytically, I will simply diagonalize the Hamiltonian numerically. The two-state

[3]In a historical analogy to NMR systems in which moving atoms average over the spatial fluctuations, such narrowing is sometimes called motional narrowing [62, 63].

Hamiltonian in matrix form can be expressed as

$$\mathbf{H} = \begin{bmatrix} E_x & g_0 \\ g_0^* & E_c \end{bmatrix}, \tag{2.49}$$

which can be extended to include any number of states. Specifically, we will include a second exciton state which also couples to the photon, but the excitons do not directly couple. This would be true of excitons originating in different bands, in separate QWs, and possibly also in different regions of a single QW with varying thickness. Generically, the two states can have different coupling strengths and different energies:

$$\mathbf{H} = \begin{bmatrix} E_{x1} & 0 & g_1 \\ 0 & E_{x2} & g_2 \\ g_1^* & g_2^* & E_c \end{bmatrix}. \tag{2.50}$$

For simplicity, I will assign some common values to these parameters for the following set of calculations:

$$E_0 = 2\,\mathrm{eV} \tag{2.51}$$

$$E_{x1} = E_0 - \Delta \tag{2.52}$$

$$E_{x2} = E_0 + \Delta \tag{2.53}$$

$$g_1 = g_2 = g_0 = 5\,\mathrm{meV} \tag{2.54}$$

or

$$g_1 = g_0 = 5\,\mathrm{meV} \quad \text{and} \quad g_2 = 0 \tag{2.55}$$

First, I will include two degenerate excitons ($\Delta = 0$) in Fig. 2.12. If one of them does not couple to the optical field ($g_2 = 0$, such as in the case of dark heavy hole excitons in GaAs or in the case of a QW placed at the node of the optical field), we recover exactly the solution to Eq. (2.36). This is shown in frames (a) and (c). However, in the case that the second exciton also couples to the optical field (frames (b) and (d)), the splitting between the upper and lower polaritons is enhanced. Physically, having N excitons all coupled to the photon with coupling constant g_0 looks identical to the photon coupled to a single exciton mode with $g_0' = \sqrt{N} g_0$. Interestingly though, we recover the same number of states, so there

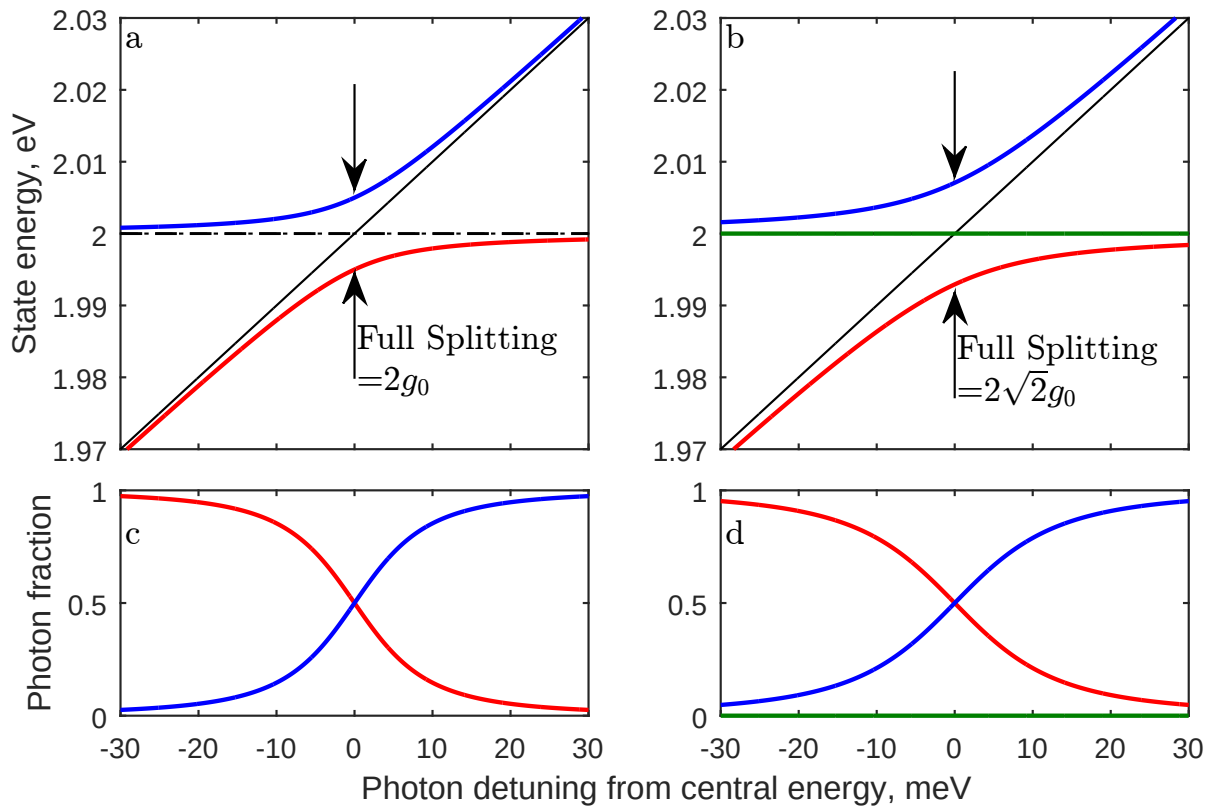

Figure 2.12: A photon mode couples to a single exciton state (frame (a)) or two degenerate exciton states (frame (b)). Solid black line: uncoupled photon; dashed black line: uncoupled excitons; red, blue and green lines: polariton states. Frames (c) and (d) show the photon fraction of the corresponding polaritons states in (a) and (b) respectively. Note the degenerate exciton in (b) enhances the coupling between the UP (blue) and LP (red), but the middle polariton (green) remains completely dark. Frame (b) includes two degenerate exciton lines that are obscured by the green middle polariton line.

will be $N-1$ states residing at the unperturbed exciton energy. This is the green line in frame (b): it looks like an exciton state that is not participating in the coupling. The photon fraction of each polariton state is shown in frame (d), and this middle state has identically zero photon fraction. It is actually still a *delocalized* exciton, having equal contribution from all the degenerate exciton states, but never coupling to the external optical field. It is

invisible to the experimenter, as it cannot emit or absorb light.

Next, I introduce a splitting into the exciton states, shown in Fig. 2.13. The first regime assumes that the splitting is large compared to the coupling (frames (a,c), $\Delta = 5g_0$). Not surprisingly, these two excitons couple separately to the optical field at distinct energy scales. In this case, it makes sense to talk about the lower and upper polaritons for each coupling separately. Indeed, the green polariton state in frame (a) is essentially an upper polariton at negative photon detuning, but transitions into a lower polariton at positive photon detuning, but in either region, one could estimate its properties simply (e.g. according to Eq. (2.41) and related equations) while ignoring the other region. This type of two-resonance coupling is in fact very common in GaAs samples, where the light hole exciton is slightly higher in energy than the heavy hole exciton due to quantum well confinement. Splittings are typically on the order of 10s of meV, which is larger than typical coupling strengths.

The more complicated regime arises when the exciton states are split on the order of the coupling strength. Frames (b) and (d) make the calculation assuming $\Delta = 2g_0$. Unlike the degenerate excitons in 2.12 (b,d), this case shows a middle polariton that has significant photon fraction, meaning it can luminesce and will have a polariton-like dispersion and interactions. The Hopfield coefficients calculated in Eq. (2.41) are not accurate for the upper and lower polaritons in this case, as the photon and excitonic components are mixed between the three polariton states rather than two.

I observed such a middle polariton in samples designed for the optical switching project. This new sample was designed with three QWs (one at each of the three optical cavity antinodes) rather than twelve QWs (four at each of the three optical cavity antinodes). The motivation for this was the reduction of the density at which excitonic absorption saturates. Upon saturation, the oscillator strength decreases, causing a collapse of strong coupling. While a high saturation threshold is ideal for polariton condensation studies, reducing it causes the polariton energy blueshift to become more nonlinear.

Certain sections of this sample exhibited a faint middle polariton state in reflectivity and luminescence data, as shown in Fig. 2.14. Modeling this data, I found it was well fit by two exciton states split by 2.9 meV, compared to the full UP-LP splitting of about 7.7 meV. The asymmetry in the middle polariton can be accounted for by assuming that the two exciton

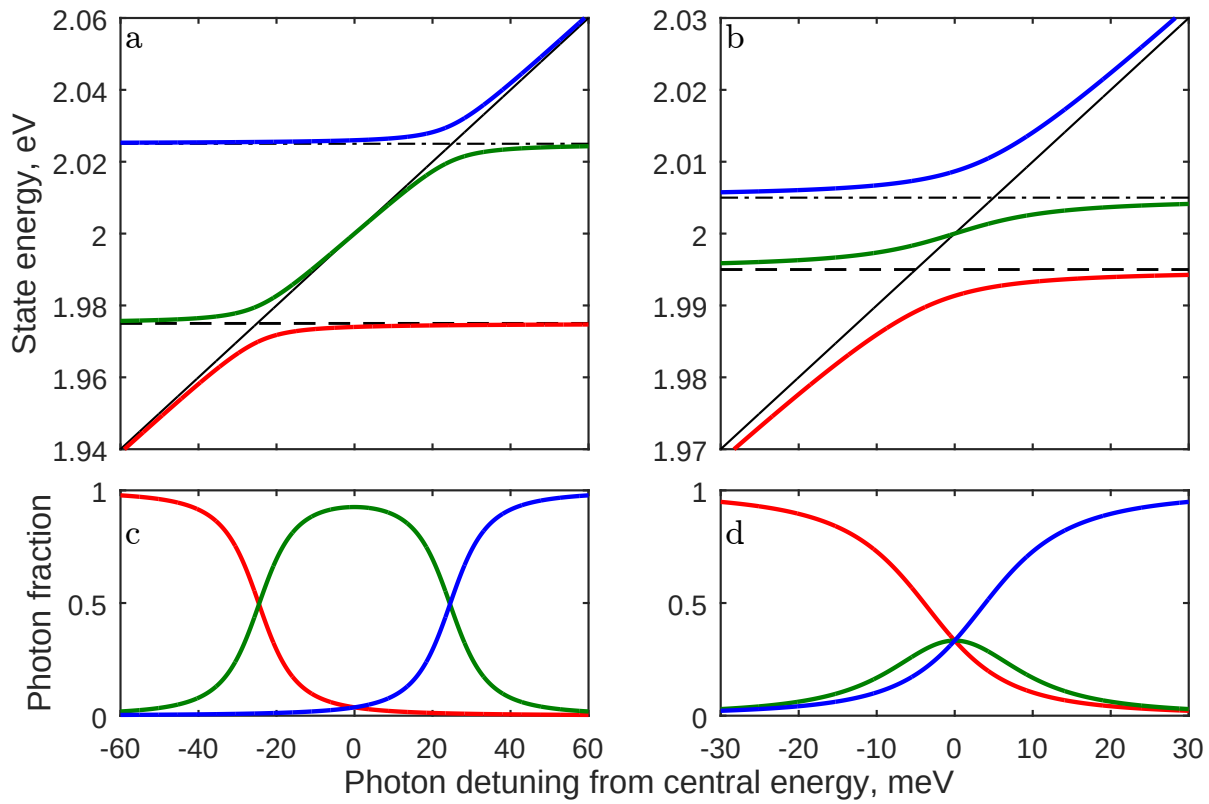

Figure 2.13: A photon mode couples to multiple non-degenerate exciton states. Here I use the same conditions as Fig. 2.12 (b,d) except that the exciton states are split. In frames (a,c), the excitons are split by much more than the coupling strength, g_0, and there are two virtually distinct regions of strong coupling around each resonance. In frames (a,c), the excitons are split on the order of the coupling strength, and the three polariton states must be considered as a single system.

states do not couple to the light equally well. In this sample with three QWs, it seems that two exciton states are nearly degenerate while the third is noticeably split. Specifically, this data is fit well by using $g_0 = 2.2\,\text{meV}$ for the lower exciton mode while the upper exciton was assumed to be $\sqrt{2}$ times stronger.

If the excitonic disorder in this sample is representative of all our samples, one can reason that the middle polariton is stronger in this sample for two reasons: first, because the

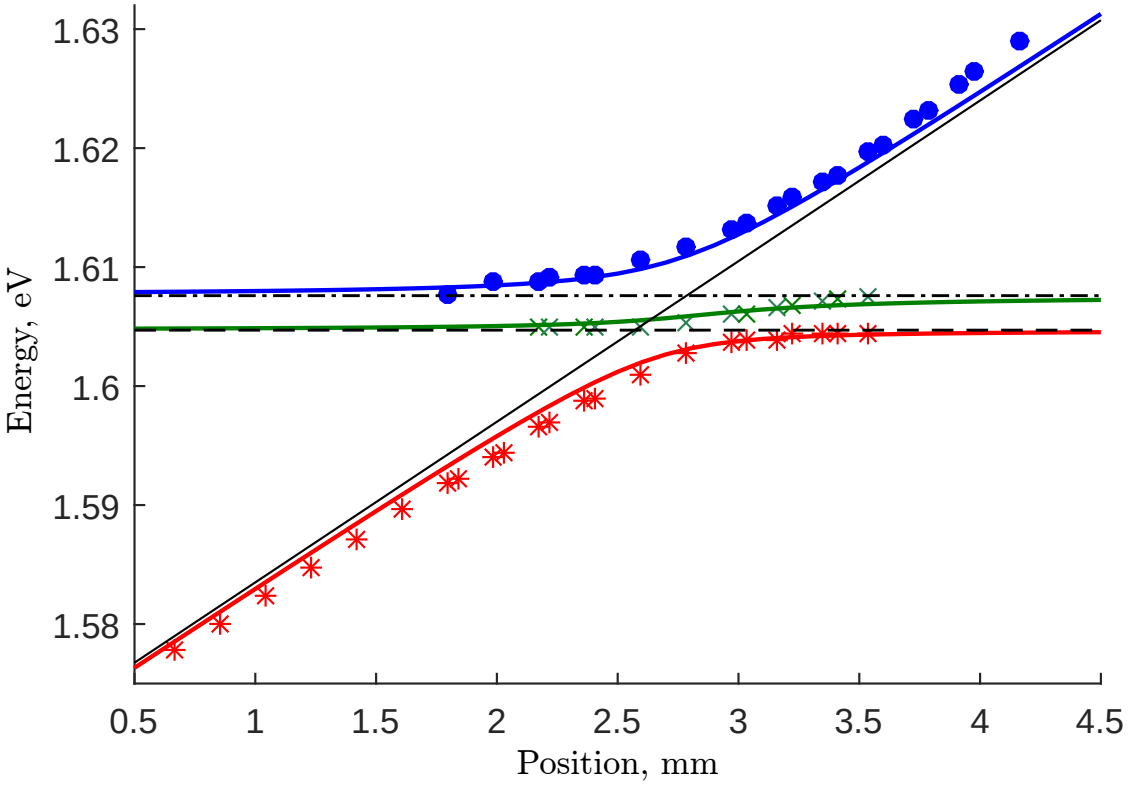

Figure 2.14: Three polariton states visible in a GaAs microcavity sample. Scattered data is from reflectivity. The red, green and blue lines are calculated polariton states arising from the photon (black solid line) and two distinct exciton states (black dashed and dashed-dotted lines). There is a slight asymmetry where the middle polariton state more closely follows the dashed rather than the dashed-dotted exciton line where both are equally detuned from the photon mode. This suggests the higher energy exciton is coupling more strongly than the lower one.

effective coupling is reduced, the disorder in the excitons is larger relative to the splitting. Second, with fewer QWs in the cavity there will be a higher probability that exciton modes will be split enough to make the middle polariton noticeable. The stochastic nature of the excitons seemed to be strongly related to observation of the middle polariton in this sample, as it was observed in some slices of the sample, but not slices taken elsewhere from the wafer.

I have shown that excitonic disorder can muddle polariton states compared to Hopfield's original derivation. However, motional narrowing from the light mass of the polariton can greatly reduce the disorder broadening of the exciton mode. Also, so long as there is a collection of exciton states with disorder less than the intrinsic exciton-photon coupling, they will contribute collectively to the UP-LP splitting. In some sense, the cleanest cases of polaritons will be samples with only single, virtually perfect QWs, *or* samples with many QWs to enhance the coupling beyond the intrinsic disorder.

2.3.6 COUPLING TO THE EXTERNAL WORLD

The most direct coupling of the polariton system to the outside world is via coupling to external photons. This enables us to directly observe the polariton system by collecting photons emitted by decaying polaritons.

Unlike cold atom studies, in which gasses of atoms are destructively measured by light scattering methods or time-of-flight measurements, the light emitted by finite lifetime polariton gasses constitutes a continuous stream of information on the state of the gas. When a polariton decays from the cavity, it becomes an external photon which conserves energy (wavelength) and momentum (in-plane wavevector). The emitted radiation field can be spatially imaged for spatial information, angularly resolved for momentum information, or spectrally resolved for energy details. We can even make correlation measurements, polarization measurements, or time-resolve the emission to collect information on the polariton gas. And all of these measurements can be made continuously on a single shot experiment.

Conversely, we can generate polaritons in a controlled state by illuminating a sample with light in the desired state. If the light is matched in energy, wavevector, etc., polaritons can be created with those same properties. In this way, one can engineer a condensate with coherence imprinted on it by an exciting laser. Exciting the polariton system in this way is referred to as resonant excitation, as the pump signal is energetically resonant to the polariton state.

However, we often generate polaritons with a pump much higher in energy than the polariton states. This is often desirable if we want to observe a condensate spontaneously

form from a thermal distribution. Physically, the high energy pump is tuned to excite higher energy exciton states or free carriers. Often it is convenient to work at a wavelength at which the DBR mirrors are relatively transparent.

By creating hot excitons or free carriers, we can safely assume that they scatter enough with the lattice while cooling down that any initial coherence from the excitation source will be lost. In general, these excitations will cool down to form a bath of excitons at high momentum. As these excitons cool towards $k_{\parallel} = 0$, they enter the polariton dispersion and can cool to form a thermal or quasi-thermal distribution of polaritons [61, 64, 65].

3.0 EXPERIMENTAL METHODS

As discussed in Section 2.3.6, the principle decay mechanism for lower polaritons is coupling to external photons, which gives us a direct way to observe and interact with the system. When a polariton decays, a photon is created external to the cavity, which can be detected using standard optical techniques. In this process, the energy and in-plane wavevector must be conserved, as well as other properties like angular momentum and phase coherence. By directly measuring light emitted from polaritons, we can thus observe these properties of the system without making destructive measurements on the population. By merit of the same coupling to the external field, we can directly create polaritons by pumping with light matched in these conserved parameters. Alternatively, we can excite electronic states at much higher energy than the polaritons. This creates a gas of carriers that can cool to form a quasi-thermal distribution of polaritons.

3.1 EXCITATION METHODS

Almost all polariton experiments to date utilize optical methods to generate populations of polaritons. While this is usually the most convenient method, one can in principle electrically inject carriers, as was recently done to make an electronic polariton laser [6].

When using optical injection, one must select the optimal method for the desired population. Nonresonant injection creates a quasithermal polariton population with a reservoir of free carriers and excitons. Typically the polaritons will have properties in accordance with a thermal-type population. Alternatively, resonant injection affords the experimenter a direct way to create and manipulate the properties of a quasicondensate, since the polaritons will

be created with the properties (e.g. in-plane momentum, polarization and energy) of the exciting light.

3.1.1 NON-RESONANT PUMPING

By pumping at the short wavelength reflectivity minima of the DBR used to make a microcavity, we can create free carriers in the quantum wells with great efficiency. Figure 3.1 highlights these spectral windows in the cavity reflectivity.

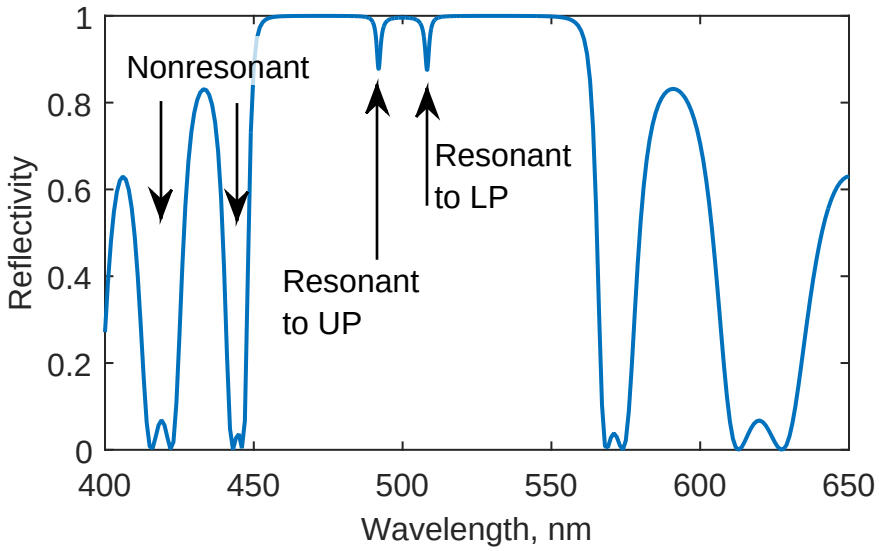

Figure 3.1: Calculated reflectivity of a microcavity structure. Nonresonant excitation is optimal when pumping hot carrier states through spectral windows of the DBR. By definition, resonant pumping requires the source to be energy tuned to the polariton states.

Nonresonant excitation is often optimal for characterization of polaritons and condensation studies. At low density, emission may only be observed from the lower (and possibly upper polariton) states. Moreover, since the excitations in the QWs must scatter many times to cool into the polariton states, we can assume that most properties of the pump, such as polarization and phase coherence, will be lost before becoming polaritons. This is particularly desirable for condensation studies where we are most interested in these properties being inherent to a general thermal state rather than imprinted by a pump laser.

The earliest application of nonresonant injection was in the search for Bose-Einstein Condensation of polaritons. The objective was to create a thermalized distribution of polaritons, then increase the density (by increasing the pump power) until the system undergoes a phase change and exhibits spontaneous onset of phase coherence. This was demonstrated (with some caveats) in 2006 and 2007 [1, 2]. Ultimately, there are some philosophical issues with the semantics of calling these states true BEC; however, the systems still demonstrate spontaneous onset of coherence and polarization emerging from a quasithermal distribution. As such, we typically call these systems quasicondensates to emphasize the fact that they are not fully equilibrated and not true BECs.

In general, increasing the density of a nonresonant pump will increase the density of polaritons. However, energy shifts of the polariton and the dynamics of cooling can be highly dependent on the pump intensity. For example, a high density of carriers can lead to a mean-field blueshift of the exciton state (essentially an exchange interaction between the fermionic components of the exciton) and also screen the electrostatic interactions of the electron and hole (thus reducing the oscillator strength of the exciton). Investigations of the cooling dynamics of polaritons also indicate that free carriers are critical to the thermalization of the very light-mass lower polaritons [65, 66]. Therefore, the quasi-thermal distribution of polaritons that arises from nonresonant pumping can actually cool more efficiently at higher density. Ultimately, these effects combined with bosonic stimulation to form a system with rich nonlinear effects.

These strong nonlinear effects therefore permit additional manipulation of polaritons. It is common practice to use a high intensity optical pump to create a local reservoir of carriers that increase the energy of the polariton state. This creates a potential energy barrier for the polaritons. In this way, we often use nonresonant pumping as a way to simultaneously inject and trap polaritons [20, 35, 67–70].

It should be noted that the cool-down dynamics of hot carriers can limit the time-response of the polariton system to fast pulses of light. For example, time resolved experiments discussed in Section 4.2 show a smearing out in time and slow decay that may be more due to the cooldown of hot excitations instead of the polaritons themselves. Essentially, the complicated scattering and finite scattering rates results in the reservoir of carriers persisting

longer than the duration of the 2 picosecond exciting laser pulse.

3.1.2 RESONANT PUMPING

Figure 3.1 also highlights the possibility to resonantly excite lower and upper polaritons with an optical source tuned to the same energy of those states. One must recognize that these states tend to be highly angle dependent and potentially sensitive to polarization. Thus, to get any light into the cavity, the excitation source must be carefully tuned for all of these properties.

Often, the high degree of finesse required to successfully couple light into the cavity resonantly is worth the effort. In principle, resonant injection allows for sculpting a population or quasicondensate with precisely the desired qualities. Using a coherent source, a condensate can be created in a monoenergetic state. Using two resonant beams, one can collide two condensates [71]. In the presence of a quasithermal background of polaritons, it is possible to use a low density resonantly injected pulse to seed a condensate [72–75], or it would be possible to stir a thermally formed condensate.

Unlike nonresonant injection, there is no cooldown time to populate the polariton states from an electronic reservoir. As such, the duration of a population of polaritons following a fast excitation pulse is limited by the intrinsic polariton lifetime. This was the foundation for the transport and lifetime measurements described in chapter 5.

3.2 OPTICAL OBSERVATION METHODS

As stated, we primarily observe polaritons by directly measuring the photons that are a product of their decay. These photons conserve information of the polaritons that spawn them, since the decay process must conserve energy, momentum and spin angular momentum.

For example, we can prove that in-plane momentum must be conserved. Figure 3.2 demonstrates the decomposition of a wavevector into its in-plane (k_x or $k_\parallel$) and out of plane (k_z) components. We recognize that the intervening layers of a DBR structure have no

impact on these conservation laws, so we directly compare the cavity layer and the final external medium.

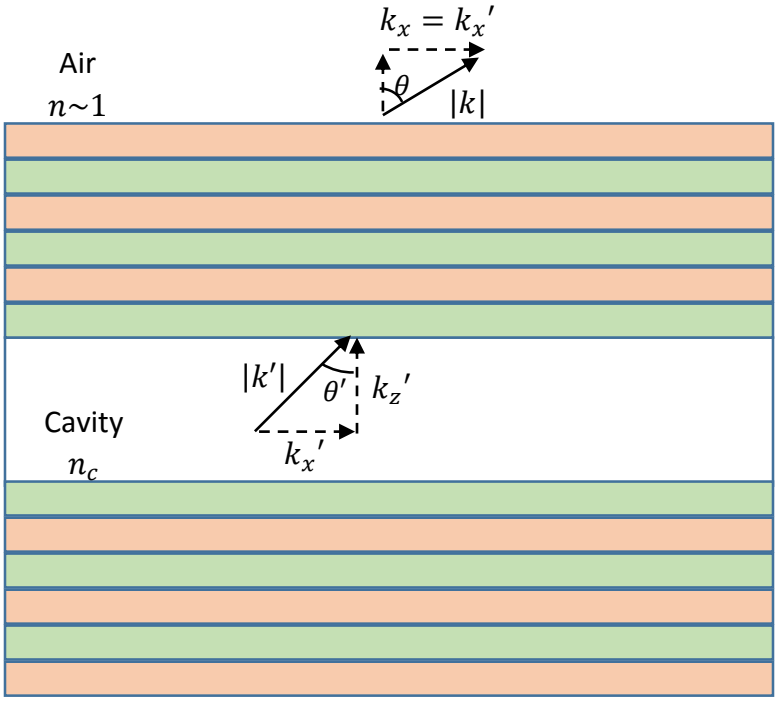

Figure 3.2: Conservation of in-plane momentum across optical interfaces. We compare the wavevector outside the sample to that of the polariton in the cavity.

Using Snell's law and requiring conservation of energy, we can deduce the relationship between the wavevectors of light in the cavity and in air. We start by defining the wavevector inside the medium as $\vec{k'}$ and wavevector in air as $\vec{k}$. We appproximate the index of refraction in air as $n_a = 1$. Using conservation of energy to relate the magnitude of these wavevectors we write

$$E = \frac{\hbar c}{n_c}|k'| = \hbar c|k| \tag{3.1}$$

$$|k'| = n_c|k|. \tag{3.2}$$

Then applying this to relationship to Snell's law and calling θ (θ') to be the angle of $\vec{k}$ ($\vec{k'}$) with respect to the normal:

$$n_c \sin(\theta') = \sin(\theta) \tag{3.3}$$

$$|k'| \sin(\theta') = |k| \sin(\theta) \tag{3.4}$$

$$|k'_{\|}| = |k_{\|}| \tag{3.5}$$

This conservation of in-plane momentum also implies that the direction of $\vec{k'_{\|}}$ is conserved, since the refracted ray lies in the same plane of incidence.

3.2.1 REAL-SPACE IMAGING

The most direct measurement of polaritons is position-space or real-space imaging, where we use standard optics to collect and image light emitted from a sample surface. All standard optical techniques for imaging and measuring light are relevant: we routinely manipulate and detect the light with waveplates and polarizers, wavelength dependent filters, CCD cameras and so on.

One typical imaging setup is to use an infinity corrected microscope objective in conjunction with an imaging lens. Infinity corrected lenses have aberrations minimized for collimating light coming from the focal plane of the lens. The microscope objective enables us to have a large numerical aperture collection and allows for moderate to high magnification. Figure 3.3 is a schematic demonstrating several aspects of a standard experiment. It includes formation of an image of the sample surface at some detector or other processing system, but it also includes methods of injecting light into the system for scattering, reflectivity or excitation based experiements.

However, due to the conservation of in-plane momentum upon emission, we often use Fourier-filtering to only image certain states. For example, we may choose to only look at $k_{\|} \approx 0$ to select the condensed portion of a gas, since the BEC should exist in the ground state. This can also be useful to cut out intense reflected laser light from the excitation method. This filtering is achieved by placing an iris or pinhole at the Fourier image (also called the $k_{\|}$ image) of the sample.

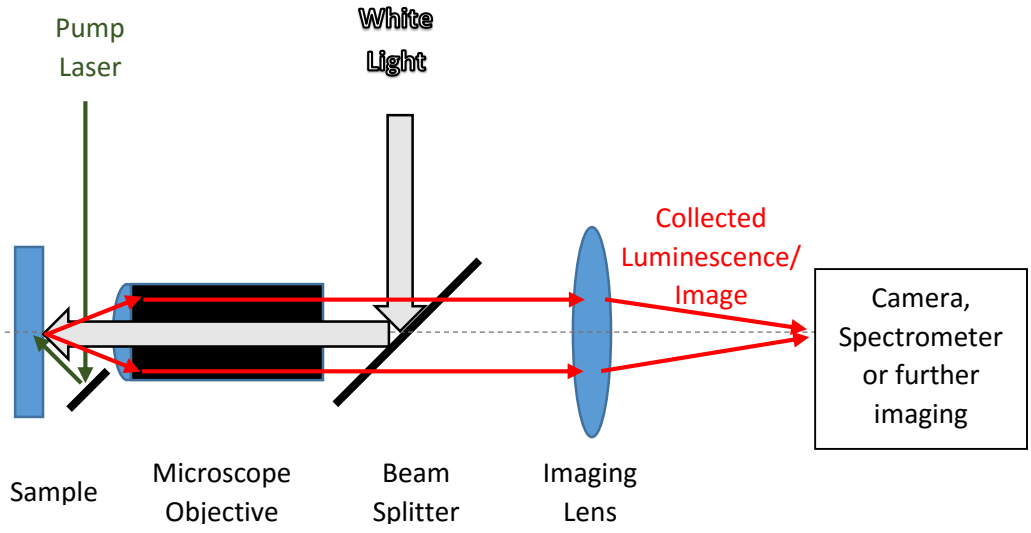

Figure 3.3: This representative experimental setup uses a microscope objective and a separate imaging lens. White light (gray arrow) can be injected through the microscope objective via a beam splitter for imaging or spectral reflectivity measurements. If the working distance of the objective permits, the pump laser (green) can be injected via a large angle. This beam can be focused and controlled independently via a separate lens. Lastly, the collected light (red) can be either luminescence or scattered light. This is imaged directly onto a sensor or further processed before collection.

Often, it is useful to trace rays when initially planning an imaging setup. A full discussion of the ray tracing calculations used in this dissertation is presented in Appendix B. In the scope of this section, we present ray diagrams to highlight the operation of simple optical setups. Figure 3.4 shows how lenses act to form an image of the sample surface in a setup similar to that in Fig. 3.3.

The nice feature of this optical design is that the final magnification can be selected by

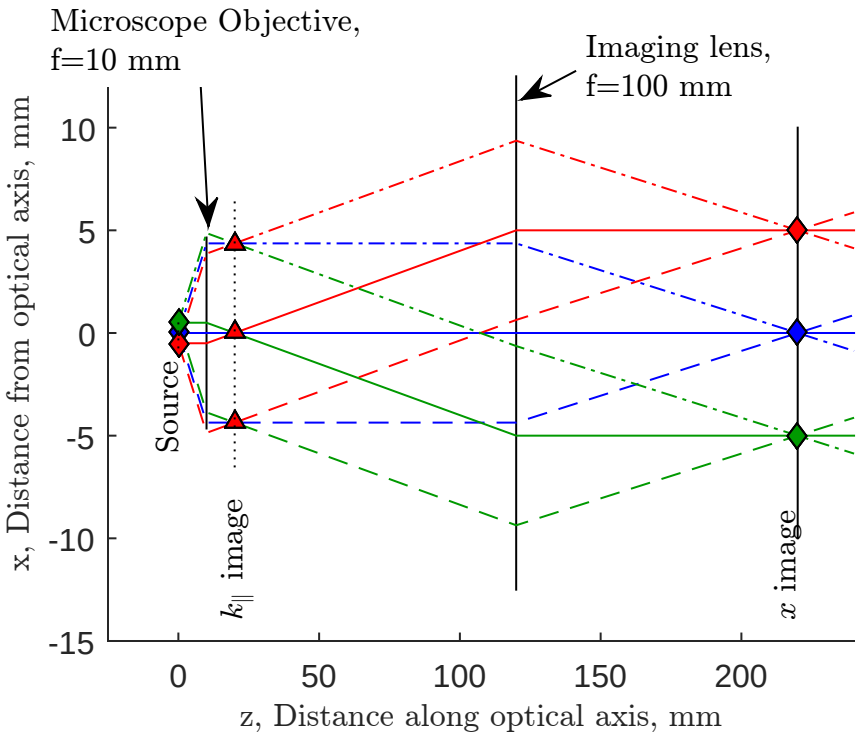

Figure 3.4: Here we image a sample using two lenses. Diamonds indicate three specific points on the sample (red, green or blue). Angle of emission is indicated by different line styles, with solid lines indicating $k_{\|} \approx 0$. Red triangles indicate points where momentum space is resolved.

simply changing the final imaging lens, as $m = -f_2/f_1$. However, such a change may require a change of the total path length from sample to final image, which can be difficult if the final image is formed on a spectrometer slit. Moreover, when using a microscope objective, it can be impractical to Fourier-filter the emission at the first momentum space image (located at $z = 20$mm in Fig. 3.4) because this would also impede other light injected through the objective, or even because the the Fourier image may be formed *inside* of the large compound lens.

Care must also be taken to ensure that light further imaged into later optics is well matched to them. In particular, spectrometers have inherent numerical aperture limitations.

Even if a good image is formed on the spectrometer entrance slit, it is very possible that some of the light can be clipped as it passes through the spectrometer. This is a particular concern when imaging a spatially extended source that spans the full numerical aperture of the imaging system, as is seen in Chapters 4 and 5.

Sometimes, it is also desirable to be able to selectively filter or image different parts of the field of view. By extending the imaging system beyond that shown in Figs. 3.4 and 3.3, we will see that we can facilitate both a real-space and momentum-space filtering. Figure 3.5 is a slightly modified and extended setup over those previously shown. This setup also uses an extended microscope objective-imaging lens distance to ensure that a momentum space image is formed in an accessible location for filtering of those images. As the final real space imaging happens after this, unwanted scattered or reflected light can be eliminated from the final image, but since these images are so far from the injection optics, the filtering won't interfere with the excitation.

It should also be pointed out that some traced rays (red dash-dotted and green dashed) slightly miss the imaging lens. These rays would be clipped using only a one inch diameter lens as shown in the ray diagram. Either we must accept that data from such a system would not be reliable at these edges of the field of view, or a larger diameter lens must be used here.

The addition of the final imaging lens also adds flexibility in selecting the field of view. This final lens can be installed on a micrometer driven translation stage, allowing the experimenter to image different regions of the field of view. Due to the conjugate nature of the microscope objective-imaging lens system, translating only one of these can easily introduce aberrations into the image. Moreover, depending on the experimental method, moving the objective can change the pump conditions.

3.2.2 MOMENTUM SPACE IMAGING

Here I discuss angle-resolved imaging of light. Since polaritons exhibit a correlation between the in-plane wavevector inside to external photons, this angle-resolved imaging allows us to measure the in-plane momentum. Thus we call these images angle-resolved, momentum-

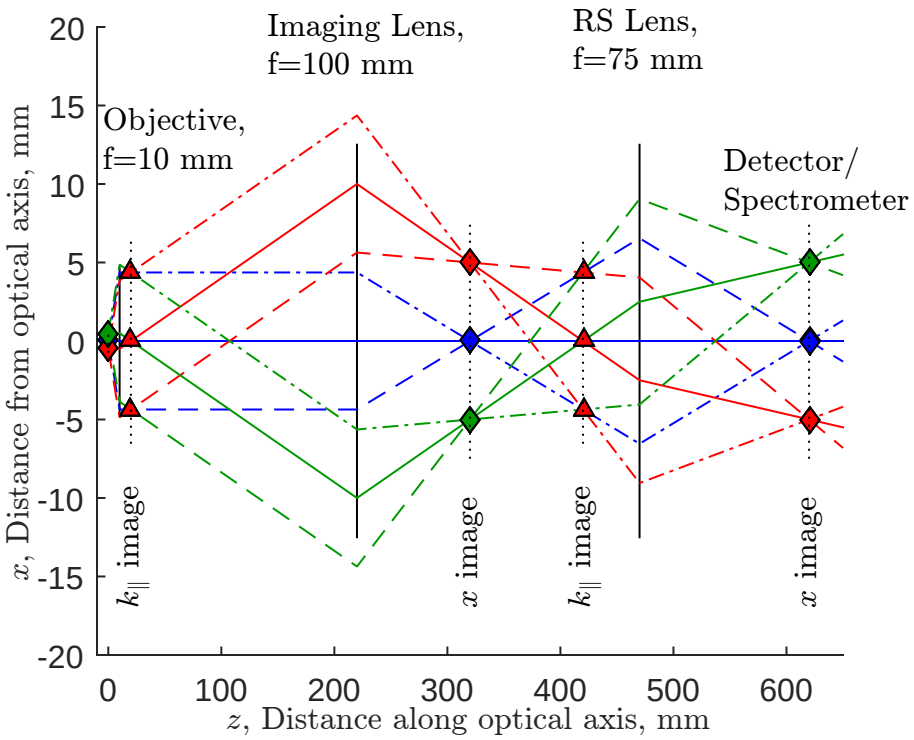

Figure 3.5: We have used a second imaging lens to re-image the first image formed in previous setups, and slightly changed inter-lens distances. Now both a real space (diamonds at $z = 320$ mm) and a momentum image (triangles at $z = 420$ mm) are accessible for filtering.

resolved, or k-space images. The general idea behind momentum-space imaging is to use a lens to Fourier transform light emitted from the sample. For a perfect lens in geometric optics, the positive focal plane is a perfect angle-transformed image of light passing through it. Figure 3.6 and 3.7 show two examples of this property. We see that it doesn't matter if a source object is placed at the object-side focal point or elsewhere. Still the angle-resolved image is formed exactly the focal length beyond the lens.

Downstream from this first lens, we recognize that we may require more optics to image this angle-resolved image plane. In principle, one could place a diffuse plate (such as ground glass) as a screen, then re-image light scattered from this. Instead, it is generally more

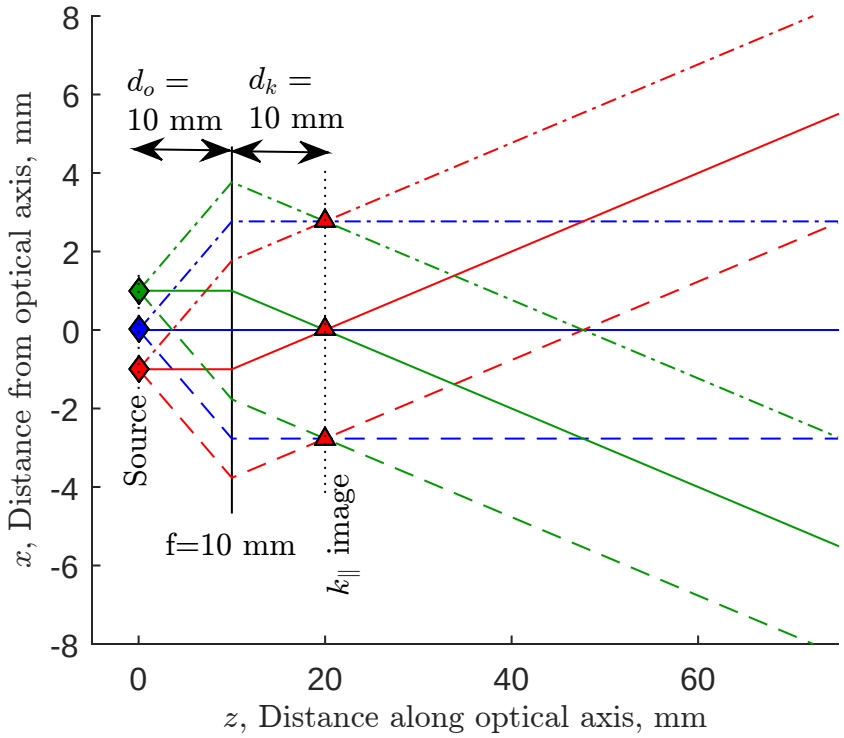

Figure 3.6: A lens forms a proper angle-resolved image at the focal plane of the lens. The magnification of the image is dependent on the focal length of the lens. When exactly the focal length away from the object, rays are collimated, with off-axis spatial points being collimated at an angle.

desirable to not use a diffuser and instead directly image the unscattered light from this plane. Figure 3.8 shows ray tracings where this can is done with a single lens.

There is a strong similarity between the designs shown in Figs. 3.5 and 3.8. Both show a series of isolated real-space and momentum-space image planes. As stated, these images permit us to easily filter out unwanted artifacts from the data, so it is desirable to do final imaging for *both* real and momentum space after these. Figure 3.9 therefore shows the most ideal momentum-space imaging setup. Since it is difficult to move most detection equipment (such as cameras, spectrometers or streak cameras), we require that the real and angle-resolved images can be projected onto the same plane. The final real-space imaging lens in

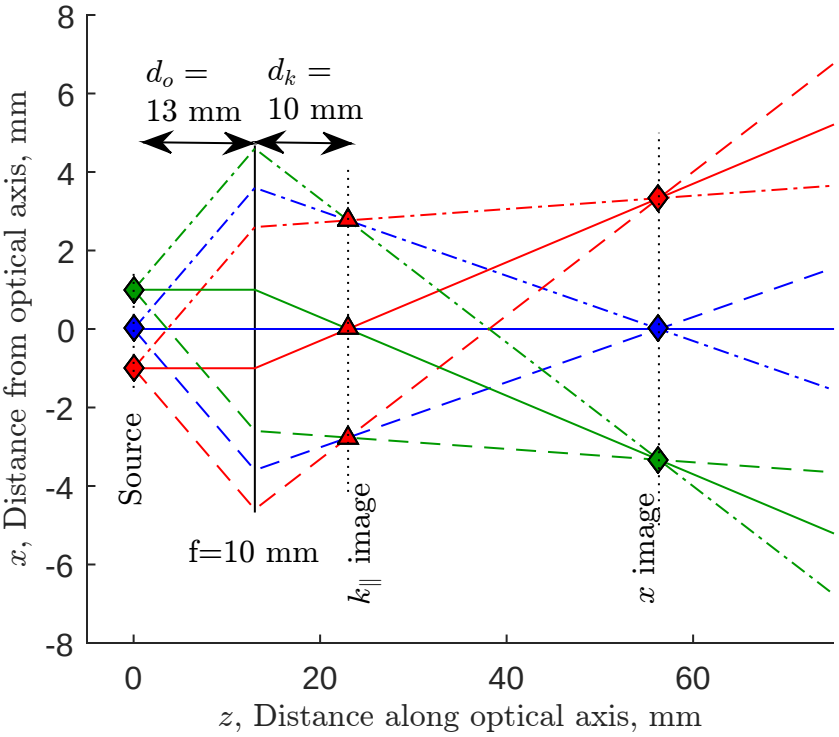

Figure 3.7: Just as in Fig. 3.6, a proper angle-resolved image is formed at the focal plane of the lens. In this case, the object is not placed directly at the focal length of the lens; rays are not collimated in the far field, but instead form an image.

Fig. 3.5 is thus replaced with a single lens or lens system to form the angle-resolved image on the sample detector plane.

The calibration of angle-resolved images is less straightforward than that of real-space. While the previous images can be calibrated using a test sample with known features or translating the sample a known distance (usually using a micrometer), angular emission is by its very nature insensitive to the sample position. The most direct way to measure this magnification is to use a calibrated iris between the microscope objective and the sample. This casts a shadow in k-space related to the diameter of the iris and the distance between the objective and sample. We require a source of illumination, which can either be properly collimated light focused through the objective, or it can be emission from a point in the

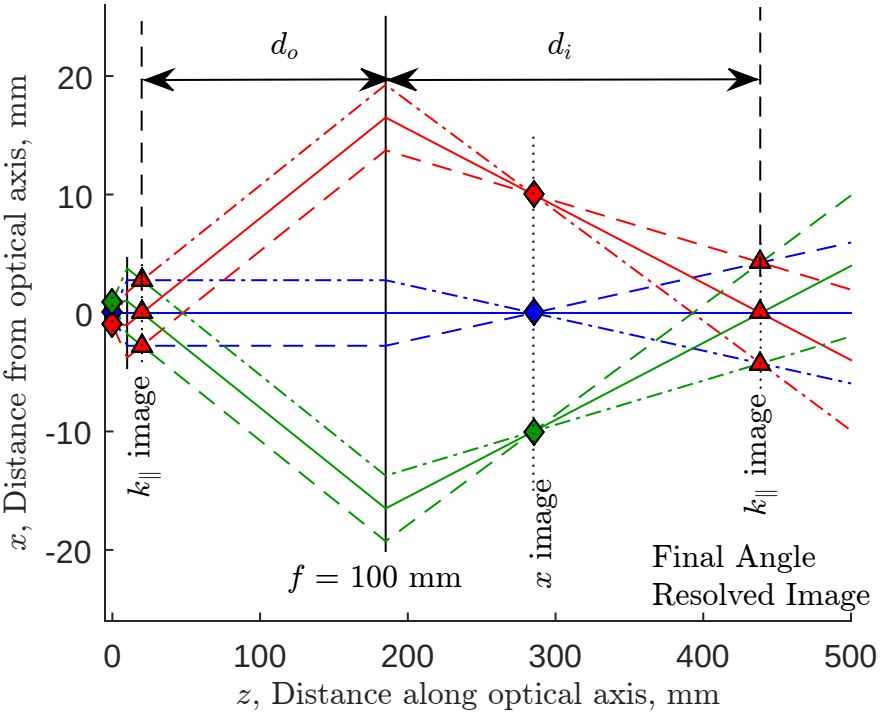

Figure 3.8: Using a second lens we can directly re-image the angle-resolved transform demonstrated in Fig. 3.6. The original angle-resolved image is treated directly as the new object, and its image is ultimately measured.

sample plane. Typically, we use an adjustable iris and measure a series of iris diameters to the aperturing in k. However, this typically requires removing the iris each time, which can shift the placement relative to the optical axis. Errors can arise from the spatial size of the illumination source (since the iris is not at a proper angle resolved plane), and from misalignment of the iris relative to the optical axis. This method can be difficult or impossible if the working distance of the objective is too short to permit inserting an iris.

A second calibration procedure currently under development is using a patterned test object, such as a grating. In this case, we can place the test pattern exactly at the imaging focus and simply scatter a plane wave from it. The diffracted light will have angular dependence on the wavelength of light and the spacing of lines or dots on the test pattern. If the groove

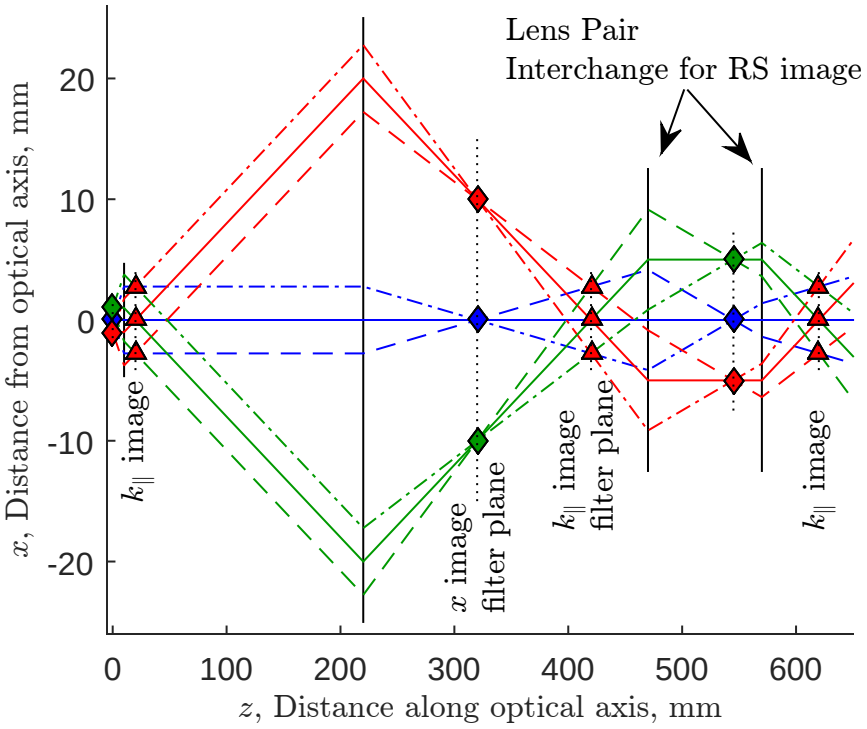

Figure 3.9: By replacing the final lens in Fig. 3.5, we can form an angle-resolved image at the same detector plane. (Note that this plot includes rays from larger x position on the source to emphasize ray tracing.) The final magnification and focus can be fine tuned by lens selection. The x-image and $k_\parallel$-image planes are ideal locations to real-space or Fourier filter the data prior to final imaging. This is easily done by pacing an adjustable iris at each plane.

density of the test pattern is low enough, several diffracted rays will be angularly-resolved, and this will immediately yield a calibration of angle-space.

Ultimately, one must note that emission from the sample is properly angle-resolved using this method, while we are typically interested in $k_\parallel$, since that is the conserved property. Since $k_\parallel = |k| \sin(\theta) \propto E \sin(\theta)$, we can only properly calibrate the k-axis for energy-resolved data. However, since the relative change in $|k|$ is small over the range of data collected from the polariton, we usually approximate the k-axis as uniform in a given image. However, it

is pertinent to realize that k-resolved data is more meaningful when spectrally resolved.

3.2.3 SPECTROSCOPY

The spectrometer is an invaluable tool for optical experiments, using a grating (typically a blazed grating ruled with 600-1800 grooves per millimeter) to spectrally disperse the light. Since spectrometers are discussed well elsewhere [76, 77] and are generally commercially available, I will not discuss much of their operation aside from that which is directly relevant to the work presented here.

Fundamentally, a spectrometer serves to disperse light into its constituent colors, allowing us to calibrate the wavelength of light in an image. However, the action of a grating generally requires that we sacrifice one dimension of an image in order to gain the spectral dimension. Practically, we select a slice of an image (possibly of an x vs y image or a k_x vs k_y image) which is then dispersed in energy (for example to form energy vs x or energy vs k_x). If one wants the spectral profile of a complete x-y image, then he must typically take a series of spectral pictures while stepping the x-y image across the spectral slit.

Often times, this sacrifice of a spatial axis is not a problem, but it must be accounted for when the underlying system is not isotropic. For example, the gradient in the photon energy inherently breaks the symmetry in our samples. Thus, real and k-space data often have broken symmetries. In many cases, it is critical to carefully align some aspect of a sample (perhaps the gradient or a structure on the sample) with the spectrometer slit.

Another important aspect of spectrometers is their limiting numerical aperture. Since the resolving power of the spectrometer increases with length, spectrometers tend to have long optical path lengths, often on the order of our entire optical setups. This means that light can disperse much more within the spectrometer and lead to clipping. Figure 3.10 demonstrates clipping that can occur within a spectrometer. Even when an image is carefully formed on the entrance slit of the spectrometer, off axis or high-NA light can easily be clipped due to the long path length.

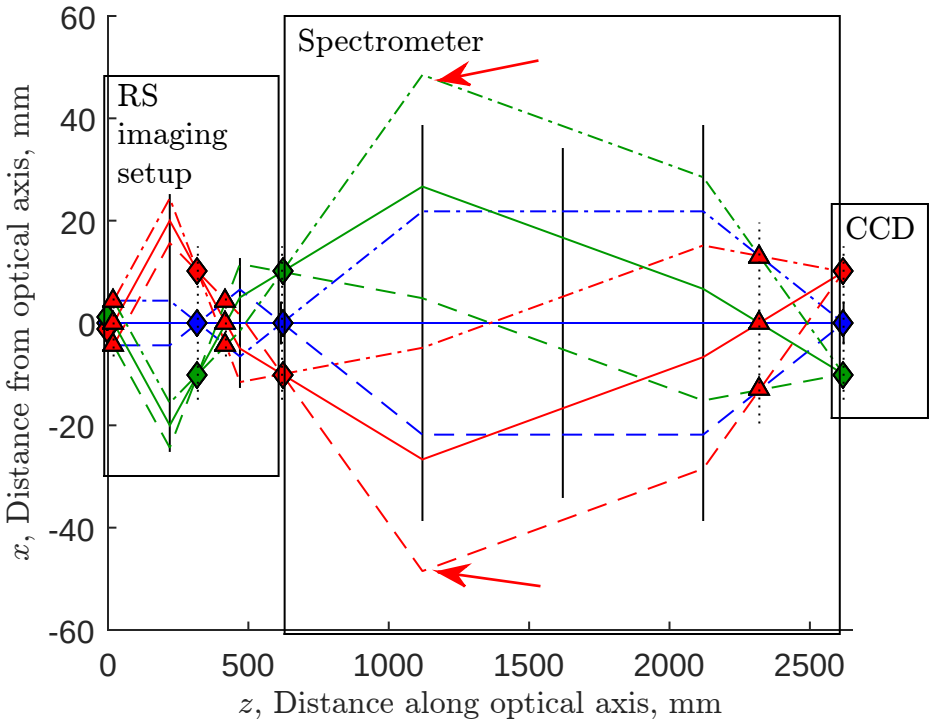

Figure 3.10: Here we take the optical setup in Fig. 3.5 and form the final image on a spectrometer slit. The spectrometer consists of two large mirrors and a plane grating. Unfolded, this optical path can be presented in this way. While all of the traced rays pass through the initial optical setup, some are clipped in the spectrometer (highlighted by the red arrows). It is important to assess these potential sources of clipping when imaging large field of view and large numerical aperture emission.

3.2.4 TIME-RESOLVED IMAGING

Depending on the chosen experiment, we may elect to work in the steady state regime in which we pump continuously with a continuous wave (CW) laser. It is important to recognize that polaritons generally exist in steady states rather than strict thermal equilibria, as the polaritons are continuously decaying. In many systems, the polariton decay can be faster than the rate at which these states are filled from excited states. Thus gain from the CW

laser into higher energy states generally leads to a steady state population of polaritons.

In what is called the quasi-CW regime, we often pump with a chopped CW laser such that the population achieves a steady state for longer than the transients in the population. This is done to minimize heating of the sample, while still effectively keeping the population in the CW state.

As the field has advanced, we have seen that the dynamics of polaritons has been increasingly interesting. For practical applications such as optical switching, the dynamics of the switch are highly relevant to the usefulness of a potential switch (see chapter 6). In terms of the underlying polariton system and its manybody effects, we can continuously resolve the evolution of the population. Thus, if we can measure the polaritons on a relevant timescale, we can characterize the evolution of the system or the constituent polaritons.

In the observable range, polaritons propagate in-plane on the order of 4% of the speed of light and have lifetimes on the order of 1-100 ps. These timescales suggest that ultrafast optical pulses (sub-picosecond) can be viable excitation sources to observe some evolution of the system. Excitation with a pulse on the order of picoseconds will be sufficient to measure particle lifetimes on the order of 10-100 ps (see chapter 5) or to probe the slow cooldown dynamics of non-resonant excitation (see chapter 4).

Here we discuss two particular methods to probe the system on the picosecond timescale. Streak cameras transform temporal dispersion to spatial displacement much like spectrometers transform spectral dispersion into spatial displacement. This is particularly useful when observing the evolution of a population following a femptosecond or picosecond excitation. Alternatively, a pump-probe technique can be used in which two pulses of light are impinged on a sample with carefully arranged time delay. Typically, the delay between pulses is controlled with a delay micrometer stage that changes the path length of one arm prior to injection.

3.2.4.1 STREAK CAMERAS

Generally a CCD camera will have exposures on the order of microseconds or longer, which limits the fastest time differential that can be observed. Fast photodiodes may have

rise times on the nanosecond scale, but still are far too slow to detect polariton dynamics on the order of picoseconds.

Streak cameras using optoelectronic systems regularly achieve time resolution on the order of picoseconds. As outlined in Fig. 3.11, these devices essentially operate in three stages. Initially, light (say a pulse, P_1) excites electrons on a photocathode, causing a cloud of electrons to be emitted. These electrons are accelerated through the streak tube toward a phosphor screen. The streak tube also has an oscillating electric field transverse to the electrons' direction of travel, which causes the electrons to deflect a distance proportional to the amplitude of the electric field. Since the electric field is time varying, electrons excited from the photocathode at a later time (say by a second pulse of light, P_2) will be deflected proportional to the strength of the field. These deflected pulses of electrons strike a phosphor screen and the emitted photons are imaged on a camera.

Streak cameras are most useful when observing a periodic signal, for example a system driven by a mode locked laser. The oscillator circuit of the streak tube is designed for a particular laser system, such as Titanium Sapphire lasers which commonly operate at 76 MHz. The streaking electric field oscillates sinusoidally in phase with the laser repetition rate. This results in the slow camera recording the image on the phosphor over many excitation cycles of the system. While this can result in unwanted jitter of the signal, the enhanced signal is typically important as the integrated intensity of a single event may be very weak.

Often the jitter in the source laser (which excites the sample system and is used as a trigger for the streak camera's oscillator) is the limiting factor in the fastest time resolution possible. However, since the displacement of the pulses is proportional to the electric field strength, streak cameras are generally equipped with different time modes: a stronger electric field will result in more dispersion and better time resolution, up to the jitter the system.

3.2.4.2 PUMP-PROBE METHOD

While pump-probe methods encompass a wide variety of techniques, the general idea is similar. A pulse of light is split into two parts. One part excites or otherwise interacts with

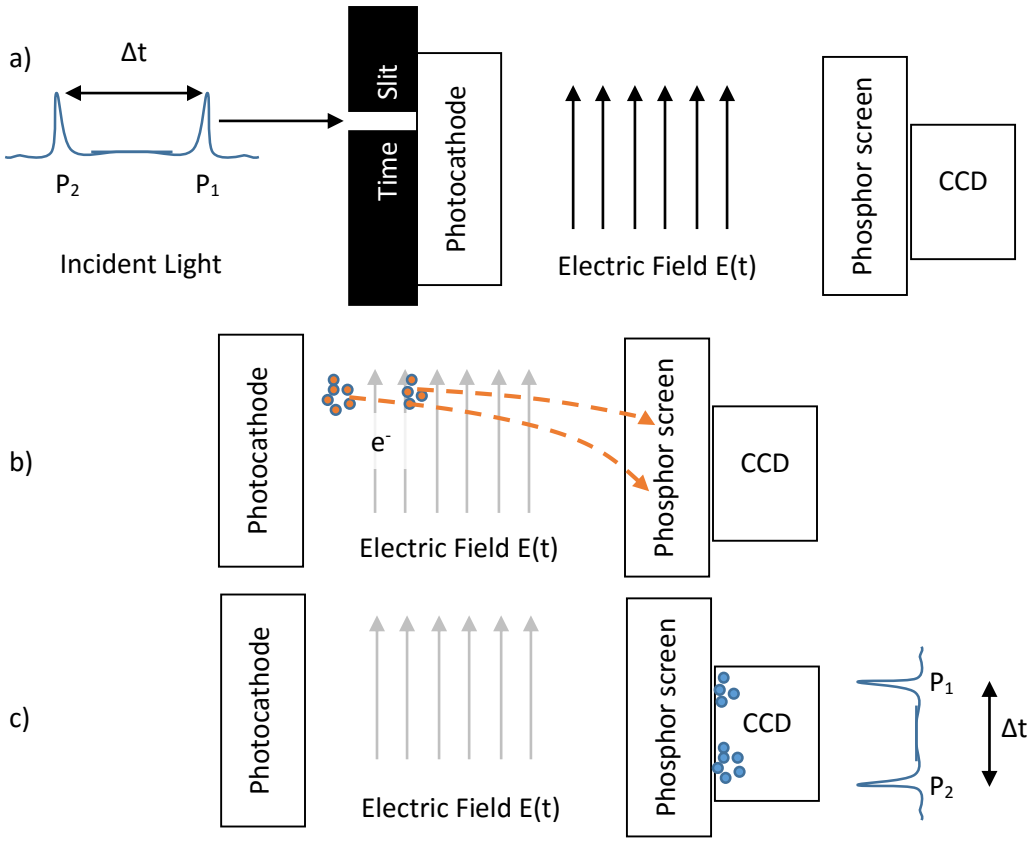

Figure 3.11: Basic principle of streak camera operation. In frame (a), light pulses enter the camera and illuminate a photocathode. This causes pulses of electrons to be photoemitted in frame (b), and these electron pulses fly down the streak tube (moving to the right). Since there is a time varying electric field transverse to the tube, the electron pulses are spatially displaced a distance related to the electric field strength and time separation between light pulses. Frame (c) shows that the displaced electrons are ultimately detected using a phosphor screen and a camera. The image from the camera shows two peaks with spatial separation determined by the relative time delay.

a sample, while the second part interacts after some controllable delay. In some cases, one or both of the pulses can be temporally or spectrally transformed. For example, one may create a population with the first pulse, then look for signs of saturation of those absorbing

states with the second pulse after various delay times. The pulses can both hit the sample at a single point, looking at local effects, or be directed at distant points as a way to probe transport or nonlocal interactions.

In terms of this dissertation, a pump-probe technique was used to demonstrate optical switching in Chapter 6. Here we used a picosecond pulsed laser to resonantly inject polaritons at very high in-plane momentum. We call this first pulse the gate pulse. Second, we used a picked off portion of that beam with a time delay to probe the reflectivity of the microcavity. We call this second beam the signal, since the population created by the first beam modulates the reflectivity that the second beam sees.

4.0 LONG LIFETIME POLARITONS AND BALLISTIC MOTION

This chapter is an adaptation of my 2013 PRB publication [34]. These results and closely related results in [35] showing novel condensation phenomena were the first publications based on new long-lifetimes samples that I finalized the design of in 2010. For many reasons, these new long-lifetime samples exhibited distinctive features with respect to traditional short-lifetime samples.

While the novel condensation results are covered in [35], the purpose of this work was to establish an accurate measurement of the lifetime of polaritons in this sample, since we are claiming that is the reason for the new physics. Here, I develop multiple indirect ways to estimate the lifetime in addition to directly showing the propagation of these polaritons on the millimeter length scale in this new sample.

There is much interest in using polaritons to create a Bose-Einstein condensate, but the lifetime of the particles in most of those experiments has been of the order of a few picoseconds, thus leading to significant nonequilibrium effects. By increasing the cavity quality, we have made new samples with longer polariton lifetimes. With a photon lifetime on the order of 100-200 ps, polaritons in these new structures can not only come closer to reaching true thermal equilibrium, but they can also travel much longer distances. We observe the polaritons to ballistically travel on the order of one millimeter, and at higher densities we see transport of a coherent condensate, or quasicondensate, over comparable distances. In this chapter I report a quantitative analysis of the flow of the polaritons both in a low-density, classical regime, and in the coherent regime at higher density. Our analysis gives us a measure of the intrinsic lifetime for photon decay from the microcavity and a measure of the strength of interactions of the polaritons.

In a typical structure, such as used for these studies, a GaAs-based microcavity is de-

signed to include GaAs quantum wells (QWs) located at the antinodes of a planar optical cavity with end mirrors which are distributed Bragg reflectors (DBRs) made from $Al_xGa_{1-x}As$ and AlAs layers. The strong coupling of the exciton to the photon through the exciton's radiative dipole matrix element leads to the formation of the upper and lower polaritons (UP and LP). Our sample is a $3\lambda/2$ microcavity containing 4 GaAs/AlAs QWs at each of the 3 antinodes. The DBRs confining the cavity are made of alternating planar $\lambda/4$ layers of $Al_{0.2}Ga_{0.8}As$/AlAs. This sample is similar to one used in previous work [2], but the number of layers in both the front and back DBRs were doubled, effectively increasing the designed Q-factor by more than two orders of magnitude and the designed photon lifetime from 2 ps to 400 ps. This is the same sample as studied in Ref. [35].

These new samples have demonstrated fundamentally different physics. At low density, the particles can propagate on millimeter length scales, allowing for polaritons to transmit information over macroscopic distances. Also, localized, moderately-high density excitation results in quasi-condensates that flow away from the pump spot. This is well understood as a gas of polaritons approaching the BEC threshold in a small region with a strong blueshift. Due to the localized blueshift, polaritons will be rapidly accelerated radially away from the pump spot, at which point they propagate mostly ballistically, carrying any coherence achieved in the high density pump spot. The distinguishing factor from previous similar results of polariton condensation [1] is that these polaritons live long enough to propagate away from the pump spot. However, at higher pumping power, the polaritons transition to form a more coherent condensate at $k_{||} = 0$ that is spatially adjacent to the pump spot and at lower energy than the lower density, extended quasicondensate [35].

4.1 LOW DENSITY: BALLISTIC PROPAGATION

The first observation of polariton photoluminescence (PL) in these samples was initially perplexing. Luminescence data in Fig. 4.1(b) shows polaritons on the LP branch propagating a long distance on the sample from the excitation spot. Looking only at this figure, it appears that the polaritons gain energy to travel uphill.

If we compare this to Fig. 4.1(c), however, we can make more sense of the data. Figure 4.1(b) was taken with small numerical aperture (NA), while Fig. 4.1(c) was taken with large NA. The NA matters because a polariton with wavevector $k_\parallel$ is a coupling of an exciton and a cavity photon both with the same $k_\parallel$; when the polariton decays, it emits a photon external to the cavity with the same wavevector. This gives a one-to-one mapping of the angle of the photon emission in the far field to the in-plane $k_\parallel$ of the polaritons before they decay into external photons. Therefore opening up the numerical aperture of the imaging system collects light from polaritons at higher $k_\parallel$. For the data of Fig. 4.1(b), the low NA restricted the polaritons observed to those with $k_\parallel \sim 0$. We see in this figure the gradient of the $k_\parallel = 0$ energy, i.e., the potential energy of the polaritons, due to the wedge in the wafer thickness discussed above. This spatial gradient of the ground state energy is the same as a force on the polaritons, since $F = -\nabla U$.

The data of 4.1(c) were taken with a lens system with a 0.4 NA, much larger than the NA used for Fig. 4.1(b). This larger acceptance angle corresponds to imaging polaritons with a much wider range of momenta. Figure 4.1(c) shows that there is a significant population of polaritons at $k_\parallel > 0$; The broad distribution of high-momentum polaritons exists at the point of creation due to the many random scattering processes which occur after non-resonant excitation. Some of the high-$k_\parallel$ polaritons flow uphill and eventually reach $k_\parallel = 0$ where they can be observed with low NA, while others flow downhill until they exit the 0.4 NA collection angle.

One critical feature to notice in these data is the sharp minimum energy cutoff on the left side of Fig. 4.1(c). The polaritons at the excitation spot partially thermalize according to the relaxation dynamics of hot carriers and excitons [65,66]. Upon reaching polariton states with very light mass and low scattering rates, the polaritons are able to travel ballistically. This explains the minimum energy observed on the right—polaritons are streaming ballistically away from the excitation spot after initially scattering into LP states. The polaritons flowing downhill immediately leave the high density excitation region and never scatter to lower energy. The polaritons flowing uphill stream until they hit a point on the sample where the $k_\parallel = 0$ state has the same energy, at which point they can no longer flow to the right, and are reflected back to the left.

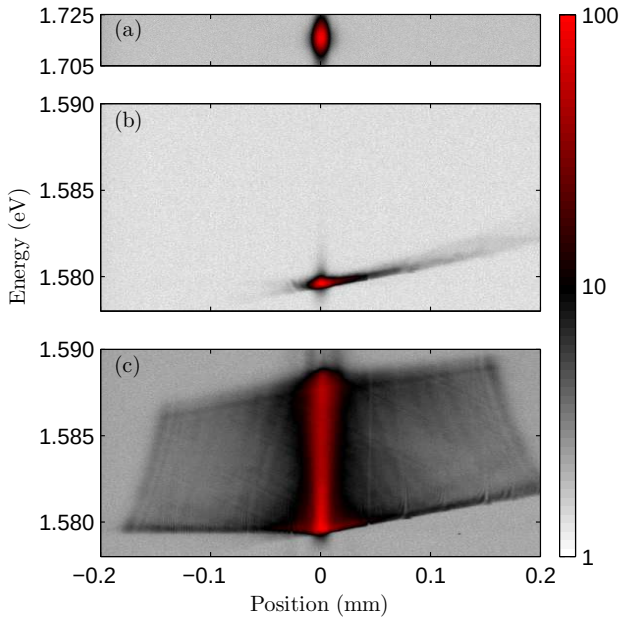

Figure 4.1: Intensity of the PL emitted from the LP, as a function of energy E and position x, recorded with an imaging spectrometer. Intensity counts are presented on a log scale to highlight motion. Data taken at a $k_{\|} = 0$ polariton detuning of -21 meV with a pump power of 500 μW at 705 nm focused to a 15 μm diameter spot size (0.28x10^6 mW/cm^2). a) Hot-carrier PL, with size indicating the size of exciton cloud. b) Lower polariton PL, spatially resolved but only collected near $k_{\|} = 0$. The bright spot is the point of creation of the polaritons; the PL at further distances shows the turn around point for each initial energy. c) The same data taken with a larger NA (larger range of $k_{\|}$). Polaritons are generated over a broad range of $k_{\|}$ at the pump spot and ballistically travel outward at constant energy. The sharp cutoff in energy on the downhill side indicates that the polaritons do not scatter once they are spatially distant from the excitation spot. The horizontal/angled cutoff at high energy is the accepted NA of the microscope objective. The cutoffs at $\pm$ 0.2 mm are due to clipping in the optics and spectrometer.

Because of the one-to-one mapping of polariton momentum to photon emission angle in the decay process, we can image the far-field PL to directly resolve the momentum space

distribution of the polaritons, just as we image the near field to observe the real-space profile. Normally, the k-space image integrates over the entire real space observed, so we must use spatial filtering to measure the dispersion relation from a single point of a spatially extended distribution. In Fig. 4.2(a) we present the far-field PL of the emission spot after spatial filtering, which was accomplished using a pinhole in an intermediate real-image plane. The spatial filter selects a region on the sample of approximately 40 μm diameter, which is slightly larger than the pump spot. The PL profile at this spot indicates the initial population before propagating away. Figure 4.2(b) shows the same data without the spatial filter.

There are several features of Fig. 4.2(b) which are complementary to the real-space data of Fig. 4.1(c). The polaritons initially at $+k_{\|}$ move uphill at constant energy while losing momentum, i.e., shifting to lower $k_{\|}$. The polaritons at $-k_{\|}$ flow downhill at constant energy and gain momentum in that direction, eventually leaving the numerical aperture of our microscope objective. There is again a clear cutoff in energy at the vertex of the excitation spot momentum dispersion parabola. The polaritons starting at $k_{\|} = 0$ are the lowest energy polaritons possible at the pump spot where the density is high enough to scatter. These polaritons stream downhill ballistically, giving rise to this energy minimum.

4.2 TIME-RESOLVED PROPAGATION

To verify that the extended polariton cloud is propagating from the point of excitation, we used a Hamamatsu streak camera to time-resolve the spatial arrival of polaritons at various points on the wafer following a pump pulse with picosecond duration. Due to the many scattering processes following the non-resonant generation, all the temporal dynamics of the cool polaritons are broadened and delayed relative to the excitation pulse. Nevertheless, by measuring the arrival times of the polaritons moving uphill, we see clearly that there is a time delay for the propagation of the polaritons as they travel across the wafer. As discussed below, this time delay is consistent with the theory for the time of flight across the sample, using the known polariton dispersion.

Figure 4.3 shows the time-resolved PL for the polaritons (solid blue line) for different

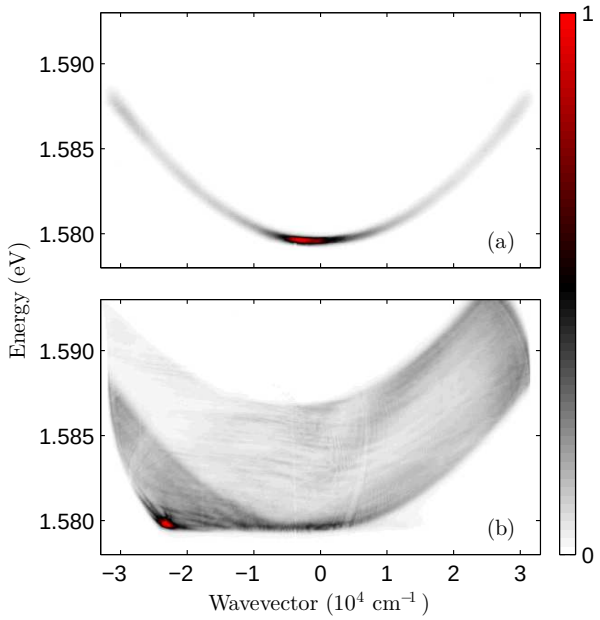

Figure 4.2: Intensity of the PL emitted from the LP, as a function of energy E and in-plane momentum $k_{\|}$, recorded using an imaging spectrometer focused on the far-field emission (Fourier plane). This data was taken under the same pumping conditions as Fig. 4.1. a) Spatial filtering is applied to an intermediate image to isolate the dispersion relation of the LP at the pump spot. Due to non-resonant excitation, polaritons are observed filling the momentum states. b) With no spatial filtering, the excitation spot polaritons are smeared in the downhill ($-k_{\|}$) direction. Polaritons at $k_{\|} = 0$ correspond to the polaritons observed in Fig. 4.1(b). Again we observe an energy minimum coinciding with the vertex of the pump spot dispersion curve, as the polaritons scatter very little after leaving the creation region.

distances x from the generation spot, following the hot PL emission (black line), which indicates the duration of the pump laser pulse. The polaritons were generated non-resonantly on the photonic side ($\delta_0 < 0$) of the wafer with a 2-ps pulsed Ti:Sapphire laser, and $k_{\|} \sim 0$ emission from individual spatial points was spectrally and temporally resolved. The polariton PL is fit with a Gaussian convolved with an exponential decay as shown with the solid red line. The details of this fit are discussed in Appendix C.

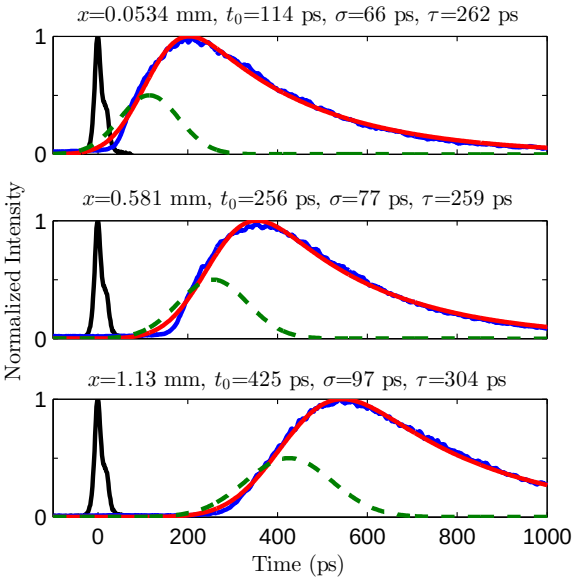

Figure 4.3: Time-resolved $k_{\|} \simeq 0$ PL from the lower polariton at three sample distances from the pump spot. Data were collected following a 2 ps, 2 mW pump pulse with $\lambda = 725$ nm focused to a 50 μm diameter spot (0.1x10^6 mW/cm^2 or 1.3 μJ/cm^2 per pulse) where the $k_{\|} = 0$ polariton detuning was -15 meV. Blue lines: intensity data of PL from propagating polaritons. Black lines: emission of the hot carriers above the stop band within 10 ps of the excitation. Red lines: Gaussian-exponential decay convolution fits to the data with the parameters given above each frame. t_0 is the central time of the Gaussian following the hot PL, σ is the standard deviation, and τ is the exponential decay time. Note that t_0 is an indicator of the travel time–we know that this value must include both the time of flight as well as the time to cool down from hot carriers to the lower polariton. Dashed green line: the unconvolved Gaussian presented as an aid to the reader to see how the t_0 parameter compares to the peak of the intensity data. The convolution with a decay pushes the peak of the fit to significantly later time than the Gaussian fit alone.

The convolution of a Gaussian and decay is an empirical fit which is sufficient for assigning an arrival time to the polaritons. Additionally, this convolution can be written in a closed

form, which makes it computationally convenient to fit the data. Aside from background and overall amplitude, the fit has three parameters: the arrival time t_0, the Gaussian broadening σ, and the decay time τ. We interpret the arrival time as the sum of two major contributions: 1) first, the hot excitations cool down to fill the polariton states at the pump spot. This cool-down time depends on the phonon emission rates. 2) The remainder of the arrival time is due to the actual time of flight (TOF) of the ballistic polaritons to reach a point on the sample where their momentum has slowed to $k_{\|} \simeq 0$, where they are observed. The decay time τ cannot simply be interpreted as the lifetime of the polaritons, since the dynamics of the hot carriers fills these states over a finite time. For example, if the time to cool down into polariton states is comparable to or longer than the lifetime of the polaritons, then the decay time will measure the lifetime of this excited population rather than that of the polaritons.

The green dashed line in each case of Fig. 4.3 is the Gaussian portion of the convolution. As seen in this figure, the peak of each PL curve is not at the fitted t_0 value, which is located at the peak of this pure Gaussian, but is shifted to a later time by the convolution with an exponential decay.

What is clear from the raw data and from the fits is that the more distant points take longer to be populated with polaritons. If the motion is ballistic in nature, then we should expect the dynamics to be explained by semiclassical particle dynamics. In Fig. 4.4 we present the time-of-flight value $TOF = t_0 - t_{cool}$ from the fits as a function of the distance from the excitation spot. The cooldown time, t_{cool}, is the time for the hot excitations (as observed in Fig. 4.1(a)) to fall down into the LP states from which they can begin to propagate. Note that t_{cool} was determined by fitting the data to theoretical models of propagation, since the data immediately at the excitation point shows unreliable t_0 values. The simplest approach to explain the data is to assume that the polariton mass is constant and the potential gradient felt by the polaritons due to the wafer thickness variation is constant, i.e., that the polaritons feel a constant force. We can envision the polaritons as starting with an initial momentum uphill and we observe them when they reach $k_{\|} = 0$. This yields the relationship that the time of flight is proportional to $x^{1/2}$, which is shown as the blue dash-dotted curve in Fig. 4.4. This works well for short distances, but the data beyond $x = 0.4$ mm show a clear upturn which deviates from this simple fit.

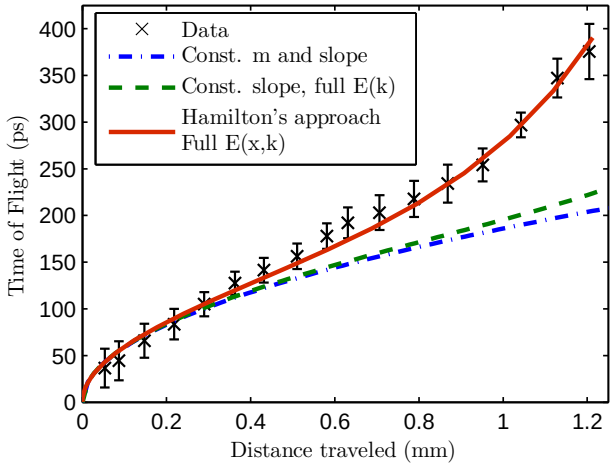

Figure 4.4: Time of flight for polaritons with different initial momenta to reach $k_{\|} = 0$ for the same pumping conditions as Fig. 4.3. Black crosses with error bars: the time of flight as determined from the time-resolved data. Blue dashed-dot line: fit of the data assuming a constant mass and constant gradient of potential energy (i.e., constant force) felt by the polariton. This model clearly fails to describe the later time arrivals. Green curve: fit of the data assuming a constant force on the polaritons but allowing for the full dispersion relation $E(k)$ of the polaritons, which have an effective mass that changes at higher $k_{\|}$. Accounting for this changing mass improves the fit only slightly. Red curve: fit calculated by numerically propagating $x(t)$ and $k(t)$ according to the full semiclassical Hamiltonian of the LP.

To go beyond this simple model, we can recognize that the effective mass approximation breaks down for polaritons at high momentum. Due to the coupling of the very light mass photon and the heavy mass exciton, the dispersion of the polaritons at high momenta deviates from the effective mass measured at $k_{\|} = 0$. This is particularly true on the photonic side where the region near $k_{\|} = 0$ may have a mass on the order of 10^{-4} times the electron mass, while larger k-values at the same spatial point have a mass on the order of half the electron mass. By using the known polariton parameters (including the coupling strength between the exciton and photon, the cavity gradient and resonance position), we can relate the distance traveled to the initial energy and therefore the wavevector of the polariton. If the gradient

of the polariton energy is approximately constant, then the force on the polaritons will be constant and the time of flight will depend linearly on the initial wavevector according to $\hbar\partial k/\partial t = F$. Including the effect of the non-parabolic dispersion relation (green dashed line in Fig. 4.4) gives a slight upturn in the time of flight at farther distances. The effect of the increasing mass is to slow the deceleration. However, this model does not yet fully fit the data.

To accurately fit the data we must take into account the fact that the polariton energy in the strong coupling region near resonance quickly transitions from the rapidly changing photonic energy to the slowly changing exciton energy, and its mass changes by orders of magnitude. Thus we should not be surprised that naive models assuming constant mass and force will fail. However, the complicated energy of the polariton $E(x,k)$ prohibits a simple analytical solution to the time of flight as a function of the initial x and k. The most adequate solution to such a problem is directly deriving the equations of motion from the Hamiltonian, $\mathcal{H}(x,k)$, based on the known polariton parameters.

The classical Hamiltonian governing the ballistic motion of the lower polariton follows the form of Eq. (2.40).

$$\mathcal{H}(x,k) = E_{LP}(x,k) = \frac{1}{2}(E_x(x,k) + E_c(x,k)) \pm \frac{1}{2}\sqrt{(E_c(x,k) - E_x(x,k))^2 + 4g_0^2}, \tag{4.1}$$

where the explicit position (x) and in-plane wavevector (k) dependence of the QW exciton and cavity photon are included. These parameters must be carefully measured for the sample in question. Specifically, I assume that both the exciton and photon are linearly dependent on x and quadratically dependent on k for the relevant ranges of this problem. Therefore,

$$E_x(x,k) = E_0 + ax + \frac{\hbar^2 k^2}{2m_x} \tag{4.2}$$

and

$$E_c(x,k) = E_0 + bx + \frac{\hbar^2 k^2}{2m_c(x)} \tag{4.3}$$

where

$$m_c(x) = \frac{E_0 + bx}{c^2/n^2}. \tag{4.4}$$

The parameters E_0, a, b, m_x, n, and g_0 are all sample specific parameters, and were chosen based on careful calibration. Parameters used for this analysis are presented in Table 4.1. Here m_0 is the bare electron mass, and the mass of the exciton is included as an order of magnitude only. These values are ultimately slightly different than parameters used in later work, but are still within the uncertainty of the calibration models.

Table 4.1: Sample parameters used for calculation of time of flight

Parameter	Calibrated Value
E_0	1.6055 eV
a	2 meV/mm
b	14 meV/mm
g_0	5 meV
n	3.5
m_x	$0.5m_0$

Using Hamilton's formulation of the dynamics of classical motion, we can choose to follow the canonical variables x and $p = \hbar k$ as a function of time. The rate of change of each depends on the form of the Hamiltonian (4.1) and the instantaneous values of x and k [54, 78]:

$$\begin{aligned} \frac{dx}{dt} &= \frac{\partial \mathcal{H}}{\partial(\hbar k)} \\ \hbar\frac{dk}{dt} &= \frac{-\partial \mathcal{H}}{\partial x}. \end{aligned} \tag{4.5}$$

The initial x is known from the experimental conditions (i.e. the point of creation on the sample). This proper value can be inferred from low density $k = 0$ energy of the LP at the point of excitation.

Ultimately, we want to calculate a time of flight for a range of energy states (i.e. initial k) to propagate against the gradient until their final in-plane wavevector is zero ($k = 0$). The total distance traveled for each initial condition will be different, and this time of flight and distance traveled corresponds to the data in Fig. 4.4.

Since this non-trivial form of (4.1) does not lend itself to analytical solutions of $x(t)$ and $k(t)$, I numerically propagated these variables. First, I started with the known initial $x = x_i$ and select a small $k_i > 0$. Using (4.5) to advance $x(t)$ and $k(t)$ for small time steps Δt, I stop the propagation when $k(t) \approx 0$. The total number of time steps thus corresponds to the time of flight, and $x(t) - x_i$ is the distance traveled. I did this for a range of k_i to plot the time of flight for all the observed energy range.

4.3 ESTIMATION OF THE POLARITON LIFETIME

The long-range motion of polaritons in these samples suggests a significantly longer lifetime than has been observed in older samples. One might look for a direct measurement of the lifetime, but for various reasons this is difficult. We expect a lifetime on the order of 100 ps, so one might imagine that we can measure the decay of the cavity emission with a streak camera. However, as discussed above, if we generate the polaritons non-resonantly, this decay will mostly be detecting the thermalization time of hot carriers as they cool and become polaritons. On the other hand, resonant excitation of the polaritons is also problematic. For a measurement of the lifetime we could imagine resonantly exciting a polariton state and measuring the PL emitted from that state. There are several problems with this. First, there will be a large amount of reflected laser light, which can be reduced but not completely eliminated. Second, the lifetime of this state will mostly be affected by the dynamics of scattering into different polariton states. Third, with resonant excitation a coherent polariton state is produced which can have superradiant emission.

Another approach would be to measure the linewidth of the cavity photon mode, which will directly give a lower limit to the lifetime. The spectral resolution of our equipment, however, is not small enough to measure a 100 ps lifetime, which corresponds to a FWHM of less than 7 μeV. We measure a line width at the limit of our spectrometer resolution of 0.05 nm (100 μeV), which implies a lifetime of at least 7 ps.

4.3.0.1 LIFETIME FROM INTENSITY VS SPACE & TIME

Due to the difficulty of applying these more direct methods of measuring lifetime, we present here our best estimate of the lifetime from two different methods based on understanding the ballistic motion of these long-lived polaritons. Note that the lifetime of the polaritons is inversely proportional to their photonic fraction for photonic detunings. The lifetime is always longer in the excitonic region of the wafer, or in high-k states which have greater excitonic fraction. We are primarily interested in the intrinsic cavity lifetime, which is half the polariton lifetime at the resonant detuning point where the polaritons have 50% photon fraction.

The transport results discussed in the previous sections demonstrate the persistence of polaritons for hundreds of ps following non-resonant excitation—as seen in Fig. 4.4, the offset time for the arrival of polaritons reaches 400 ps. In addition to measuring the TOF in the above data, we also have measured the overall intensity reaching $k_{\|} = 0$ at various positions across the sample. Each final point corresponds to the number of polaritons that have survived the time of flight. We expect an exponential decay due to leakage of the photon mode through the mirrors, so the final population should be $n(t_0) = n(0)\exp(-t_0/\tau_i)$ where t_0 is the time of flight for that datapoint and τ_i is the lifetime of that state.

Figure 4.5 was determined by the following process: 1) the intensity $I(x)$ at $k_{\|} = 0$ was found for a range of distances x from the generation spot. Because of the gradient of polariton energy, each of these positions had a different energy. 2) The initial intensity $I_0(E)$ as a function of energy was found at the generation spot, from k-space data such as shown in Fig. 4.2(a). The higher energies correspond to higher momenta; these momenta drop to $k_{\|} = 0$ as the polaritons travel uphill. 3) The ratio $I(x)/I_0(E(x))$ was plotted as a function of the time-of-flight value t_0 found for each value of x. If we assume that the lifetime is approximately constant for polaritons in a certain energy range, then fitting this plot to an exponential decay gives the lifetime. The result of this lifetime fit gives a polariton lifetime of 200 ps, as shown in Fig. 4.5. We note that this lifetime includes all processes which remove particles from a ballistic path, including scattering from disorder. In addition to showing that the cavity lifetime is long, this measurement also shows that the disorder is very low.

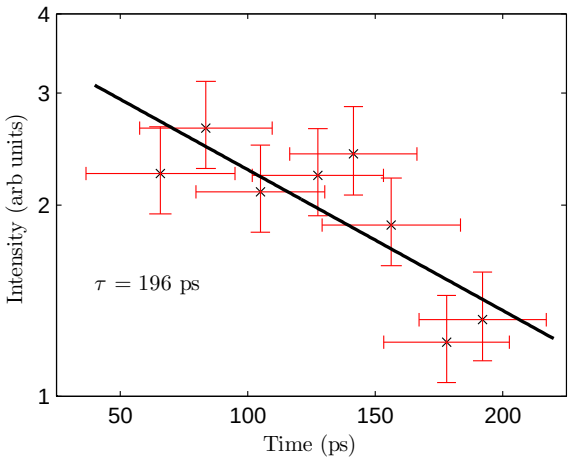

Figure 4.5: Lifetime of the polaritons based on the normalized intensity versus time of flight. The time values for the red crosses (data points) are the time-of-flight data presented in Fig. 4.4. The intensity values of these data points are the intensity detected at the point and time of measurement, normalized by the intensity at the same polariton energy taken from k-resolved data under the same conditions as Fig. 4.2(a) except that the pump spot detuning was the same as the time-resolved conditions. Since the emission at each spatial point corresponds to a single initial $k_{\|}$-state at the pump spot, this ratio gives the loss during the spatial propagation due to radiative emission and other scattering processes. The solid black line is a fit of a single exponential decay with lifetime of 200 ps (uncertainty is discussed in the text).

Since the calculation of the initial intensity ($I_0(E)$) was based on steady-state excitation rather than pulsed excitation, the normalization $I(x)/I_0(E(x))$ should only be considered qualitative at best. The different excitation conditions can in principle lead to rather different cooldown dynamics for hot carriers to fill the polariton states. In addition to this, there is large time- and intensity-uncertainty in the data plotted in Fig. 4.5. This large uncertainty is compounded from the fitting methods used in Section 4.2: the initial cooldown dynamics broadens and adds an intial cooldown time which is difficult to isolate. Therefore, this lifetime estimate should be taken only as an order-of-magnitude estimate. Still, a lifetime on

the order of 100 ps is a vast improvement over the few-ps lifetime assumed for first generation samples.

Of course, the polariton lifetime is not constant, but depends on the energy of the polaritons due to the dependence of the photon fraction on the detuning. Over the range of energies used in Fig. 4.5, we estimate that the photon fraction changed from about 90% to 75%. The fit value for the polariton lifetime of 200 ps therefore represents a cavity lifetime of about 150-180 ps.

4.3.0.2 LIFETIME FROM CW INTENSITY

An alternative way to measure the lifetime of the polaritons is to track the intensity change in k-space. The fit of the Hamilton's method theory in Fig. 4.4 gives $k(t)$ for each polariton energy. Therefore we can convert $I(k)$ to $I(t)$ for a given energy in data like that of Fig. 4.2 and extract a lifetime for each polariton energy from a fit to an exponential decay. This is shown in Fig. 4.6. Here the photon fraction ranges from about 95% at lowest energy to 85% at highest energy.

Over the range of detunings with reliable fits, these results show a lifetime of 200 $\pm$ 120 ps in a region where the polariton is mostly photonic. While we are unable to extract a trend of lifetime vs initial wavevector that clearly matches up with detuning dependent liftime or scattering trends, it is clear that these data support the conclusion that the cavity mode has a lifetime on the order of 100 ps.

It should be noted that, while this lifetime measurement is not directly a time dependent measurement, it is in some ways a more satisfactory means of observing the lifetime than resonant or non-resonant time-reolved observations. The CW density is lower than would be achieved in pulsed experiments, which circumvents superradiant effects that could affect resonant pumping or saturation and renormalization that could arise from resonant or non-resonant pulsed pumping. Also, since these polaritons propagate over such large distances, care would have to be taken in time resolving the motion of a single shot to keep the entire population within the imaging window, especially since gaining time resolution requires sacrificing a spatial, momentum or spectral axis.

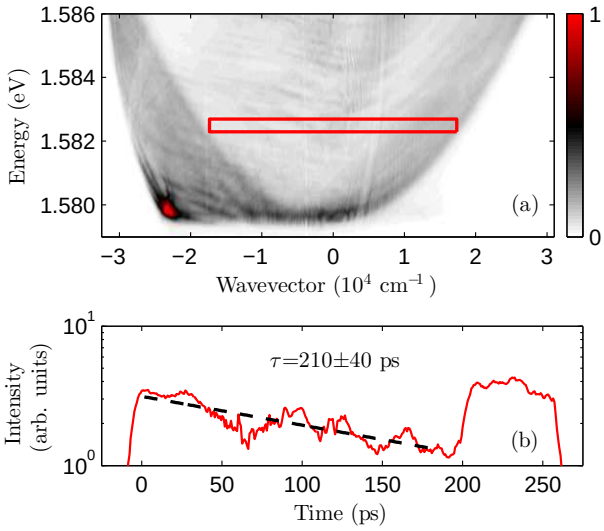

Figure 4.6: Lifetime measurement based on the steady state k-resolved PL intensity data. (a) The same data as presented in Fig. 4.2(b), with a highlight showing a selected detuning to generate an intensity profile. (b) The intensity profile for the selected detuning. Dashed line: fit to a single exponential decay in time. Note that the polaritons travel uphill and come back down. We therefore restrict the fits to times before the polaritons have returned back to the same place, which corresponds to $k_{\|}$ equal but opposite the initial $k_{\|}$. The polaritons moving downhill from the generation point are ignored due to noise in the data and the fact that they are observed for a short period of time which renders the fits unreliable. The time calibration in this plot is generated using the $k(t)$ prediction at each energy based on the initial conditions and applying Hamilton's method, as used for the fit of Fig. 4.4.

Since the polariton is moving at low density under the influence of a constant force, we can reason that the calibration of time of flight from lost $k_{\|}$ is deterministic according to the method introduced in section 4.2. Thus, careful calibration of momentum-space data and knowledge of the polariton parameters, such as coupling strength and energy gradients, yields a sufficient calibration of time of flight.

Furthermore, by recognizing that we are measuring momentum along the gradient (say

k_y) with a finite window in momentum in the perpendicular direction (k_x) determined by the spectrometer slit, we can be confident that we are properly tracking all the polaritons as they move away from the point of generation. This is in contrast to similar measurements in real space where we must recognize that the polaritons are not moving in 1-d along the spectrometer slit, but instead spreading out into 2π. In k-space, we know that k_x is conserved, and so those polaritons that initially fall into our collection will still be collected after being accelerated by the gradient in the y-direction. This argument is valid as long as the cavity gradient is properly aligned with the spectrometer slit.

4.4 HIGHER DENSITY: COHERENT FLOW

As the density is turned up, the polaritons experience a blue shift of their $k_{\|} = 0$ energy. This comes about due to exciton-polariton repulsion and possibly also to some degree due to a shift of the lower polariton branch due to phase space filling, which reduces the oscillator strength that gives the Rabi splitting between the upper and lower polariton branches. The excitons are produced by the same off-resonant pumping process that generates the polaritons—hot free carriers first form into excitons, and then some fraction of the excitons scatter down into exciton-polariton states. In many cases the exciton population can be 20 times greater than the polariton population [66]. The exciton population does not move long distances like the polaritons in these samples, however, because the exciton mass is about 10^4 times larger than the lower polariton mass. The exciton cloud diffuses at most a few microns from the laser excitation spot. This has been used [70, 79] to create user-controlled potential barriers for polaritons. In many works with short-lifetime polaritons, the exciton cloud is assumed to be everywhere that the polaritons are, and is called the "exciton reservoir," but in our long-lifetime samples, the polaritons can move very far from the exciton cloud.

In the experiments reported here, the polaritons are in an unbounded geometry—they can flow away from the excitation spot in the two-dimensional plane of the microcavity. It is therefore problematic to define Bose-Einstein condensation exactly. In a two-dimensional unbounded system, there is no "true" condensation [25, 80]. Rather, the fraction in low-

energy states near the ground state increases rapidly as the density increases, for a constant temperature, until a large fraction of the particles are in states with kinetic energy much less than $k_B T$. This is often called the "quasicondensate" [81]. The quasicondensate has many of the properties of a "true" condensate but has imperfect phase coherence.

In the case of a steady-state system with generation, decay, and flow away from the point of creation, the ground state of the system is not localized to just the region where the particles are created. As we have seen in the previous sections, the polaritons can travel ballistically hundreds of microns away from the creation spot. We therefore expect that the ground state will be a state that extends far from the creation point even while having a single energy [82].

Figure 4.7 shows the real-space and k-space energy distribution of the polaritons under similar conditions as Figs. 4.1 and 4.2, namely off-resonant excitation on the photonic side of the wafer, but with increasing excitation density. Two changes are notable as the density increases. One is that the energy of the polaritons shifts upward. This energy shift corresponds to the shift of the ground state of the polaritons at the point of creation due to their repulsion from the exciton cloud, discussed above. The second notable feature is that the energy distribution of the polaritons changes from a broad range of energies (Cf. Fig. 4.2(b)) to a single energy. This is due to the interactions of the polaritons in the excitation region, which allow them to thermalize. Even though they never perfectly thermalize when they are mostly photonic in character, as is the case here, they still have enough interactions to redistribute their energy distribution strongly toward the ground state. As seen in Fig. 4.7(h), they move at the same energy several hundred microns away from the laser excitation spot. Although the polaritons far from the exciton cloud probably have very weak interaction with each other, they still maintain the same energy. This extended, mono-energetic state is the effective ground state of the steady-state system, as discussed above. The two bright spots at $k_\parallel \sim \pm 1 \times 10^4$ cm^{-1} in Fig. 4.7(d) correspond to the velocity which the polaritons have after accelerating away from the exciton cloud, trading all of their potential energy for kinetic energy (cf. Ref. [82]). The polaritons moving uphill, with initially positive $k_\parallel$, slow down and eventually pass through $k_\parallel = 0$, which corresponds to the turnaround point seen in Fig. 4.7(b4). After passing through $k_\parallel = 0$, they have turned

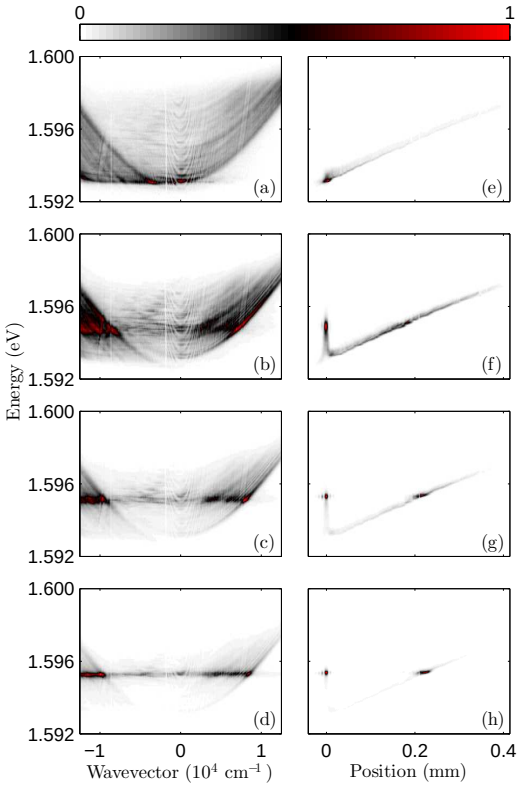

Figure 4.7: Frames (a)-(d): k-resolved PL from the polariton population at pump densities of 0.14x10^6, 12x10^6, 17x10^6, and 20x10^6 mW/cm^2. Data collected using a pump laser with $\lambda = 705$ nm focused to a 15 μm diameter pump spot where the LP detuning was -8 meV. Frames (e)-(h): $k_{||} \sim 0$ real-space-resolved emission at the same densities. Note that at the lowest density ((a) and (e)), all of k-space is occupied at the emission spot and the polaritons roll uphill and downhill as discussed above. At higher pump power, renormalization occurs at the pump spot, and a larger occupation builds up in the $k_{||} = 0$ state at the pump spot. The high occupation of a single state is seen as a monoenergetic line in k-space and two spots in the low-NA, real-space data, corresponding to the excitation spot and the turnaround point 200 μm away. In real space only two spots are observed because the polaritons in between, as well as those traveling downhill, are outside the angle of emission being imaged.

around and are moving in the opposite direction.

This monoenergetic quasicondensate acts as coherent wave. One way to see that the state is coherent is to simply note the spectral narrowing, to a peak with width about 0.2 meV. This width is actually broadened somewhat by the time averaging in our experiments. Fluctuations of the laser power lead to fluctuations of the exciton cloud potential energy height, which determines the polariton ground state energy. Another way to see the degree of coherence is by an interference measurement. Figure 4.8(a) shows the spatial pattern which is the result of interfering the $k_{\|} = 0$ emission from the creation spot with the $k_{\|} = 0$ emission from the turnaround spot 200 μm away. Figure 4.8(b) shows the fringe contrast as a function of delay time. This shows that the coherence time of the propagating ground state is approximately 40 ps, with an offset given by the propagation time $t_{prop} = 140$ ps from the creation spot to the turnaround spot. We believe that this interference measurement is also somewhat degraded by fluctuations of the pump laser power, which cause not only fluctuation of the energy of the polaritons due to the change of the potential energy of the polaritons due to the exciton cloud density, but also fluctuations of the spatial position of the turnaround point, i.e., the point with $k_{\|} = 0$ energy equal to that at the creation point.

This quasi-coherent flow can be easily understood as the propagation of a macroscopically occupied single wavefunction according to the system Hamiltonian. A simple approximation is to model the evolution of the quasicondensate using a 1D Schrödinger equation. While this involves approximations (for example, outflow to the sides will give a shorter effective lifetime), it makes the problem manageable and can recreate the major features of the observed real-space distribution, and allows us to make another constraint on the polariton lifetime.

To model this system we work in the effective mass regime for the lower polariton and model the spatial potential as a linear gradient with a Gaussian peak due to the exciton cloud at the excitation spot, as is visible in Fig. 4.7(f) and (g). This gives the general Gross-Pitaevskii equation

$$i\hbar\frac{\partial\psi}{\partial t} = \left(-\frac{\hbar^2}{2m}\frac{\partial^2}{\partial x^2} + U_0 e^{-x^2/\sigma^2} + Fx + U|\psi|^2\right)\psi - \frac{i}{2\tau}\psi + G(x), \tag{4.6}$$

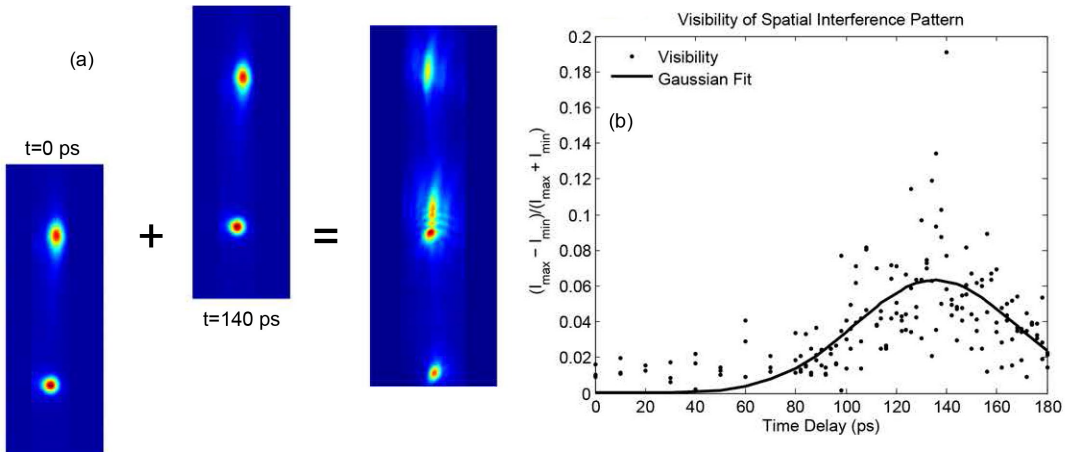

Figure 4.8: Interference measurements conducted by overlapping PL from the pump spot with time-delayed PL from the turnaround point in the medium density regime. These data were collected using a 34 mW pump laser with wavelength of 705 nm focused to a 25 μm diameter pump spot (7x10^6 mW/cm^2) where the LP detuning was -4.5 meV. Frame (a) shows the real space luminescence from the individual points and a sample interference pattern. Frame (b) plots the visibility of the fringes as a function of delay time. The fact that the greatest visibility is seen at 140 ps makes perfect sense as this is the propagation time for the polaritons to travel 200 μm from the pump spot to the turnaround point. The high scatter and overall low visibility of the fringes is primarily due to the instability in the pump laser, which leads to instability of the blueshift peak on which the polariton quasicondensate is formed and therefore causes both the condensate energy and turnaround point to fluctuate.

where U is the polariton-polariton interaction potential, τ is the polariton lifetime, and $G(x)$ is the localized polariton generation term (which can, in general, depend on the local polariton density, since a condensate of polaritons stimulates conversion of excitons into polaritons). The slope F is measured from the observed polariton gradient at low density, the Gaussian peak height U_0 is measured as the condensate emission energy, and the Gaussian peak width σ is determined from the pump spot size. The effective mass m can be found from low density k-space data (i.e. the curvature of the dispersion seen in Fig. 4.2(a)),

and we can justify using this effective mass because the mass changes minimally over the narrow energy range of this matter wave. In the low density limit, the polariton-polariton interaction is negligible, and this equation becomes simply a 1D Schrödinger equation with generation and decay.

The eigenstates of the system can be generated for a discretized space by numerically diagonalizing the 1D Schrödinger equation. Once we have a real-space representation of the eigenstates, it is trivial to decompose a matter wavepacket into constituent eigenstates and evolve it. The finite spatial grid and window leads to quantized states in the downhill direction where there is really a continuum, but artifacts created by this can be minimized if we ensure the space simulated is large enough that the state spacing is small compared to the energy range occupied by the condensate.

Using this prescription, we can evolve the motion of a pulse of a matter wave in real space and k-space with any lifetime. We can easily compare the characteristics of different lifetime particles by simply changing lifetime and evolving again.

Simulations with three different lifetimes are presented in Fig. 4.9 with comparison to an observed intensity profile with low-NA acceptance. Comparing the simulation results to k-space data also gives good agreement, indicating a good confidence in the simulation parameters such as effective mass and Gaussian peak width. As seen in Fig. 4.9, changing the lifetime has a strong effect on the relative height of the turn-around intensity peak to that at the generation spot. A very short lifetime will cause the uphill peak to vanish entirely, as polaritons decay before reaching that point, while a very long lifetime can make the uphill peak intensity comparable to the generation point intensity. The lifetime found here, 113 ps, is an underestimate of the polariton lifetime, because the effective lifetime for this model will be shorter due to outflow of the polaritons in the full 2D system, away from the 1D path considered here.

Related effects have been seen before with short-lifetime exciton-polaritons. If the laser generation spot is made very small, then there can be separation of the polariton motion and the exciton cloud even if the polariton lifetime is short. Refs. [31, 83] shows peaks at $\pm k$ which corresponded to acceleration away from the exciton cloud, as here. The group of J. Bloch [84] has shown mono-energetic propagation of a quasicondensate in a 1D quantum-

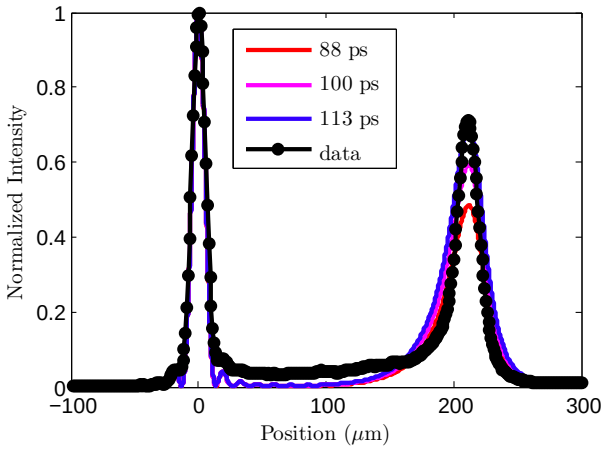

Figure 4.9: Comparison of GrossPitaevskii-simulated state evolution with observed real space intensity. The data were collected under the same conditions as Fig. 4.8. Note that three simulated lifetimes are presented for comparison.

confined wire, and Baumberg's group has seen similar behavior [70] with single laser spots in a 2D unbounded system.

Just as the resonant or non-resonant scheme can affect the polariton condensate formed, Richard et al. [83] demonstrated that the pump spot can also change the features of the condensate. It has been observed that a small excitation region can give rise to a condensate at finite k [82]. It is typically the case that polariton condensates form in regions where there is substantial renormalization, since the high carrier density that allows the condensate to form also causes a real blueshift of the polariton. Therefore, it is not surprising that a condensate of small size which is formed on top of a hill will flow outward.

4.5 HIGH DENSITY: TRAPPED CONDENSATE

While the quasicondensate described in the previous section does not exhibit a sharp threshold, as expected for a 2D system, at higher density we observe a much sharper threshold

transition to a trapped condensate with much greater coherence. This has been reported elsewhere [35]. In this case the polaritons scatter into a much lower energy state and localize at the energy minimum formed between the renormalized peak and the uphill gradient.

Although a first glance at the potential energy profile felt by the polaritons would indicate that they are not truly trapped, since the potential energy minimum shown in Fig. 4.7 is only in one dimension, there exists a process by which the polariton may in fact self trap, leading to a true 2D confinement. Two terms in the above Gross-Pitaevskii equation should be altered to take into account the interaction of the polaritons and the exciton cloud. First, the generation rate of polaritons can be written as $G(x)(1+\alpha|\psi(x)|^2)$, where α is a parameter, to take into account the fact that high polariton density will stimulate conversion of excitons into polaritons in regions where both exist. Second, the exciton cloud height U_0 can be written as $U_0(1-\beta|\psi(x)|^2)$, to take into account the fact that stimulated conversion of excitons will drop the potential energy height felt by the polaritons, since polaritons repel each other more weakly than excitons repel polaritons. The modified Gross-Pitaevskii equation is then

$$\begin{aligned} i\hbar\frac{\partial\psi}{\partial t} &= \left(-\frac{\hbar^2}{2m}\frac{\partial^2}{\partial x^2}+U_0(1-\beta|\psi|^2)e^{-x^2/\sigma^2}+Fx\right. \\ &\quad \left.+U|\psi|^2\right)\psi-\frac{i}{2\tau}\psi+G(x)(1+\alpha|\psi|^2). \end{aligned} \tag{4.7}$$

This highly nonlinear equation can have self-trapping solutions near the exciton cloud.

When multiple laser spots are used, an externally generated trapping potential can be created. Then even when the polaritons are generated in a region of the wafer where they are more exciton-like, they can undergo Bose condensation to a trapped state very much like the one reported in Ref. [35]. The increased lifetime of the polaritons allows for better thermalization of the polariton gas and truly equilibrium condensate theory to apply [79].

4.6 CONCLUSIONS

With increased reflectivity on the mirrors in these new high Q-factor microcavity structures, the polaritons demonstrate qualitatively different phenomena. Even in the low density regime

we observe clear signs of polaritons propagating much farther than in previous samples with or without 1D waveguide structures which promote long-range motion. At higher density we observe long-range, monoenergetic outflow which can be interpreted as a quasicondensate due to the Bose statistics of the interacting polaritons. The outflow from this condensate carries its coherence over a long distance.

These phenomena are a direct result of the increased lifetime of the polariton, and they also give us indirect ways to estimate the polariton lifetime. More direct methods of measuring the lifetime are diffucult due to the very narrow linewidth of the cavity photon and high reflectivity of the cavity. However, by looking at the decay of the polaritons with distance in real space and k-space we have estimated the lifetime of the polariton to be greater than 100 ps, about an order of magnitude longer than previous samples.

5.0 SLOW REFLECTION AND A DIRECT LIFETIME MEASUREMENT

This chapter is based on my 2015 Optica publication [36]. I make a more direct measurement of the polariton lifetime, improving on the results of Chapter 4. This yields the most direct measurement of a polariton lifetime ever presented in the literature, and clearly shows these samples to have a cavity mode lifetime an order of magnitude higher than the next best samples. Since the data is taken when the polaritons are turning around, it is also a novel demonstration of photons (dressed as polaritons) decelerating as massive particles, turning around and traveling backwards.

In the studies reported here, we have resonantly injected polaritons into a cavity and tracked them in time and space over millimeter distances as they experience a force due to a gradient of cavity width. Their motion is well described by the equations of motion for a moving mass under a constant force, that is, a parabolic trajectory. This can be called "slow reflection," as the photons gradually decelerate, turn around, and go back the other direction. From these measurements we can accurately measure the lifetime of the polaritons in our samples to be 180 ± 10 ps, over an order of magnitude longer than reported in previous works. This corresponds to a cavity leakage time of 135 ps and a cavity $\mathcal{Q}$ of 320,000. Additionally, we have shown the same effect with polaritons generated by direct two-photon excitation of the polariton states, allowing the possibility of modulation of two-photon absorption by a polariton condensate.

5.1 INTRODUCTION

Since the initial observation of exciton-polaritons in a strongly coupled microcavity in 1992 [85], a wide range of quantum many-body effects have been observed in polariton fluids such as Bose-Einstein Condensation [1,2], and superfluidity exhibiting quantized vortices [19] and solitons [14]. Most of these results have been interpreted in terms of nonequilibrium Bose gas theory, because the thermalization of the polaritons has been limited by their short cavity lifetime, on the order of 10 ps. Our recent results [34,35] have indicated that we can now produce structures which allow much longer lifetime, of the order of 200 ps. Here we report on accurate measurements of this lifetime using a unique method in which we inject polariton pulses at finite momentum into a microcavity and track their motion in time and space. This allows us to observe "slow reflection," in which renormalized light slows down to zero velocity, turns around, and goes back the other way. In addition to providing a measure of the lifetime, the long-distance propagation seen here allows the possibility of beam-like polariton-interaction experiments and all-optical switching methods over long distances.

As the technology of microcavity polaritons is now well established, much attention has turned to increasing the lifetime of the polaritons, to allow better thermalization and to allow propagation over longer distances. The lifetime of polaritons is a function of the intrinsic photon lifetime of the cavity and the fraction of photon in the polariton states. As amply discussed elsewhere [64], a polariton state $|P_k\rangle$ is a superposition of an exciton state $|e_k\rangle$ and a photon state $|\gamma_k\rangle$,

$$|P_k\rangle = \alpha_k|\gamma_k\rangle \pm \beta_k|e_k\rangle, \tag{5.1}$$

where α_k and β_k are the k-dependent Hopfield coefficients. The $\pm$ signs indicate that there are two superpositions, known as the upper and lower polaritons; in the experiments reported here we focus entirely on the lower polariton branch. At resonance, $\alpha_k = \beta_k = 1/\sqrt{2}$, while far from resonance the polariton can be nearly fully photon-like or exciton-like. This implies that the k-dependent lifetime τ_k of the polaritons is given by

$$\frac{1}{\tau_k} = \frac{|\alpha_k|^2}{\tau_{\rm cav}} + \frac{|\beta_k|^2}{\tau_{\rm nonrad}}. \tag{5.2}$$

For polaritons in our GaAs-based samples, the rate of nonradiative recombination τ_{nonrad} is negligible, so the lifetime is essentially entirely determined by the photon fraction and the cavity lifetime. In early polariton experiments [1,2,4,86], the cavity lifetime was on the order of 1 ps while the polariton lifetime was at most 10-15 ps, even well into excitonic detunings. This implied that polaritons would only scatter a few times on average before decaying. In recent experiments [67, 84], the polariton lifetime has been extended to about 30 ps.

Our previous work [34,35] has given estimates of the polariton lifetime in new samples of the order of 100-200 ps, allowing polaritons to propagate hundreds of microns to millimeters within the cavity and to show a sharp transition to a superfluid state. Because of the propagation of the polaritons to long distances away from the excitation spot, the configuration of those experiments made it difficult to get an accurate measure of the lifetime. A measurement spatially restricted to the laser excitation spot would give a severe underestimate of the lifetime, because the polaritons do not stay put—they feel a force due to the cavity thickness gradient that pulls them to one side, leading them to travel hundreds of microns from the excitation spot. Therefore, to accurately measure the lifetime, a measurement must track the polaritons in space as they move. The measurements reported here do just that. These measurements confirm the earlier estimates of the lifetime but considerably reduce the uncertainty.

5.2 METHODS AND RESULTS

The sample was arranged such that the gradient was aligned with the streak camera time slit, and then polaritons were injected at a large angle such that they moved directly against the gradient. The experimental setup is shown in Fig. 5.1. We used an objective with a wide field of view in addition to a large numerical aperture. A resonantly injected picosecond pulse of polaritons was tracked as it entered the field of view, turned around and traveled away. This occurs because the sample has a cavity thickness variation that leads to an energy gradient of the polariton. In simple terms, one can think of the motion of the polaritons as governed by energy conservation with the following Hamiltonian, which is just the same as

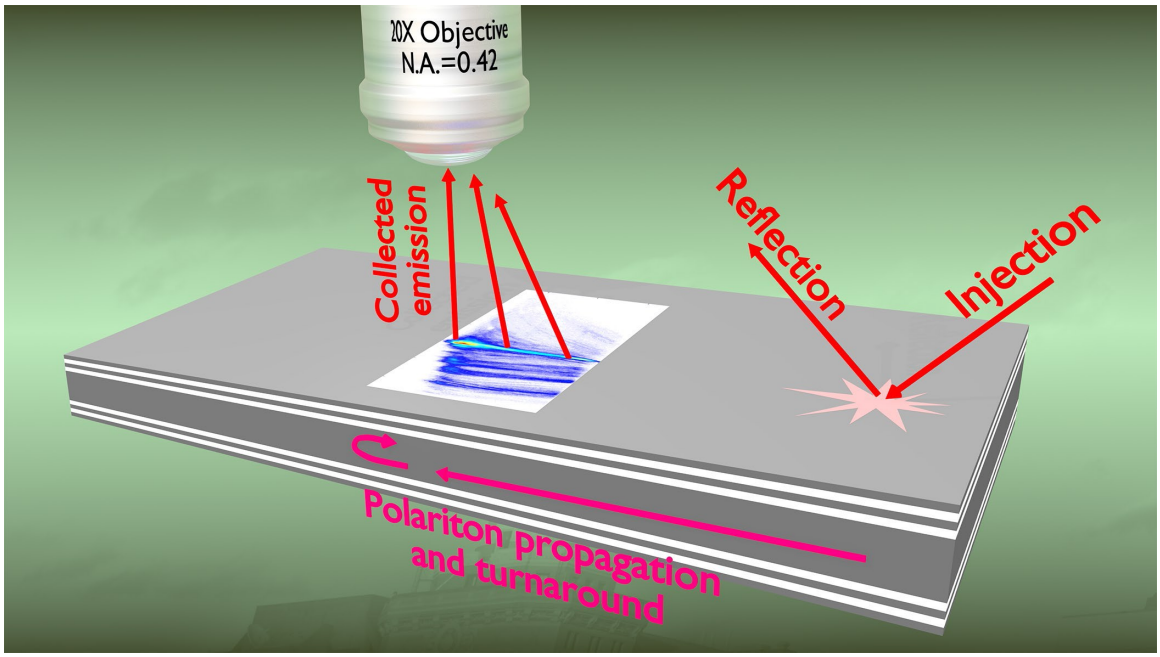

Figure 5.1: Diagram of experimental setup. The sample is a microcavity polariton wafer consisting of an active GaAs/AlAs cavity layer sandwiched between distributed Bragg reflectors (DBRs) on the top and bottom. See Appendix D for full details of the structure. A picosecond pulse is focused onto the sample at a large angle far outside the field of view of the collection optics. The orientation and wavelength of the beam are selected such that resonantly created polaritons flow directly uphill against the gradient and just turn around in the field of view. Emission from returning polaritons is not shown. Image is not to scale.

that of a massive object moving in a potential gradient:

$$H = \frac{\hbar^2 k^2}{2m_{\text{eff}}} - Fx. \tag{5.3}$$

Here m_{eff} is the effective mass of the lower polariton branch that we observe to be approximately equal to 5×10^{-5} times the vacuum electron mass in these experiments. The mass of the polariton depends on detuning and therefore changes with position and wavevector, but over the range we observe data it varies by only about 10%, so we can approximate it as constant. The force F is given by the gradient in space of the $k = 0$ cavity resonance energy,

and is approximately equal to 10.5 meV/mm for the section of the microcavity studied here. We will refer to "uphill" as moving toward higher cavity resonance energy (narrower cavity width) and "downhill" as moving to lower energy (wider cavity width). Work presented in Ref. [34] demonstrates how one can extend analysis to account for varying effective mass and potential.

This experimental setup utilizes the fact that the polaritons in these high-$\mathcal{Q}$ samples flow over a great spatial distance and change in-plane momentum rapidly. The lifetime of shorter-lived polaritons is more difficult to directly observe by streak camera measurements due to the overlap of any emission with the injecting laser. When polaritons are resonantly excited, they initially have the same energy and polarization as the laser photons. The initial polariton population therefore will have the same characteristics as the exciting laser and cannot be separated from it. Observing any other state (for example by looking at cross-polarized emission) will inherently measure the scattering time of the polaritons to enter that state. In this experiment, we rely on the fact that low-density, photonically detuned polaritons will flow ballistically from the point of injection to the point of detection in order to separate the observed luminescence from the reflected laser. To the extent that this motion is ballistic, integrating the population over the observed spatial region will directly yield the population decay of the polaritons. Unlike the case of observing luminescence from a different energy or polarization state than the initial population, this method directly follows the decay of a single population rather than relying on an average over many k-states.

The momentum of the injected polaritons is controlled by the angle of the laser which generates them. The angle of incidence used here was $\sim 42^\circ$, corresponding to an initial wavevector of 5.5×10^4 cm^{-1}. After propagating uphill for over two millimeters, the polaritons enter our spatial field of view and optical collection angle. Observing the polaritons far from the injection point reduces collection of scattered light from the laser excitation, and injecting polaritons at a large angle ensures that the reflected laser is outside the collection angle of the lens. As the polaritons flow against the cavity gradient they lose momentum, effectively exchanging in-plane kinetic energy for confinement energy, similar to a ball rolling uphill exchanging kinetic energy for gravitational potential energy. Since polaritons have a one-to-one relationship of in-plane momentum to external angle of emission for emitted photons,

watching the luminescence change emission angle while the gas of polaritons propagates gives us a direct observation of their slowing. Because the entire process is energy conserving, the injection laser, the ballistic polaritons and the emission all have the same wavelength. Once the polaritons reach a turn-around point, they flow back downhill and the emission angle increases to the negative direction.

Figure 5.2 shows time-integrated real space emission intensity from the microcavity near the turn-around point of the polariton gas. The coordinates in this image are such that the injection point is at roughly (0,0), and the force due to the cavity gradient is nearly directly toward $-x$. While polaritons were injected primarily in the $+x$-direction, the initial narrow spread of momenta in the y-direction led to a spread in real space after propagation over a long distance. At roughly $x = 1.7$ mm, polaritons are seen entering the field of view, which also corresponds to the acceptance angle of the optics. The brightest streak, directly vertical at y=0 mm, is the trajectory of the most intense part of the injected population which was peaked at zero momentum in the y-direction. Other bright streaks can be seen arcing to $\pm y$, and all states reach their respective turn-around points at $x \leq 2.25$ mm. The fact that there are bright streaks in this image rather than a smooth cloud suggests that the injection of the polaritons into the cavity occurs unevenly in momentum space. The asymmetry of the cloud between $-y$ and $+y$ may be due to a slight misalignment between the cavity gradient and the injection direction.

To measure the lifetime, the bright jet of polaritons was time-resolved using a Hamamatsu streak camera. To facilitate this, the sample was initially installed such that the gradient was aligned with the horizontal time slit on the streak camera. This enabled us to track a single jet of polaritons while they propagate against the gradient, turn around, and travel backwards, as shown in Fig. 5.3(a). The trajectory of the polaritons is easily seen in the data, which in this region is well described by a parabolic fit, as expected for the Hamiltonian (5.3), which is equivalent to that of a ball moving with a constant force due to gravity. These data directly demonstrate the in-plane velocity and acceleration of the polaritons during their trajectory. As discussed above, we estimate that the polariton mass changes by roughly 10% over the range of these data, which would lead to a minor correction to this model (as treated in Ref. [34]). One should note that this region of observation is already more than

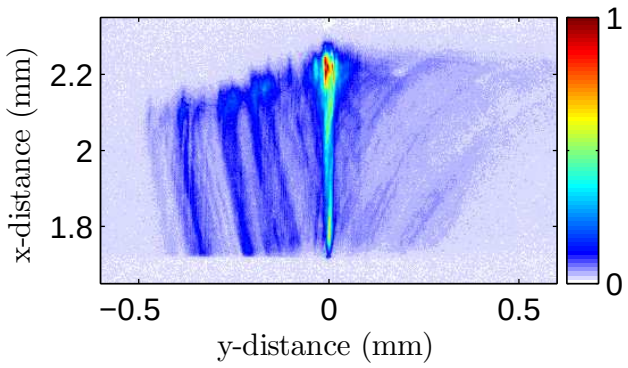

Figure 5.2: A time-integrated observation of passing polariton pulses. Coordinates are such that the point of injection is defined at (x,y)=(0,0), and the gradient is approximately toward $-x$. Polaritons approach this field of view from below and turn around at $x \approx 2.2$ mm before flowing back to $-x$. The sharp cutoff at x =1.7 mm is due to clipping in the spectrometer.

a millimeter and nearly 200 ps from the injection point, indicating that these polaritons are propagating farther and persisting longer than those in earlier samples, even without confinement in 1D structures, such as used in Ref. [84].

A simple analysis of this data yields the polariton lifetime after summing in the spatial dimension, as shown in Fig. 5.3(b). The data are well fit by a single exponential with a lifetime of 180 ± 10 ps. For the region of the sample observed in Fig. 5.3, the detuning of the polariton corresponds to the lower polariton approximately 75% photonic. (Although the polaritons move long distances, their detuning does not change much because they stay at the same energy.) From this we estimate that the cavity photon lifetime is approximately 135 ps, which corresponds to a $\mathcal{Q}$-factor of over 320,000.

It should be noted that this lifetime measurement may still be an underestimation of the lifetime. Close inspection of Fig. 5.2 reveals that individual jets of polaritons are still spreading out from the central jet. A population with some spread in initial momenta perpendicular to the cavity gradient must spread out horizontally while propagating uphill. The fraction of polaritons that move out of our field of view will lead to an underestimation of the lifetime. This error can be compensated for by using a narrower time slit to cut out

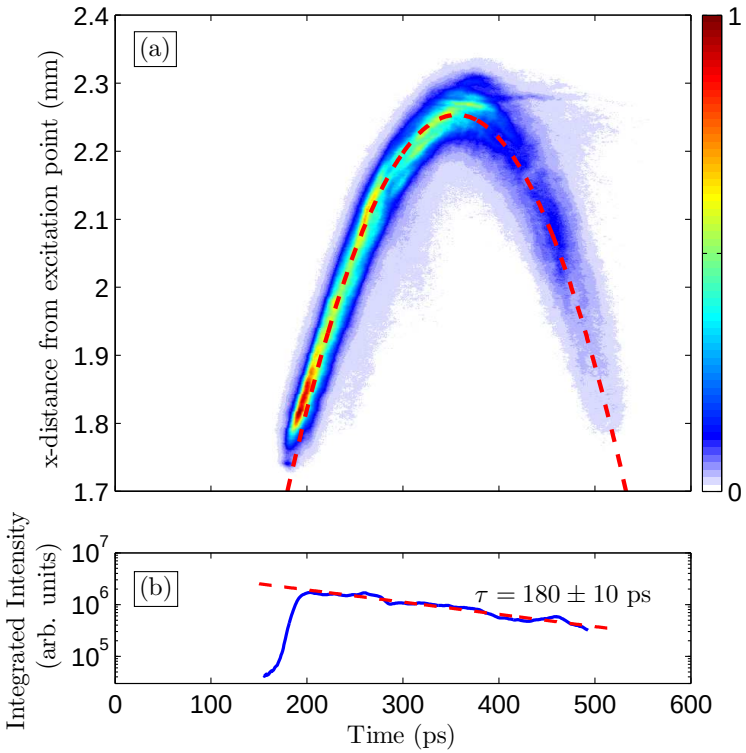

Figure 5.3: A time-resolved observation of passing polariton pulses in the region of -150 μm $\leq y \leq 150$ μm of Fig. 5.2. (a) Intensity vs x-distance vs time of the propagating polaritons. The dashed red line is a fit to the polariton motion as they feel a constant acceleration of 36 mm/ns^2. This acceleration is in good agreement with the expected value based on the known cavity gradient and the effective mass. (Over this 500μm region, the effective mass of the polariton is nearly constant.) (b) The polariton intensity of (a) summed in the x-dimension to highlight the exponential decay of the population. The data are well fit by a single exponential decay with lifetime of 180 ± 10 ps.

adjacent jets at early times; however, narrow slit widths can result in errors that will either underestimate or overestimate the polariton lifetime if the entirety of the main jet is not aligned with the time slit. In this experiment, data was collected over a range of slit widths from 50 to 300 μm with consistent results.

5.3 CONCLUSIONS

Polaritons can be viewed as "renormalized photons," especially in the region of the cavity where the the detuning makes the polaritons mostly photon-like. As mentioned above, the behavior we have seen here can thus be viewed as a type of "slow light," or "slow reflection," in which the photons decelerate from $\sim 3.5\%$ of the speed of light to a full stop and then go back the other way. This behavior is expected for light in a wedge-shaped cavity, without any need for the excitonic part of the polaritons. However, it has been hard to directly observe, because one must have very high $\mathcal{Q}$ and fast time resolution to track the motion of the photons. These measurements show that the photons can truly be viewed as having effective mass and feeling a force.

With such long distance propagation and long lifetime, it is now possible to construct experiments in which two or more beams are used and caused to interact. This could be used to directly measure the polariton-polariton interactions and also for schemes of optical gating using polaritons, as presented e.g. in Ref. [87].

6.0 ALL-OPTICAL SWITCHING WITH POLARITON TRANSISTORS

This chapter is based on my 2012 Applied Physics Letters article [37] in which I demonstrate the operation of a microcavity reflectivity switch. In Sec. 6.1 I give an overview of alternative switching schemes using the microcavity polariton system.

I present a method for all-optical switching using a microcavity exciton-polariton system. Unlike many other switching methods, we use a single wavelength for both the signal and gate, which means that this method could be used for cascading optical logic gates and amplification within an all-optical circuit. We resonantly pump the sample with a laser beam at an angle and probe the sample with a laser beam at normal incidence. Upon saturation of the exciton-polariton states, the normal-incidence resonance increases in energy to permit the probe beam to be transmitted through the sample. Experimental results demonstrate successful switching using a GaAs/AlAs microcavity. Switching times on the order of ten picoseconds and on/off ratios on the order of 10:1 have been observed, and we present design options to improve upon these.

While a great variety of schemes have been developed for optical switching over the past two decades, many methods utilize different wavelengths of light, and most require high intensity gate beams [88–91]. Some promising techniques can use relatively low gate-to-signal power ratios, but require gaseous media rather than solid state components [92–94]. A goal for switching schemes is to use relatively low gate powers to modulate signals of the same wavelength and similar (or greater) intensity. If the modulated signal is more intense than the gate, this allows for amplification of signals. Having both at the same wavelength eliminates the need for converting between wavelengths. Also, it would be ideal to use a solid state device for optical switching for robustness. With these considerations in mind, the microcavity polariton system is a good candidate for optical switching, as we will see

below.

The exciton-polariton system has been studied extensively over the past 15 years, mostly out of interest in bosonic statistics. The superposition of a bare exciton and photon creates a light-mass quasiparticle that obeys bosonic statistics. At low temperature and high density, these composite bosons can exhibit phase transitions such as Bose-Einstein condensation and superfluidity [1, 2, 29]. The strongly coupled light-matter system is also proving useful as a low threshold coherent light source [4]. For a more thorough review of exciton-polaritons, see [30] or [64]. The unique dispersion relation of polaritons and saturability of the exciton oscillator strength allows for our switching method.

Our sample consists of a microcavity with embedded GaAs/AlAs quantum wells (QW). The microcavity is made of two planar $\lambda/4$ distributed Bragg reflectors (DBR) made of alternating layers of $Al_{0.2}Ga_{0.8}As$/AlAs and a $3\lambda/2$ cavity. The QWs are located at the antinodes of the electric field in the cavity. This sample is similar to one used in previous work [2], but the number of quantum wells was reduced from 4 per antinode (12 total) to 1 per antinode in an attempt to reduce the carrier density threshold required to make the strongly coupled polariton convert into weakly coupled carriers and photons. This sample shows a Rabi splitting of 4.5 meV (cf. 7.5 meV in the previous sample [2]) at an exciton resonance of 772.25 nm.

Time-dependent experiments were conducted with a picosecond Ti:Sapphire laser. The output of this single light source was divided into both the signal and gate. Luminescence and reflectivity were time resolved on a Hamamatsu streak camera. The sample was held at T=10K in a gaseous helium bath.

The excitons are confined within the QWs, just as the cavity photon is confined between the DBRs in the same direction. Both are free to propagate in the plane, but the two have drastically different effective masses: $m_{\text{Exciton}} \approx 0.5 m_0$ and $m_{\text{Photon}} \approx 4 \times 10^{-5} m_0$, where m_0 is the free electron mass. This results in the coupled polariton having a dispersion relation that depends on the energy difference (detuning) between the bare exciton and photon, as shown in Fig. 6.1. Since the in-plane wavevector of an external photon must be conserved upon entering the cavity, the steep dispersion relation contributed by the photon enables us to create a population by resonantly pumping at an angle to the sample while not being

resonant with the normal-incidence polariton.

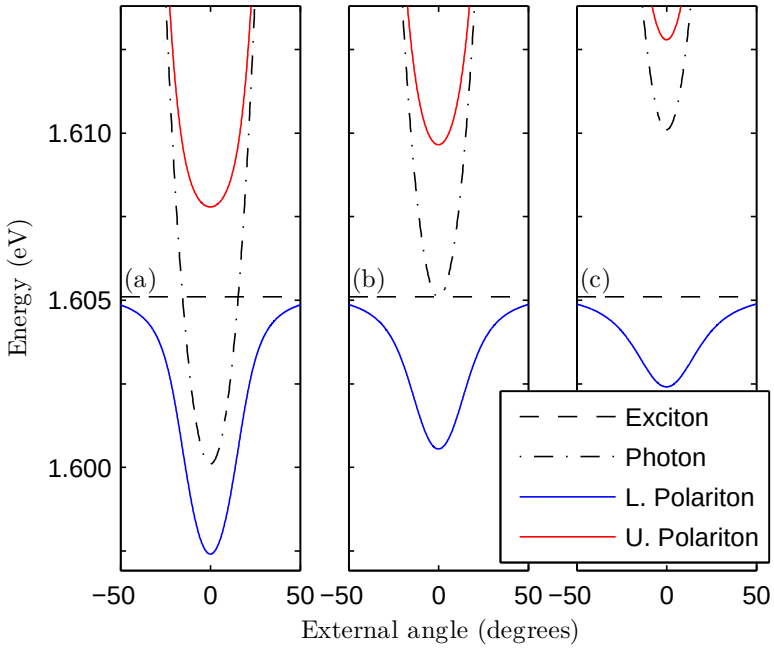

Figure 6.1: Polariton dispersion relation for various detunings: photonic (a), resonant (b) and excitonic (c). This shows that care must be taken when selecting a pump angle and energy for resonant injection.

Figure 6.2 demonstrates the simple geometry used in executing switching with a polariton microcavity. Fig. 6.2(a) shows the sample in the "OFF" state, which is probed by reflecting a signal beam normally off the sample. Due to the high reflectivity of the microcavity, virtually all of the light is reflected. Upon pumping the sample with a gate (as shown in Fig. 6.2(b)), the signal beam is transmitted though the sample.

While the topic of many-body effects of high density carriers in the microcavity is a complicated problem, it has been shown that in resonantly pumped polaritons, a transition to weak coupling occurs, in which the oscillator strength of the exciton is greatly reduced [26,95]. This can arise from free carriers screening the electric field that binds electrons and holes into excitons. Moreover, phase space filling and exchange between the fermionic electrons and holes that constitute the exciton can saturate the absorption at the exciton resonance.

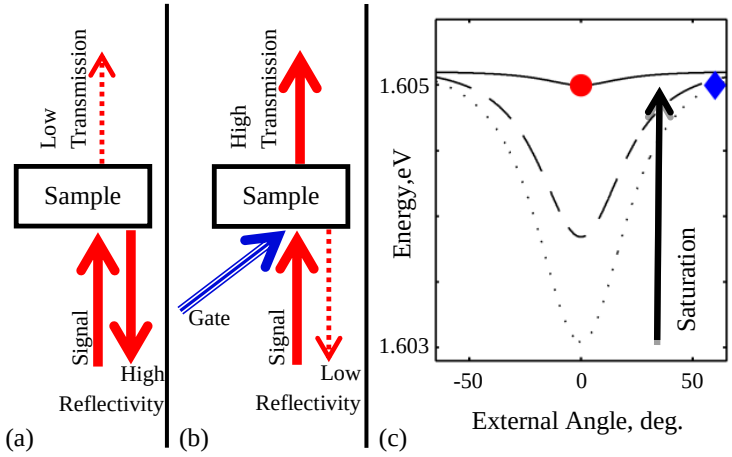

Figure 6.2: Geometry of polariton switching. Solid red indicates bulk of signal power incident, reflected and transmitted in (a) and (b). Blue line indicates gate beam in (b). The "ON" state corresponds to an increase in transmission or a decrease in reflection of the signal. Part (c) demonstrates that a progressive loss of coupling causes the lower polariton dispersion relation to flatten out. The dotted line indicates the polariton dispersion at low carrier density, such as when the gate is off or very weak–the switch is in the "OFF" state. The dashed line shows a partly saturated polariton dispersion following a moderate gate pulse—the state is still "OFF". The solid line indicates an almost fully saturated dispersion after a strong gate pulse–the signal is now resonant with the polariton and the state of the system is "ON". The red and blue symbols indicate the energy and angle of the signal and gate beams, respectively. The gate beam is resonant to the polariton at low density, while the signal will only be resonant at high density.

At very high density, carriers in the QWs can also lead to a shift in the index of refraction which causes the bare cavity photon energy to change [26].

We may observe the low-density resonance of the upper and lower polariton (UP and LP) by looking at the reflectivity. Upon creating a population of carriers either through resonant or nonresonant pumping, we can observe the resonance directly through luminescence from the sample. In this way we can observe the effect of the gate, which pumps the sample at an

angle, by looking at the photoluminescence (PL). Likewise, we observe the state of the switch by probing the sample with a normal incidence signal beam. This geometry is demonstrated by the signal (red circle) and gate (blue diamond) angles of incidence in Fig. 6.2(c).

Since the energy shift of the polariton is dependent upon the exciton density, we must create a real population of excitons to achieve modulation. At zero density the signal beam will be reflected since it is not resonant with the normal incidence polariton (red circle compared to dotted line in Fig. 6.2(c)). Upon creating a real population of excitons in the QW with the gate pulse (blue diamond) which is resonant to the polariton at an angle, the interaction leads to a decrease in oscillator strength. At moderate densities, the decrease in oscillator strength can partially uncouple the exciton and photon (dashed and solid lines). At densities above the saturation threshold, the polaritons can completely uncouple into the bare exciton and photon modes.

Switching will occur when the signal becomes resonant with the zero-angle LP, such as on the solid line in Fig. 6.2(c). At this point, the transmission of the sample will increase dramatically for the signal, corresponding to an "ON" state. Since the sample used for this experiment was grown on a GaAs substrate, we were not able to measure this increase in transmission and instead observed the decrease in the reflection of the signal.

Note that the signal and gate are degenerate in energy. As stated, both the signal and the gate were generated by the same picosecond laser source. This means a series of switches could create a logic circuit in which the output of one becomes the input of the next. This could prove vital for practical device applications.

Taking care to overlap the signal and gate pulses in space and time, we can observe a modulation of the signal by turning on and off the gate beam. It is necessary to use sufficient power and work at a proper detuning to renormalize the normal-incidence polariton energy up to or above the energy of the gate laser. Figure 6.3 shows the renormalization of the lower polariton resonance as a function of time. In some circumstances, the polariton can renormalize to higher energy than the laser, such that the signal will become resonant after a short delay following the gate pulse (compare Fig. 6.3(b) from (a)). The time-resolved photoluminescence shows a direct measurement of the shifting energy of the polariton. As the population of excitons and carriers decays in the quantum wells, the polariton energy

decreases to its low density value.

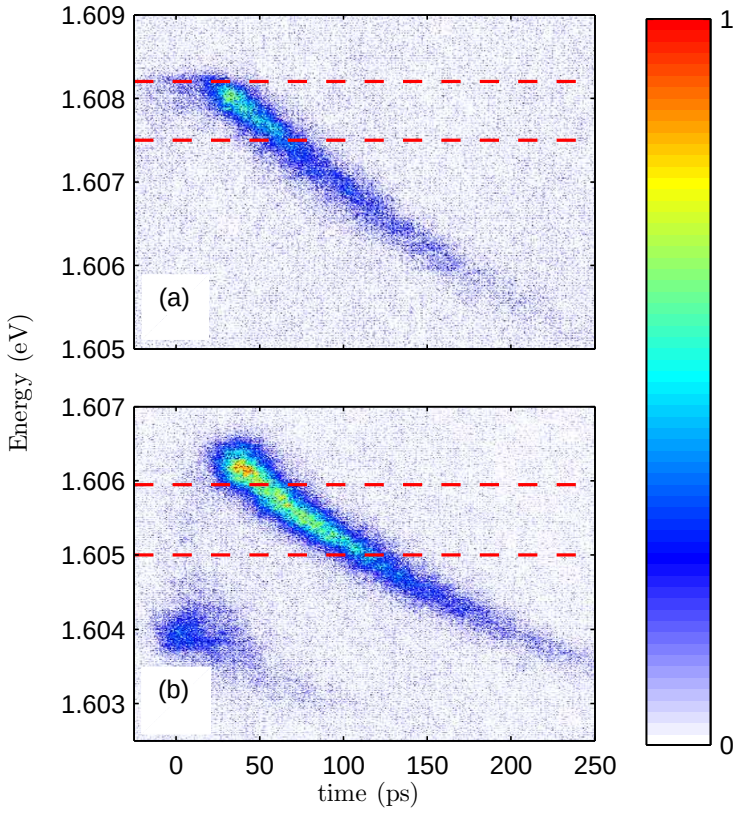

Figure 6.3: Time-energy resolved PL after a picosecond gate pulse. The red dashed lines indicate the FWHM of the exciting laser. Only PL near zero emission angle is observed. Data in part (a) were taken at 2 meV excitonic detuning, 100 mW average pump power and 10 μm spot radius at a 15^o angle. Data in (b) were taken at roughly 0-3 meV photonic detuning, 350 mW average pump power and 10 μm spot radius at a 20^o angle. Note that in frame (a) the maximum energy PL falls within the FWHM of the laser, but in frame (b) PL is observed at higher energy than the exciting laser at early times.

As an aside, we note that the data in Fig. 6.3(b) deserve further study. Here we observe saturation to shift the lower polariton energy higher than the exciting laser. If we consider the polariton at resonant detuning, then we know that the LP is lower in energy than the bare exciton or photon states. Upon saturation, the only available states are the bare exciton and photon states. This sudden increase in energy of the lowest momentum state (observed

PL is limited to normal incidence) indicates each quantum of excitation has increased in energy above the energy per quantum put into the system by the gate. At early times this corresponds to a laser cooling process when the bulk of the PL is emitted at higher energy than it was absorbed at. We find a sudden shift of the resonance with pump power, which may indicate a Mott ionization transition [96]. It may be useful to continue studying the dynamics of this polariton renormalization to better understand the Mott transition where excitons transition into free electrons and holes.

Figure 6.4 demonstrates picosecond modulation of the reflectivity of the microcavity following a gate pulse. Figure 6.4(a) is the reflectivity of the signal without the gate. Reflectivity is near 100% as the probe laser is far from being resonant with the unoccupied lower polariton. In Fig. 6.4(b), the signal follows a gate pulse (approx. 40 ps delay, 15 ps pulse FWHM) which has shifted the resonance of the polariton. At the time when the signal is resonant with the decaying polariton energy (compare Figures 6.3(a) and 6.4(b)), the reflectivity is reduced. The clear reduction in the reflected signal corresponds to an ON:OFF ratio of approximately 9:1 at the energy of the polariton or a modest ratio of about 3:1 when integrating over the pulse temporally and spectrally.

Note that this modest on-off ratio could be improved by many means. By using a transparent substrate we would see that the reflective on-off ratio R_{off}/R_{on}=9:1 would correspond to a transmissive ratio of T_{on}/T_{off}=180:1, given that R_{off} in these samples is 99.5%. Also, the picosecond pulse is spectrally broad compared to the polariton, meaning that the entire signal is not modulated simultaneously. Using slightly less reflective mirrors could match the cavity linewidth to the signal better, while also allowing more signal transmission so that less power is lost per gate.

In this sample it takes approximately 10-20 ps for the polariton energy to transit the FWHM of the signal, which puts a limit on how quickly signal pulses can follow one another. This time scale can be engineered by changing the reflectivity of the mirrors to modify the Q-factor of the cavity since the decay of the polaritons (Figure 6.3) is strongly dependent on the lifetime of the cavity mode. We note that it is not necessary for the entire population of excitons to decay away for the switch to resume its "OFF" state. It is only necessary for the density to drop enough for the polariton energy to fall below the laser wavelength.

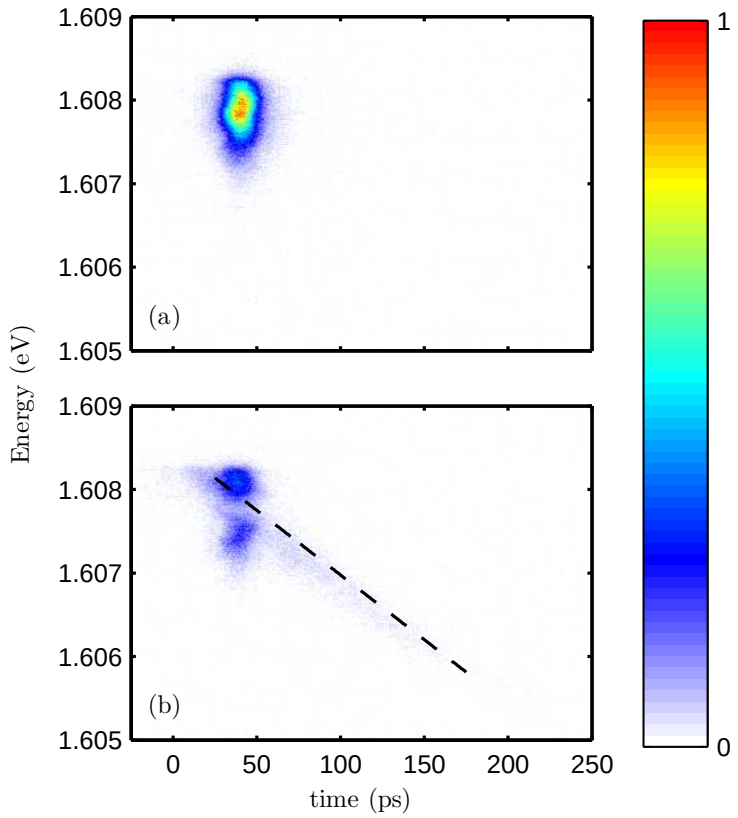

Figure 6.4: Time-energy resolved reflectivity of the signal pulse without the gate (a) and with the gate (b) modulating the sample. These data correspond to the same pumping conditions as Fig. 6.3(a). The dashed line in (b) indicates the decaying LP energy observed in 6.3(a).

Therefore repetition periods of tens of picoseconds are possible.

We have demonstrated that all-optical switching is feasible based upon saturation of a solid state microcavity polariton. Both the signal and the gate have the same wavelength, making this switch reasonable for logic circuits or other cascaded applications. We have achieved on:off ratios of 9:1, and theoretical switching times of 100 GHz. This technology can be extended to materials other than GaAs, and design changes are possible to improve functionality and efficiency. While these experiments were conducted at 10 K, alternative

materials with higher exciton binding energies may permit room temperature functionality.

6.1 ALTERNATIVE POLARITON-SWITCHING METHODS

Starting in 2010, the polariton community has made a concerted effort to develop a viable optical switch. Such distributed, competitive efforts have resulted in a diverse array of possible switching schemes. Here I outline a few works that have advanced the polariton switching field.

Since the initial drive in the field of microcavity polaritons has primarily been condensation studies, it is not surprising that some early switching schemes were focused on switching condensates. For example, a polariton condensate transistor was demonstrated in 2012 [69, 97]. This technique is based on generating and modulating thermal populations of polaritons by non-resontantly generating an exciton barrier. This is interesting in condensate studies, and such optically generated barriers are standard. However, non-resonant excitation requires the system to utilize two wavelengths, one for the polariton and a second for the excitation and barrier gating. Switching schemes reliant on two wavelengths are generally less cascadable, as the output of one switch must be wavelength switched to act as a new gate.

The first demonstration of proper single-wavelength optical switching in microcavity polaritons was also in 2012 [98]. This design uses the optical bistability of the polariton absorption. Much like our design, these researchers used one pump laser split into a signal and a gate. The gate beam was at high angle and resonant to the polariton at the point of injection. The signal was at lower angle, and therefore blue detuned from the polariton absorption at the point of injection. Due to the strong nonlinearity of the system, at high density of polaritons, the signal beam *is* resonant to the LP and will sustain a steady state population. Optical bistability is the fact that the absorption of the signal beam can turn on if the beam is intense enough to self-trigger the required high density of polaritons. In this two-beam switching scheme, the gate pulse seeds a population to trigger the absorption of the signal beam.

The scheme of Amo et al. was improved in 2013 [87] when researchers demonstrated cascaded switches capable of executing compound logic operations. In this scenario, separate nodes are constantly illuminated by signals below the bistable absorption threshold, resulting in essentially zero absorption. However, if one node is turned on by an external gate beam, polaritons are created at that node. Those polaritons then propagate in plane to the next node, and can act as a gate to switch that node on. Ballarini et al. show that these gates can function in 'AND' or 'OR' operation depending on the intensity of the illumination. Such as system truly allows for on-chip optical logic.

There are many other proposed and demonstrated polariton optical switches, as well. As early as 2010, Lew et al. proposed using Tamm-plasmonexciton-polaritons to make polariton integrated circuits [99]. These quasiparticles are the coupling of surface-plasmons to polariton states and can be manufactured by depositing metal films on the surface of a microcavity [100,101]. In this case, the plasmon-polariton states can be tuned by electrically biasing the metallic patterning, and under the proper conditions, excitations will be confined to these wires.

The group of Dr. Deveaud-Plédran has extensively demonstrated the operation of spin-sensitive switches in certain polariton geometries [102,103]. Since polaritons have two distinct spin states (inherited from the photon polarization), it is well know that switches will exhibit spin sensitivity [98]. However, in systems with asymmetry and reduced dimensionality, it is common for these otherwise degenerate spin states to be split in energy. This gives rise to interesting dynamics and state control [104].

Recent work has been done to understand the dynamics of spin flip operations in polariton systems [105]. This would make polaritons more interesting as optical memory elements.

7.0 PROGRESS ON ROOM TEMPERATURE MICROCAVITIES

In this chapter I will discuss the progress in the field of developing room-temperature microcavity polaritons. Great strides have already been made in developing GaN and organic-based samples, and polariton lasing has been demonstrated in these systems [8, 10, 106].

In order to advance the applicability of microcavity polaritons to technological uses, I have made efforts to develop room temperature strong coupling in microcavities based on materials other than GaAs. In particular, we have developed collaborations with scientists specializing in ZnSe and MoS_2.

7.1 SURVEY OF LITERATURE

Several systems have already been used to demonstrate room-temperature strong coupling microcavities including organics [107], high bandgap semiconductors [108, 109] and novel atomically thin materials [110]. In each case, an excitonic mode of the material is coupled to the photonic mode of a cavity. Unlike GaAs, however, these material systems generally do not lend themselves to producing monolithic photonic crystals. The lattice mismatch in GaN and ZnSe-based systems is too large to grow high quality DBRs without inducing large disorder in the excitonic material. Organics and dichalcogenides are easier to work with, as they can be directly deposited onto dielectric mirrors, but the excitons in these materials often exhibit relatively high disorder, efficient nonradiative decay, or interactions with the surrounding environment. Organics also have a tendency to degrade in light and atmosphere.

7.1.1 GaN

Gallium nitride is a direct-gap semiconductor. It should be noted that GaN proved to be of huge technological importance at the end of the 20th century. The 2014 Nobel Prize was awarded to I. Akasaki, H. Amano and S. Nakamura for advancing the technology of GaN to the point that it could be used to make blue LEDs. With this efficient blue light emitter, we can now generate white light an order of magnitude more efficiently than with incandescent or fluorescent lighting [111].

The high bandgap of GaN also results in tightly bound, high-oscillator strength excitons, making them of interest to strong coupling applications. While somewhat analogous to the bandgap hierarchy of AlAs-GaAs-InAs, the AlN-GaN-InN system lacks the convenient lattice matching of their lower-bandgap relatives. This has proved to be a headache for MBE and metalorganic chemical vapor deposition (MOCVD) specialists working to grow crystalline heterostructures of these materials.

As early as 2003, strong coupling was observed in GaN-based microcavities at cryogenic temperatures [112]. Initially, the MBE-grown quasi-bulk GaN was used in the active region and the cavity was a hybrid design of a bottom GaN/AlN DBR and a top SiO_2/Si_3N_4 dielectric DBR. The bottom mirror had to be made mostly out of GaN to minimize disorder in the crystalline layer of GaN in the excitonic region.

Improvements to this basic sample design resulted in demonstration of strong coupling at room temperature in 2006 [109] and room temperature polariton lasing in 2007 [8]. Polariton lasing was also demonstrated in a GaN quantum well based structure with an impressive Rabi splitting of 56 meV. Unfortunately, the GaN in these samples exhibits a high level of disorder [113]. This is not at all improved by the fact that the active layer must be grown on top of the bottom GaN/AlN DBR, which is highly strained. Sample designers must trade off the optical enhancement from increasing the periods in the DBR against the disorder that the accumulated strain causes.

The strain-induced disorder of the DBRs can in principle be reduced by a few means. Firstly, it could be greatly improved by producing lattice matched DBRs [114], but certain alloys in the AlN-GaN-InN family show miscibility issues. Otherwise, epitaxial liftoff and

subsequent deposition of dielectric DBRs has shown to be a viable way to avoid growing the GaN on top of disordered DBRs [9].

7.1.2 ZnSe

Similar to GaN, zinc selenide and related materials are another direct-bandgap semiconductor family with tightly bound excitons. These II-VI materials also suffer from a lack of lattice matched components, meaning all the problems of the previous section are still a factor.

It should be noted that the first evidence of polariton Bose Einstein Condensation was observed in CdTe, part of the (Cd,Zn)(S,Se,Te) family [1]. Even earlier in 1995 strong coupling was demonstrated at 10 K in CdZnSe QW microcavities with SiO_2/TiO_2 dielectric DBRs [115]. CdTe is the lower bandgap material in this family, making it suitable for cryogenic microcavity applications. However, in the early 2000's there were demonstrations of other II-VI microcavities.

Pawlis et al. developed a microcavity with CdZnSe QWs embedded in a ZnSe cavity that had a Rabi splitting of 44 meV at room temperature [108]. This cavity was was grown by MBE on a GaAs substrate that was subsequently chemically wet etched, permitting dielectric mirrors to be deposited directly on both sides of the ZnSe cavity layer.

While this sample is clearly strongly coupled at room temperature, the quality of the splitting is not superb–the upper and lower polaritons are roughly as broad as the Rabi splitting, meaning that the states are only barely resolvable. This broadening comes from a combination of the cavity quality, the interactions with phonons (26 meV at room temperature), and the disorder in the exciton states. Disorder in both the exciton and photon modes is made much worse because of the aggressiveness of the GaAs etch [116]. Following the etch, cracks can easily be seen throughout the ZnSe epilayer with domains on the scale of 200 μm. The unevenness of the surfaces and residual strain will respectively increase the disorder of the cavity and exciton modes.

The next improvement in II-VI microcavities came from improving the method of removing the substrate. By employing a MgS sacrificial layer developed at Heriot-Watt University [117, 118], Curran et al. fabricated CdZnSe QW based microcavities [119]. The MgS is

rapidly etched chemically, and the residual strain in the structure causes the epitaxial layer to peel off of the substrate. The epilayer must have a handle attached before the etch (wax is typically used) to give it structural support. After the etch, the layer can be adhered to a new substrate by van der Waals forces. The idea is to use a substrate that already has a dielectric DBR coating, and a second DBR can be deposited on the new surface [120]. The initial samples were developed to function at room temperature and consisted of the active cavity layer sandwiched between two dielectric mirrors rather than applying an optical coating directly to the epilayer [119]. Other samples prepared by this group were a metal-dielectric hybrid cavity and an actively tunable thickness cavity using a freestanding mirror and a nanopositioner to tune the air-gap separation between the two mirrors [120].

Each of these cavity designs is interesting, but not without drawbacks. It is likely that sandwiching the epilayer between two macroscopic substrates results in large strains in the sample or included dust particles detuning the cavity thickness. The metal-hybrid cavity is inexpensive and easier to make, trading one dielectric DBR for a simple silver 50 nm thick, but all metallic mirrors are lossy, reducing the maximum $\mathcal{Q}$ of the cavity. Moreover, since we generally want light to exit the cavity, either the metallic mirror must be thin enough to be transparent, or the dielectric mirror must not be much better than the metallic mirror. If the rate of absorption in the metallic mirror is much greater than the rate of leakage through the dielectric mirror, then there will be virtually no luminescence from the sample, and it becomes difficult to interact with the system. The last design offers unique flexibility as the cavity thickness can be actively changed while ensuring that the excitonic properties are unchanged. This is unlike our monolithic structures in which changing the detuning requires us to sample different regions of the excitonic material, which can therefore sample extra disorder. The downside is that the air gap typically requires a much longer cavity length, greatly reducing the enhancement of the optical mode at the location of the excitonic material, which reduces the overall coupling strength. Also, the increased cavity length reduces the free spectral range, which at some point will cause multiple cavity modes to simultaneously couple to the exciton mode.

A more recent advance in ZnSe microcavities relies on the careful growth of monolithic semiconductor DBRs. Although the II-VIs are not generally lattice matched, very thin ($\sim$ 1

nm) layers of strained alloys can be grown iteratively to postpone relaxation and dislocation formation. Sebald et al. created a semiconductor DBR out of ZnMgSSe layers (41 nm) alternated with superlattices (46 nm) of MgS and CdZnSe [121]. In this case, the ZnMgSSe and supperlattice layers act as the $\lambda/4$ DBR layers. The thin layers do not fully relax and create crystal defects, but since they are much smaller than the wavelength of light, the refractive index of the $\lambda/4$ layer is essentially an average of the indices of the two materials. However, this method requires extreme expertise and MBE deposition of quaternary alloys.

The 2000's saw a marked decline in the number of groups actively working with ZnSe, primarily because the technology of GaN was advancing well. Initially, both were sought after as prospective green and blue emitters, but GaN has since solved the problem of blue LEDs. Although this work was initially promising, the work by Pawlis et al. and Curran et al. was primarily ignored by the greater polariton community that preferred to work with clean GaAs or the more exciting GaN and organic materials. However, I believe that ZnSe still has the potential to produce excellent microcavity samples, which is why we have pursued a collaboration with Angelo Mascarenhas and Kirstin Alberi, who specialize in MBE of ZnSe at the National Lab of Renewable Energy. I discuss the status of this project in detail in Sec. 7.2.1.

7.1.3 ORGANIC SEMICONDUCTORS

Organic semiconductors are generally interesting as inexpensive optical materials. Rather than being crystalline bulk materials like GaAs, GaN or ZnSe, organics tend to be aggregates or polymers of molecules. While this typically results in larger inhomogeneous broadening of the exciton due to fluctuations in the environment, organic semiconductors can often be made by spin coating or dipping a substrate into a polymer solution rather than meticulously and slowly growing by MBE.

Of particular interest to microcavity polariton studies is the very high oscillator strength of excitons in organic semoconductors. Unlike crystalline, inorganic semiconductors in which the excitonic Bohr radius is much larger than the crystal lattice constant (this regime is referred to as Wannier-Mott excitons, with binding energies on the order of 10 meV), or-

ganic excitons are generally localized to a single molecule (or adjacent molecules). Such tightly bound excitons are called Frenkel excitons if the electron and hole reside on the same molecule, or charge transfer excitons if they sit on adjacent molecules. Since the oscillator strength of an exciton scales according to the overlap of the electron and hole as in Eq. (2.8), the tightly bound Frenkel excitons generally have much stronger interactions with light than Wannier-Mott excitons. For example, the oscillator strength of tetra-(2,6-t-butyl)phenol-porphyrin zinc (4TBPPZn, a material used in early organic microcavities) has an oscillator strength three orders of magnitude higher than individual InGaAs QWs [107].

Additionally, the new organic molecules can be engineered to suit a variety of needs. For example, the functional groups of a particular molecule can be changed to tune the absorption spectrum.

With the great variety of organic molecules and the ease with which they can be produced, organic-based microcavities have received a good deal of attention over the past 15 years. In 1998, Lidzey et al. demonstrated strong coupling at room temperature in an organic-based microcavity [107, 122]. The impressive Rabi splitting in this system, 160 meV, is 32 times larger than the splitting observed in inorganic microcavities of the same era.

Over the next decade, many other groups worked extensively with a variety of organic molecules to produce more microcavities [123–132]. Many of these results are presented at room temperature and in general the Rabi splitting is larger than any observed for GaAs or CdTe-based microcavities. More recently, researchers have achieved coherent light emission (polariton lasing) from organic microcavities [10], and some have even made claims of true Bose-Einstein condensation at room temperature in these samples [27].

In spite of the strong oscillator strength of these organic materials, the disorder and non-radiative losses of the exciton greatly influence the spectrum and dynamics of the polaritons. For example, the disorder has been shown to potentially lead to localization of polaritons [133] and influence the dynamics of thermalization [134]. The properties of the underlying organic material can also give rise to quenching effects [135]. Of particular interest to condensation studies and nonlinear optics is the Auger effect, by which one exciton annihilates by imparting energy to another electron or exciton [136]. This effect becomes more apparent at higher density, thus limiting the usefulness of some organic semiconductors

for condensation and nonlinear optics applications. Additionally, organics tend to degrade, especially when exposed to light and oxygen. When not being used, these samples are generally stored in darkness and under vacuum or inert atmosphere. Many groups studying organic microcavities are forced to constantly study different regions on the sample, as small regions of the organic can easily be damaged after being excited by intense laser pulses.

7.1.4 DICHALCOGENIDE SEMICONDUCTORS

A novel material more recently investigated as a potential for microcavity polaritons is the family of dichalcogenide semiconductors, such as MoS_2. Much like graphene, these materials form 2D layered sheets that stack together to form 3D bulk material via van der Waals forces between the layers. Like graphene, few or single layers can be exfoliated with tools as primitive as scotch tape. Unlike graphene, a single unit cell layer of this material is three atomic layers thick–the transition metal (in this case Mo) is sandwiched between two nonmetals (in this case S). The particular difference between the dichalcogenides and graphene results in the opening of a bandgap in the former while the latter has no intrinsic gap between the valence and conduction bands.

One particularly interesting property of MoS_2 and related materials is that the bulk material exhibits an indirect bandgap, while a monolayer (meaning three atomic layers, i.e. S-Mo-S) becomes direct. The exact mechanism for this transition is still under debate in the community, but it is clear that reducing the number of monolayers leads to an increase in the indirect exciton energy while the direct exciton stays relatively unchanged [137]. The quantum yield in PL from monolayer MoS_2 is four orders of magnitude greater than that in the bulk [138]. Indeed, this direct exciton of MoS_2 has an extremely high oscillator strength; a monolayer of this material is an excellent absorber and emitter of light and therefore an excellent candidate for microcavity polaritons.

MoS_2 was embedded in a tunable microcavity in the weak coupling regime at cryogenic temperatures in 2014 [139]. The cavity was shown to nicely enhance the luminescence but was insufficient to achieve strong coupling. This same group is currently working to show strong coupling by stacking multiple monolayers within the cavity (just as we stack multiple QWs at

the optical antinodes) [140], and many theorists are busy studying dichalcogenide-polariton condensates [141]. Already an MoS_2 microcavity was shown to be in strong coupling at room temperature [110,142]. However, all of these experimental demonstrations of strong coupling are tenuous at best. The Rabi splitting of the upper and lower polaritons is roughly equal to the broadening of each state, which means that the anticrossing is barely observable. Unfortunately, MoS_2 excitons interact strongly with the environment surrounding the monolayer: early results show much stronger luminescence from regions of suspended monolayer MoS_2 compared to monolayer MoS_2 resting directly on a substrate [138]. This can be understood because the tightly confined electron and hole will have some evanescent component extending into the surrounding substrate (or vacuum), and also the electromagnetic attraction between the two will extend into these surrounding materials. Therefore, inhomogeneity in the surrounding materials can greatly influence the MoS_2 excitons.

One current problem with monolayer MoS_2 is the existence of a strong trion state, or a bound state of two electrons and one hole [143]. This state is redshifted from the bare exciton by roughly 20 meV. This means that emission from the trion can easily interfere with standard polaritons formed from the standard exciton-cavity interaction, or that the full system may be treated as multiple excitons mixing with the photon. However, the physics of the exciton-trion-photon system will be more complicated than my theory covered in Section 2.3.5, since the trion is a fermion, while the exciton and photon are bosons. While some evidence suggest that the trion states do not couple as strongly to the exciton mode [140], the existence of the trion clearly influences the relaxation dynamics and the emission properties of the polariton states. One current interpretation is that the strength of the trion is due to the tendency of monolayer MoS_2 to pick up charge from the substrate it sits on. It is possible to actively tune the strength of the trion by applying a gate voltage to the MoS_2 [143].

To develop our own MoS_2 microcavities, I have cultivated a collaboration with Ken Burch at Boston College, as discussed in Section 7.2.2.

7.1.5 OTHER SEMICONDUCTOR AND DIELECTRIC STRUCTURES

It should be noted that other demonstrations of microcavities in strong coupling have been made with careful nanostructuring semiconductor materials to either enhance the excitonic binding energy or improve its oscillator strength with the optical mode in different cavity designs. For example, GaN nanowire microcavities have been used to achieve room temperature polariton lasing [144, 145]. ZnO nanowires embedded in a cavity showed strong coupling [146] and polariton lasing [145] at room temperature.

Interestingly, strong coupling has even been seen in isolated nanowires: the Fresnel reflections at the boundaries of nanowires can give rise to appreciable Fabry-Pérot or whispering gallery[1] optical modes. Strong coupling was observed in a ZnO nanowire in 2006 [147], and in 2014 polariton lasing was claimed to persist as high as 455 K [148]. Similarly, a GaN nanowire clad by GaInN quantum wells exhibited strong coupling of those exciton modes to the whispering gallery photon mode [149].

Recently, planar, quasi-bulk ZnO layers embedded in microcavities have also demonstrated polariton lasing [150]. Room temperature condensation in these structures is even claimed over a wide range of detunings, changing the photon fraction from 32% to 76% [151].

Theoretical work has shown that nanostructuring the dielectric mirrors can greatly enhance the confinement of the optical photon. This would enhance the Rabi splitting, possibly even enough to keep GaAs- or CdTe-based excitons in strong coupling at room temperature [152, 153]. Such dielectric mirrors require great expertise to produce, but one type–a "photonic woodpile"– was demonstrated as early as 2004 [154].

7.2 PROJECTS

To develop room temperature microcavity samples of our own, we have sought out several collaborations with groups specializing in the fabrication of promising materials, such as ZnSe, MoS_2, and porous silicon.

[1] A whispering gallery mode describes a wave that constructively propagates around the perimeter of a geometry, i.e. an integer number of wavelengths must fit within the circumference.

At this point, characterization of any preliminary samples is done primarily with white light reflectivity. Any luminescence data is collected by following non-resonant excitation by a 405 nm diode laser.

7.2.1 ZnSe

We have begun to collaborate with Angelo Mascarenhas and Kirstin Alberi, who specialize in the MBE of GaAs and ZnSe at the National Renewable Energy Laboratory (NREL). Using MBE, Kirstin is able to grow CdZnSe QWs embedded in a ZnSe spacer on a GaAs substrate. By chemically etching the GaAs substrate, we can isolate the active cavity layer. Here at Pitt, we are developing the expertise to deposit SiO_2/Si_3N_4 dielectric DBRs using plasma enhanced chemical vapor deposition (PECVD). The outline of the procedure is shown in Fig. 7.1.

The basic design is to use a $Cd_xZn_{1-x}Se$ strained QW embedded in a bulk ZnSe cavity spacer. The Cd reduces the bandgap of the CdZnSe alloy such that the bulk ZnSe acts as a barrier for the electrons and holes. Therefore, tuning x, the mole fraction of Cd, and the QW thickness allows some variation of the confinement energy and the actual transition energy of the exciton. However, predicting the exact confinement properties is difficult due to the lattice mismatch between CdSe and ZnSe.

The best method to fabricate the first microcavity is to fabricate a series of test QW structures, then observe the QW experimentally. The cavity thickness and PECVD DBR designs will then be adjusted to suit the experimentally determined exciton spectrum.

I characterized initial QW samples, produced by NREL in Spring 2015, to characterize the exciton energy and effects of the GaAs etch procedure. These samples were produced while running initial tests of the newly assembled II-VI MBE system at NREL. Quasi-bulk ZnSe samples (few hundred nm) were grown for characterization of the crystalline purity, which can be done by x-ray diffraction (XRD) at NREL, and optical characterization at both NREL and Pitt.

Here several samples were prepared with different Cd concentrations or other conditions. Of this initial set of QW samples listed in Table 7.1, each had 10, 5 nm thick QWs

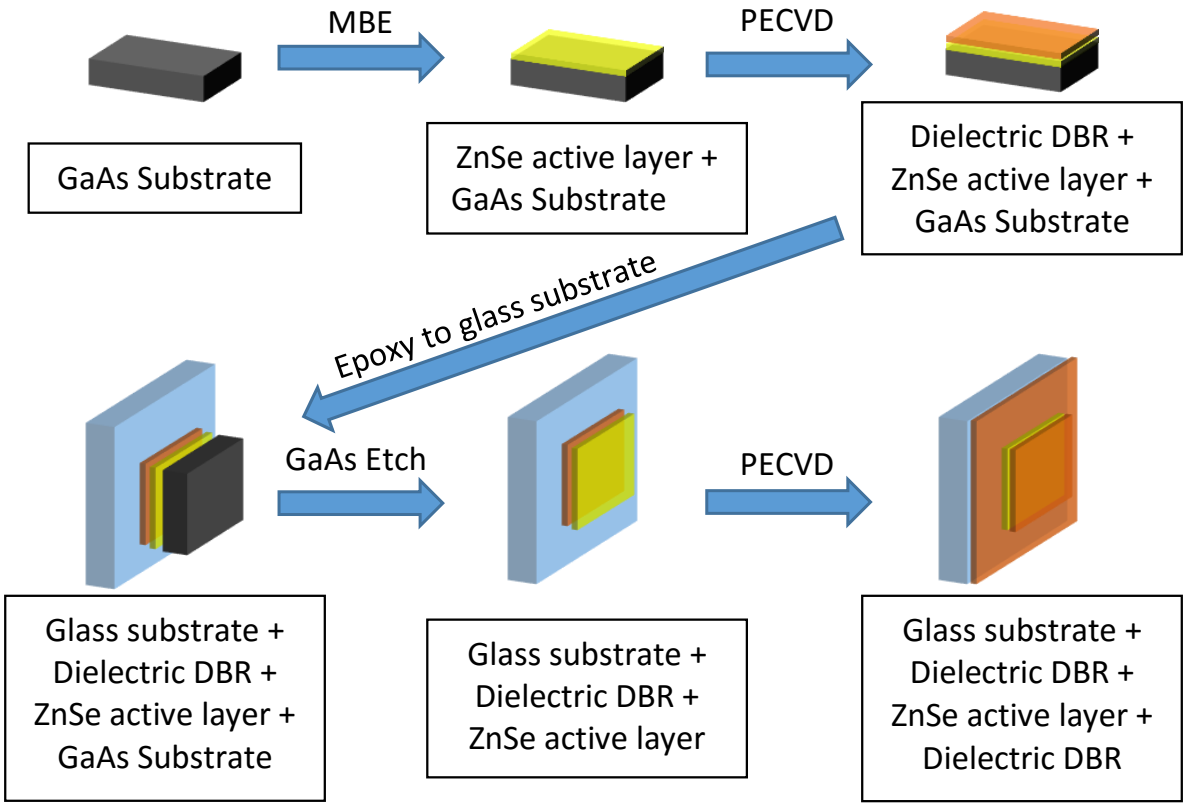

Figure 7.1: Planned procedure to produce ZnSe-based microcavities. MBE will be done at NREL, while PECVD and etching procedure can be done at either Pitt or NREL.

($Cd_xZn_{1-x}Se$) separated by pure ZnSe barriers according to the data in the table. Additional parameters, such as broadband illumination during MBE growth or automated shutters controlling the growth timings, are listed in the table as well. Samples appended with an "LO" are epoxied to glass substrate and have had the GaAs substrate etched off.

At Pitt, I characterized these samples at room temperature by exciting them with a 405 nm laser and observed the luminescence over the full visible spectral range. Figure 7.2 shows a comparison of a lifted-off sample compared to the same MBE growth without the GaAs substrate removed. Aside from the strong QW emission at around 2.4 eV, there is a high energy peak around 2.7 eV and a broad low energy continuum around 1.75 eV. The

Table 7.1: Spring 2015 ZnSe Samples

Sample	Cd conc.	Barrier (nm)	Notes
26_0054	19%	5	Manual timing
26_0054-LO	19%	5	Manual timing, removed GaAs
26_0063	28%	25	Automated timing, broadband illumination
26_0064	28%	25	Automated timing
26_0064-LO	28%	25	Automated timing, removed GaAs

high energy peak is bulk ZnSe luminescence, and in this case the liftoff process appears to suppress this peak.

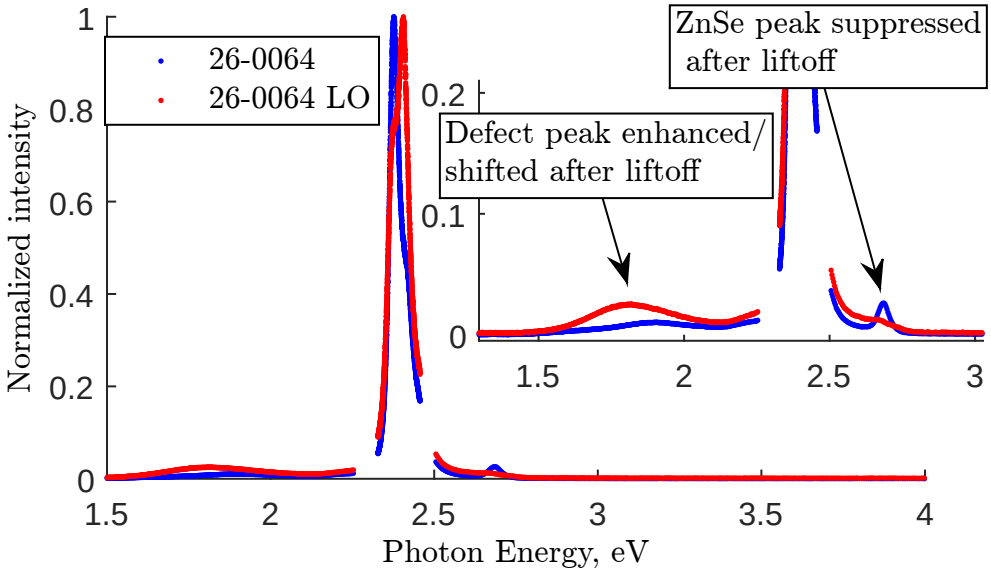

Figure 7.2: Normalized luminescence from two CdZnSe QW samples. Both samples were from the same growth, but 26-0064 LO has been epoxied to a glass substrate and the GaAs has been removed. Peak heights are qualitative only, as the data around 2.4 eV is from a higher resolution system; this gives rise to the gaps in the data around the 2.4 eV peak. Inset: closeup to show detail of low-intensity peaks.

The broad low energy peak is more mysterious. We suspect it is due to defects at the ZnSe-GaAs interface. In this data, it is seen to be enhanced relative to the QW emission. It is unclear if that enhancement is due to the fact that we are exciting and observing through the substrate side in the LO sample as opposed to the top ZnSe surface in the non-LO sample. Further work must be done to understand the origin of this state, and what affect the liftoff has on it.

Figure 7.3 compares the QW emission between these samples. Clearly the samples with higher Zn content in the QWs (26_0054) exhibit higher energy exciton emission, since higher content of Zn correlates with an increased bandgap of the alloy. Interestingly, all spectra show a resolvable second exciton line separated by about 40 meV. This is possibly due to the different confinement energies arising from the light and heavy hole confinement in the QWs [155]. Possibly one way to increase the splitting between these states would be to decrease the QW thickness, as long as the finiteness of the barriers doesn't become an issue.

The liftoff process (LO samples) does not seem to have destroyed the QWs, but it has made a minor change to both the 26_0054 and 26_0064 samples. In the 26_0054 sample, there is a slight redshift of the exciton mode, but less than the linewidth of the mode. The 26_0064 sample, on the other hand, seems to have no spectral shift of the modes, but there is a change in the relative emission strengths between the two exciton modes.

A set of sample designs is presented in Appendix E.

7.2.2 DICHALCOGENIDES

A second major effort for us has been in MoS_2 microcavities by collaborating with Ken Burch and Marcel Hoek at Boston College. They have been developing expertise in the exfoliation and manipulation of MoS_2. With our guidance in designing dielectric cavities, they have taken the lead on PECVD deposition of the DBRs and spacer layers and the exfoliation of MoS_2, while I am taking the lead on optical characterization and improving the cavity design.

The basic design for microcavity fabrication is outlined in Fig. 7.4. A dielectric DBR is produced tuned to the MoS_2 resonance. Specifically, the structure must terminate with a low

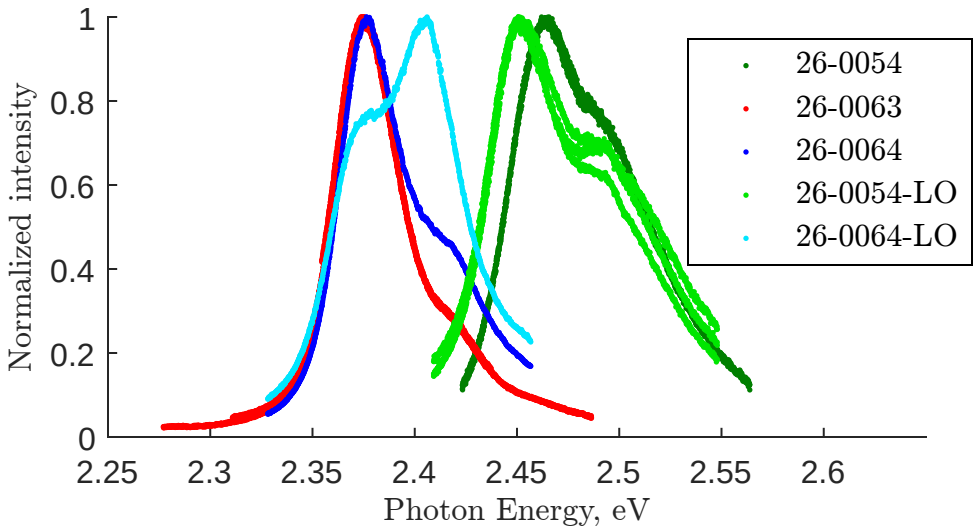

Figure 7.3: Normalized luminescence from all CdZnSe QW samples at room temperature. The higher Cd concentration samples show a higher energy exciton, as expected since all QWs are the same thickness. All samples exhibit a two-lorentzian profile, suggesting that multiple exciton states will be relevant to the polariton. The excitons have a FWHM of roughly 35 meV. Three spectra from different spatial points on 26-0054-LO are plotted to demonstrate disorder across the sample.

index material such that an antinode of the optical mode sits at the surface of the structure. The most clear description of this is that this top layer acts as half a cavity spacer, and if the same structure were inverted and placed on top, the two $\lambda/4$ layers together would form a $\lambda/4$ cavity. The following step is to exfoliate monolayer MoS_2 on the surface of this half-cavity.

Upon exfoliation of MoS_2, the Boston College group must optically investigate the exfoliated material to locate monolayer pieces. Exfoliation generally yields a jumble of pieces, some of which may be single layer. Since adding the top mirror will (by design) obscure or modify the emission from the exciton, monolayers must be identified prior to deposition of the top mirror.

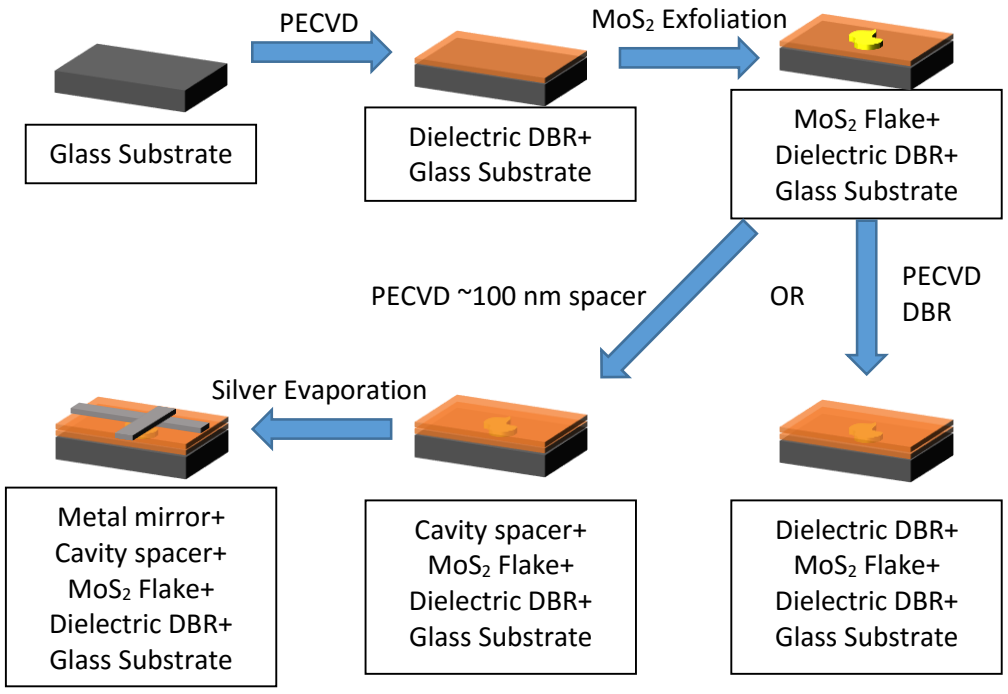

Figure 7.4: Planned procedure to produce MoS_2-based microcavities. Photolithography can be used to pattern silver on the surface to aid in locating the monolayer flakes. PECVD, exfoliation and evaporation will be done at Boston College. Optical design, characterization and study will be done at Pitt.

At this point there are two separate possible designs. The simplest design requires using PECVD to deposit the other half of the cavity and a DBR in one growth. This produces the highest quality cavity possible, as both mirrors are dielectric.

The second fabrication design is to use a metallic mirror instead of a dielectric one. This requires first depositing the second half of the dielectric cavity spacer via PECVD, then evaporating silver on the surface. If the top mirror is a conductor and only about 100 nm from the MoS_2, this would allow us to easily gate away the excess negative charge in the dichalcogenide to reduce the effect of the trion state (we could easily use ITO glass as the substrate). Unfortunately, the inclusion of metal reduces the maximum Q-factor of the

cavity. Even thick silver is only about 96-98% reflective. Moreover, since any loss through thick silver would be absorbed, we must interact with the polariton through the back DBR. This requires that the back DBR be roughly no more reflective than the metallic mirror. For comparison, our long-lifetime GaAs microcavities utilize DBRs that are 99.99% reflective.

One additional bonus of using metal as the top mirror is that it can also easily be patterned as a means to locate the small flakes. This is emphasized in Fig. 7.4 by the silver being deposited as a cross where the small MoS_2 flakes are.

The first trial of MoS_2 microcavities used the metal-dielectric hybrid cavity design. Figure 7.5 shows the characterization of the exfoliated flakes. This charaterization must take place prior to the top mirror deposition, and was done at Boston College. It is clear that there are many flakes in this vicinity that are likely not monolayer. Once the top mirror has been applied, it is difficult to know whether you are looking just at multilayer flake, or actually a monolayer flake that is just not strongly coupled to the cavity.

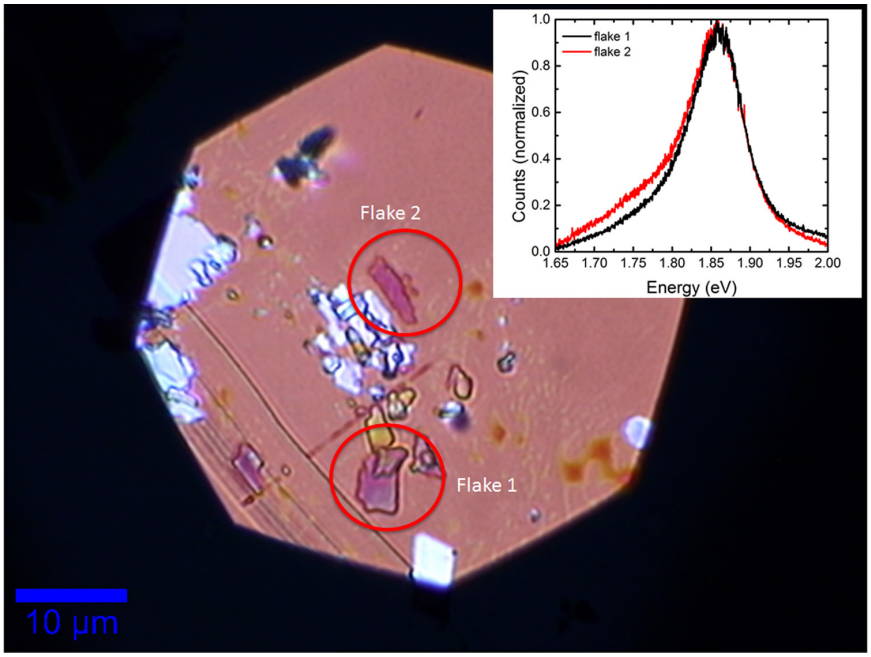

Figure 7.5: Characterization of MoS_2 prior to cavity finishing. The blue scale bar measures 10 μm. Inset: the luminescence spectrum of the two located flakes.

These data also highlights just how small these exfoliated flakes tend to be. Certainly we

won't see long range motion of polaritons in these samples. The small flake size also requires a redesign of optical imaging design to achieve high enough magnification to clearly resolve these few micron flakes–something not necessary when working with 2D semiconductor samples.

Figure 7.6 shows the silver patterning on the microcavity to enable us to locate the MoS_2 flakes. However, even with such a detailed map, it can be difficult to optically locate these precise flakes when they are surrounded by so many other flakes and defects in the silver.

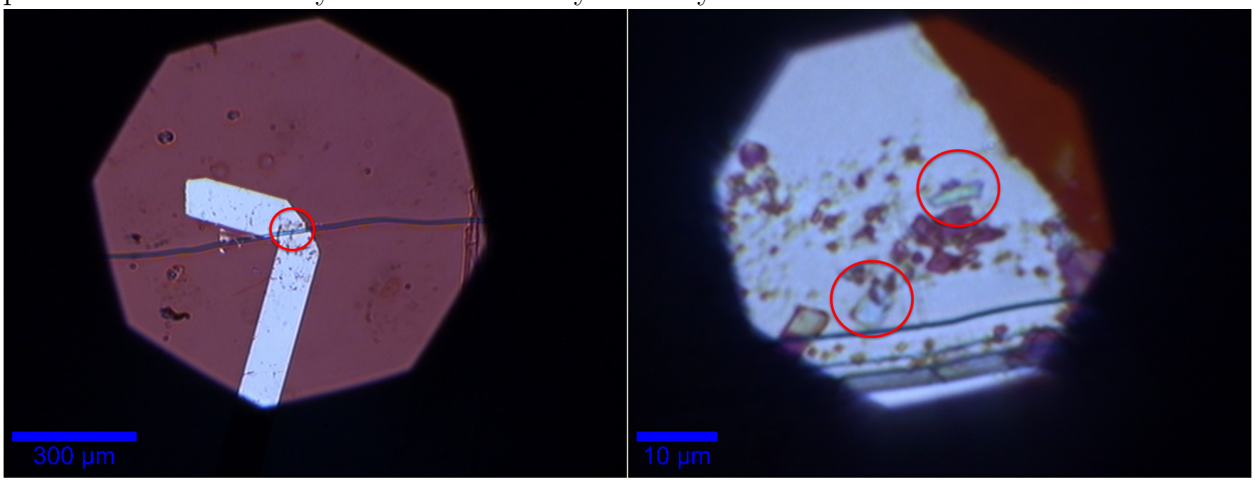

Figure 7.6: These maps were provided by Marcel Hoek at Boston College to assist in locating the MoS_2 flakes after deposition of the top mirror. He deposited the silver (blue-silver color) in a pattern to aid in locating micron-scale flakes. The horizontal line seen running through the left frame is a crack in the mica substrate. The scale bars measure 300 μm and 10 μm in the left and right frames, respectively.

Initial optical studies of this sample showed no sign of strong coupling around the flakes of MoS_2. It is likely that this is due to the large detuning of the cavity mode from the exciton, as seen in reflectivity data shown in Fig. 7.7. The dielectric DBR can easily be characterized by measuring reflectivity away from the silver, as seen in the lower magnification frame of Fig. 7.6. Data are presented from multiple points for both the DBR and the full cavity to give indication of the variation due to disorder.

A cavity mode is clearly seen at 686 nm, over 50 meV detuned from the expected exciton mode. This is larger than the Rabi splitting observed in a higher $\mathcal{Q}$ cavity [110], so we

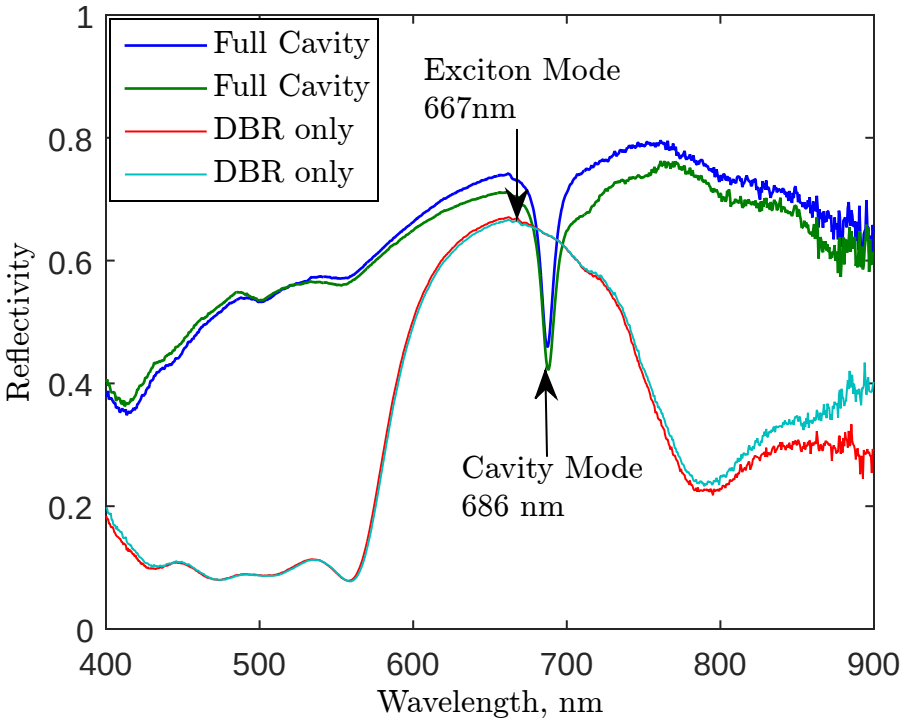

Figure 7.7: Reflectivity of the MoS_2 metal-hybrid cavity. Reflectivity is measured for the full cavity as well as the dielectric DBR away from the metal capped cavity. The cavity mode is highlighted with an arrow, as is the exciton mode as measured prior to deposition of the silver mirror shown in Fig. 7.5.

should expect this metal-hybrid microcavity to be in weak coupling.

To further investigate these samples, I excited them with a 405 nm laser and observed broadband luminescence. Unexpectedly, there was broadband luminescence that seems to be filtered by the cavity reflectivity, as seen in Fig. 7.8 (a). The peaks in the PL are consistent with the reflectivity minima of the stopband, and closing the Fourier-filter iris sharpens these peaks appropriately. Likewise, frame (c) shows this broadband PL when viewed through the DBR. However, when exciting at the location of the expected MoS_2 flake, shown in frame (b), there is an extra peak at roughly 600 nm. It is possible this is luminescence from high energy excitations in the MoS_2, or perhaps the flakes are damaged or doped during the

fabrication steps following the initial characterization (Fig. 7.5).

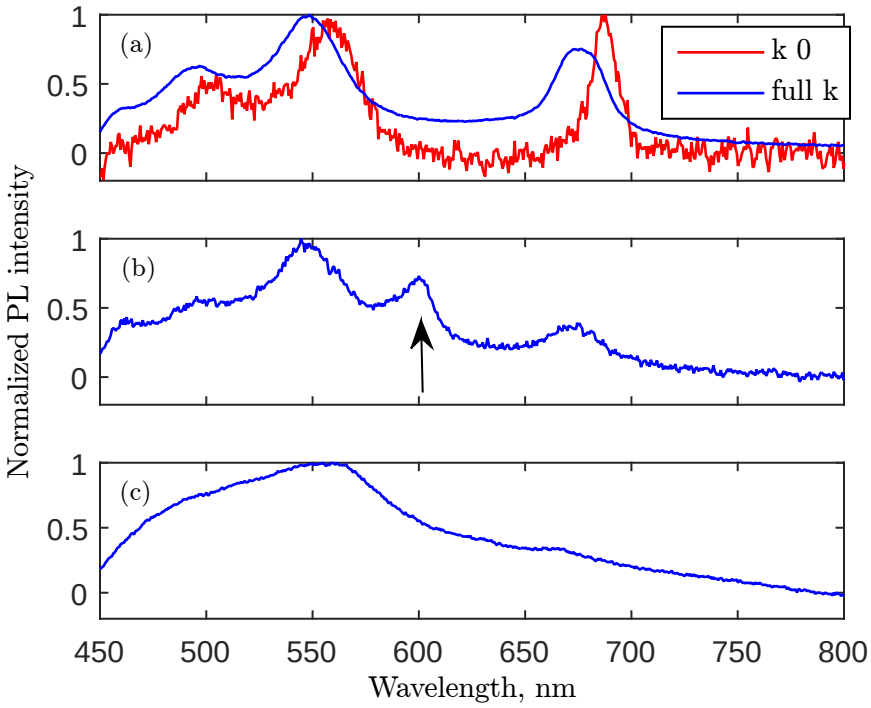

Figure 7.8: Room temperature luminescence from the MoS_2 metal-hybrid cavity while excited by a 405 mn laser. Frame (a): typical PL from the cavity region. Interestingly there appears to be broadband luminescence from the sample, possibly due to some contaminant. The luminescence is consistent with being filtered by the cavity, as the peaks in the PL match with the minima in the reflectivity in Fig 7.7 and are highly angle-dependent. Frame (b): PL suspected to be from the MoS_2 flake. There is extra luminescence around 600 nm (arrow). Frame (c): typical PL from the DBR region.

Currently, we are working on developing fully dielectric cavities to improve the light-matter coupling. Also, it is critical that we properly tune the cavity to the exciton mode to optimize the coupling.

8.0 CONCLUSIONS AND OUTLOOK

The objective of this dissertation was to develop and improve applications of microcavity polaritons. I have demonstrated a polariton optical switch based on the strong nonlinear response of the polariton system [37]. Such an application is of great interest if the polaritons can exist at room temperature, so I worked extensively with collaborating scientists toward designing microcavities that would be in strong coupling at ambient conditions.

When polaritons are strongly coupled at room temperature, a wide range of other applications becomes interesting. The extensive history of studying bosonic condensation effects indicates that rich quantum coherence effects arise in polariton systems. The light mass of polaritons means that coherent effects should be achievable at room temperature, so these new samples promise to bring both the study and application of quantum many body effects out of cryostats and into our hands and pockets. Even if thermodynamic equilibrium is not achieved for a true BEC, for example, polariton lasing[1] promises to give us coherent light generation at lower threshold densities than standard lasing.

By developing samples with greatly enhanced lifetime, we enable technology reliant on the propagation of polaritons over macroscopic distances on-chip [34–36]. For example, a polariton integrated circuit can be implemented, in which some optical transistor nodes are gated by the polaritons created at previous nodes [87].

While the thrust of this research has been to improve the applicability of microcavity polaritons to technology, the advances to polarion lifetimes and operational temperature are also a great boon to basic science research as well. With a longer lifetime, polaritons are better thermalized and form extended condensates qualitatively distinct from previous

[1]The term polariton lasing indicates the spontaneous buildup of coherence, but does not imply thermodynamic equilibrium. Therefore polariton lasing is a less strict condition than BEC.

studies [20, 35]. Unlike in short lifetime samples, we can observe a thermal gas of these polaritons spatially separated from the non-thermal cloud of carriers that create them when using non-resonant excitation. The key here is that the polaritons actually live long enough to propagate far from the exciton cloud, and this has allowed us to study the polariton-polariton interactions without the presence of a dense, more strongly interacting gas of uncoupled carriers [156].

Room temperature samples promise to bring these condensation studies to room temperature. Researchers could study novel quantum phenomena at ambient conditions rather than being forced to consume costly and limited liquid helium. The flexibility of not having to keep samples in a cryostat will also greatly increase the ease of experimental design.

One particular future application of these long-lifetime samples is to study the interactions between macroscopically occupied condensates. David Myers is currently developing methods to fabricate wires out of the 2D polariton structures. He plans to study many phenomena of polaritons propagating macroscopic distances down these wires. In particular, by crossing two wires and launching polariton condensates down them, we hope to see entanglement of the two condensates.

My primary interest at this point is the completion of these room temperature samples. Both the ZnSe and MoS_2 based samples seem close to completion. Once these samples are complete, it would be interesting to demonstrate optical switching and condensation effects at room temperature. As many researchers working with other materials have shown, there is great interest in polariton lasing. If a microcavity can be electrically pumped to generate coherent light, it would be immediately applicable. Basically weak coupling microcavities, vertical cavity surface emitting lasers (VCSELs) are already widespread. If it can be shown that room-temperature coherent emission is achieved at a lower electrical pump threshold in the strong coupling regime than the weak coupling regime, polariton lasers could easily show up in handheld and pocket electronics within a few years.

APPENDIX A

TRANSFER MATRIX FOR PLANAR STACKS

A.1 TRANSFER MATRIX THEORY

As thoroughly discussed in section 2.2, our samples are primarily designed with planar stacks of thin films. These make up the DBR mirrors as well as the cavity spacer. Such layers are on the order of $\lambda/4$ in the DBR layers and up to a few λ in the cavity spacer; thus accounting for typical refractive index values these layers are on the order of 100 nm. Traditionally, there is either a bulk region of excitonic material in the cavity layer, or there can be quantum wells, which are typically on the order of a few nm. A single structure may contain as few as a few layers, or as many as hundreds of layers.

Numerical calculations based on classical electrodynamics are sufficient to predict the wavelength-dependent reflectivity of such structures. Since our structures are planar, a simple plane-wave decomposition of the optical field makes the calculation straightforward, and as we will see, the transfer matrix method is inherently not memory intensive. While the optical properties of the stack rely on the thickness and refractive index of each layer, the entire action of the stack can be described by a 2×2 matrix for each wavelength calculated. In fact, the action of each layer can be represented as a 2×2 matrix, with the full structure thus being the product of these matrices, which is generally done iteratively. Ultimately, we usually compute reflectivity as a function of wavelength, which corresponds again to iterative computations. Thus, these calculations generally run quickly, even on a moderately low powered or older machine.

While this method is covered in many sources, this derivation will follow the presentation by S. L. Chuang [44]. We start by selecting a particular wavelength for which we want the reflection and transmission amplitudes. For simplicity, we will assume that this wave is polarized with the electric field transverse to the plane of incidence (transverse electric, or TE). The most intuitive decomposition of this general wave is

$$E_y = (Ae^{-ik_x x} + Be^{ik_x x})e^{ik_z z}. \tag{A.1}$$

Here E_y is the electric field, and A and B are the complex amplitudes of the forward and backward traveling waves. As already see, k_z will be conserved since it is the component of the wavevector parallel to the surface, but k_x is going to be material dependent.

Here we consider a structure like that presented in Fig. A1. The structure consists of M layers, and each region can potentially have forward and backward propagating waves. Thus, the problem initially consists of $M+2$ A coefficients and also $M+2$ B coefficients. For consistency, we note that these coefficients are defined either at the *start or end* of each layer. The exact choice is immaterial, but we will see that the action of propagation is broken into two parts.

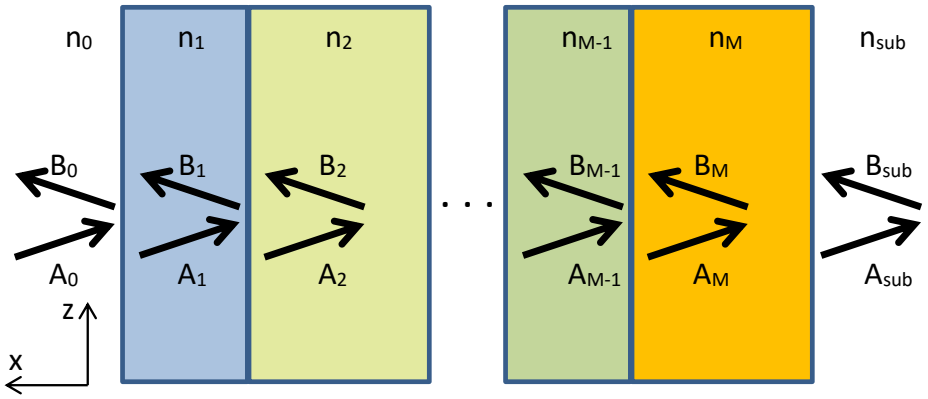

Figure A1: A stack of thin films is illuminated from the $+x$-direction. Each layer will have a forward propagating (A) and backward (B) propagating waves, with A_{sub} and B_{sub} indicating those coefficients in the infinitely thick substrate. Ultimately, we fix $B_{sub} = 0$ since there cannot be a backwards propagating wave in that region.

At each interface, there will be Fresnel style reflection due to the index mismatch at the boundary. Continuity of the tangential electric and magnetic fields allows for the field amplitudes to be computed just across a boundary. We relate the coefficients just across a boundary in region $i+1$ (A'_{i+1}, B'_{i+1}) to those just before the boundary in region i (A_i, B_i):

$$\begin{bmatrix} A_i \\ B_i \end{bmatrix} = \mathbf{T_{i,i+1}} \begin{bmatrix} A'_{i+1}, \\ B'_{i+1} \end{bmatrix} \tag{A.2}$$

$$\mathbf{T_{i,i+1}} = \begin{bmatrix} (1+P_{i,i+1}) & (1-P_{i,i+1}) \\ (1-P_{i,i+1}) & (1+P_{i,i+1}) \end{bmatrix} \tag{A.3}$$

where

$$P_{i,i+1} = \frac{\mu_i k_{i+1,x}}{\mu_{i+1} k_{i,x}} \tag{A.4}$$

where μ is the relative permeability of the material, which is close to 1 in most materials in the optical regime. $k_{i,x}$ is the wavevector in the direction of propagation in the material, which is calculated from the conserved k_z and the fixed frequency of the light. $\mathbf{T_{i,i+1}}$ is technically the matrix which propagates backwards across an interface.

The remainder of the propagation in layer i to relate A'_i to A_i is the straightforward accumulation of phase due to propagation through the medium:

$$\begin{bmatrix} A_i \\ B_i \end{bmatrix} = \mathbf{P_i} \begin{bmatrix} A'_{i'} \\ B'_i \end{bmatrix} \tag{A.5}$$

$$\mathbf{P_i} = \begin{bmatrix} e^{-ik_{i,x}l_i} & 0 \\ 0 & e^{ik_{i,x}l_i} \end{bmatrix} \tag{A.6}$$

Typically, we group the propagation into and through a layer together, since this flows with the logic of adding physical layers to a structure. Thus we define the backwards propagation matrix as

$$\mathbf{B_{i,i+1}} = \mathbf{T_{i,i+1}} \mathbf{P_{i+1}} \tag{A.7}$$

$$\mathbf{B_{0,M+1}} = \mathbf{B_{0,1}} \mathbf{B_{1,2}} \mathbf{B_{2,3}} \cdots \mathbf{B_{M-1,M}} \mathbf{B_{M,M+1}} \tag{A.8}$$

This matrix $\mathbf{B_{0,M+1}}$ directly relates the external fields. Ultimately, we are interested in the reflection and transmission amplitudes, $r = \frac{B_0}{A_0}$ and $t = \frac{A_{M+1}}{A_0}$ respectively. In the linear regime, the magnitude of A_0 is irrelevant; and moreover, since we are interested in the response to light incident from one side, we can set the backward propagating wave in region $M+1$ equal to zero. Physically, this is permissible if the final region is so thick that any light reflected from a later interface will not coherently interfere with the waves under study.

$$\begin{bmatrix} 1 \\ r \end{bmatrix} = \mathbf{B_{0,M+1}} \begin{bmatrix} t \\ 0 \end{bmatrix} \tag{A.9}$$

$$\tag{A.10}$$

Thus the reflection and transmission amplitudes are calculated directly from the elements of $\mathbf{B_{0,M+1}}$:

$$\mathbf{B_{0,M+1}} = \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} \tag{A.11}$$

$$t = \frac{1}{b_{11}} r = \frac{b_{21}}{b_{11}} \tag{A.12}$$

Typically, all we want is the magnitude of the reflectivity (R) or transmissivity (T), which are the magnitude-squared of r and t respectively.

Also, while the above derivation is specific to TE polarization, the same method can be applied to TM polarized light. One can use the duality principle to cyclically permute E, H, μ and ϵ, where ϵ is the dielectric constant in that medium. The only significant change is in the $P_{i,i+1}$ term which embodies the effect of index contrast at the interface. For the TM polarization, we replace it with

$$P^{TM}_{i,i+1} = \frac{\epsilon_{i+1} k_{i+1,x}}{\epsilon_i k_{i,x}}. \tag{A.13}$$

A.2 TRANSFER MATRIX APPLICATION

When fine tuning designs, it is important to account for the refractive index of the constituent materials. Various models or experimental values can be used, but if the user is interested in accuracy over a broad spectral range, then dispersive models must be used. Historically, our group has used a model accounting for temperature, wavelength and mole fraction of AlGaAs alloys [157]. Such a model was useful because we tuned the alloy composition and characterized samples at room temperature and 4 K. For other uses, a good starting point is the website `http://refractiveindex.info/`. This site lists models or tabulated values for some hundreds of materials and cites the source of each. While the user must still be weary of bad models, it is often quite helpful.

In practice one can implement any refractive index model in this calculation, but numerical error can wreak havoc in the cases of thick layers of absorptive material (such as metal). Accurate reflectivity results can be achieved by modeling thin or thick films of metal, for example, but calculations of the optical field profile can become unstable as the film thickness is increased. It is good to be aware of this when observing results with such materials.

Practically speaking, we calculate reflectivity in an iterative way. First, we design a structure with known layer thicknesses and refractive index. In practice, the refractive index of each layer is dispersive and must be computed or looked up for each wavelength. For this reason, we recompute the structure design for each wavelength that we compute R.

In the attached code, this sample design is stored and built in code such as the layer.m code contained in section A.3.1.

This code is called from the actual reflectivity calculator, which is called Reflectivity2015.m and presented in section A.3.2. For each wavelength, this code calls layer.m to generate an array of thicknesses and refractive indices. The calculator then extracts the r and t amplitudes for this wavelength, then the process iterates for the next. For the case where oblique incidence is desired, one must include the angle of incidence and compute both TE and TM polarizations. This is done using ReflectivityTETM2015.m in section A.3.3.

Once a feature of interest (such as a cavity mode or polariton state) has been identified in the reflectivity spectrum of a structure, we can investigate the optical field profile of

that mode directly. As discussed above, r and t are typically calculated from the backward propagation matrix. Once the r coefficient has been determined, then one can *forward* propagate the A_i and B_i coefficients. This is done using ReflectivityTETM2015.m in section A.3.5. Typically we plot the optical mode profile in addition to the refractive index of the structure, as in Fig. 2.7. For example, this is useful in judging if QWs are placed properly at the antinodes of the electric field for optimal coupling.

A.3 EXAMPLE TRANSFER MATRIX CODE

In the following sections, we present a selection of code which can demonstrate the transfer matrix calculation. In practice, we maintain a library of refractive index functions and lookup tables as well as structure functions. This code and description is likely only of interest to those wanting to implement the code.

A.3.1 STRUCTURE DESIGN

The following code is used to 'build' a structure. This is a generic example of a design function, but typically we keep a library of these 'layer.m' files specific to the structures grown and sent to us as well as those in development. This code is either kept in a library of structures added to Matlab's path (for example, grown structures) or in the current directory (for example, when designing and editing the structure).

layer2015_GenericCavity.m

```
1 %this simple design demonstrates the basics of defining sample
2 %design
3
4
5 function l=layer2015_GenericCavity(lambda,DesignParams)
6 %This function generates an array of thicknesses and indices
7 %which represents sample design.
```

```
8
9  %%%Notes that may be helpful:
10 %1)Distances are in angstroms
11 %2)at the end one dummy layer is added and one layer of air is
12 %added to the start of the structure for computing reflectivity
13 %(WHY??? For mode plot? 10/19/2015, Mark: I think it just
14 %vestigial.  Just have to change both Reflectivity and
15 %PlotSingleMode)
16
17 % Often, we may use tabulated or computed refractive indices.
18 % For example, we historically used a model for AlGaAs that
19 % accounted for temperature and alloy mole fraction:
20 % nGaAs = AlGaAs(T,0,lambda);%(pure GaAs)
21 % nAlGaAs = AlGaAs(T,.2,lambda);%(20% aluminum alloy)
22 % Or we may use an oscillator model to account for exciton
23 % effects
24 % nWell = QuantWell(lambda);
25 % Or we could use a fixed, nondispersive index model
26 % nGlass = 1.5;
27
28 %relevant design handles for this design:
29 n1=DesignParams.n1;%index of DBR material 1
30 n2=DesignParams.n2;%index of DBR material 2
31 nc=DesignParams.nc;%index of cavity material
32 qc=DesignParams.dc;%multiplier of cavity thickness
33 qt=DesignParams.dt;%multiplier of overall thickness
34 m=DesignParams.mc;
35 %Cavity Order (ie. m*lambda/2, our samples use m=3)
36 mF=DesignParams.mF;%periods of front DBR
37 mB=DesignParams.mB;%periods of back DBR
```

```
Lam0=DesignParams.lam0;%DBR design wavelength
%Note that any of the above values need not be defined in
%DesignParams.  In that case, they could be static values
   changed
%directly by modifying this code.  However, including them as
%vriables here allows us to vary them in later code.


%simple cavity and DBR design
% Lam0=500;%DBR design wavelength
LamC=Lam0;%Cavity design wavelength
% m=3;%Cavity Order (ie. m*lambda/2, our samples use m=3)
%thicknesses in Angstrom
d1=Lam0*10/(4*n1);
d2=Lam0*10/(4*n2);
dCav=qc*m*LamC*10/(2*nc);


%initialize the structure with an external air layer.  In
%principle this does not have to be air

%l=[thickness in angstrom, index];
k=1;
l=[20000,1];
k=k+1;
%k will track what layer we are on.  We will use matlab's
%on-the-fly memory allocation since that affords great
%flexibility in sample structure

%top DBR is
```

```
l(k,:)=[d1,n1];
k=k+1;
for i=1:mF;
    l((k:k+1),:)=[d2,n2;d1,n1];
    k=k+2;
end

l(k,:)=[dCav,nc];
k=k+1;


for i=1:mB;
    l((k:k+1),:)=[d1,n1;d2,n2];
    k=k+2;
end
l(k,:)=[d1,n1];
k=k+1;
%the substrate can be air or another material.  Be careful! if
%the substrate is not the same material as layer 1, we must
%account for the effect that the index mismatch has on the
%E-field and intensity of light in that material.
l(k,:)=[20000,1];


l(:,1)=qt*l(:,1);
%add buffer layer.  This is important to either the
   Reflectivity
%code or the PlotSingleMode code, I think.
l=[0,0;l];
```

```
96
97 end
```

A.3.2 REFLECTIVITY CALCULATOR

This code calculates the reflectivity of a structure for some vector of wavelengths. Typically the following code is stored in the matlab path and called from some other script, such as ReflectivityScript2015.m.

Reflectivity2015.m

```
1  %This code computes the reflectivity & transmissivity of a
2  %microcavity vs wavelength for selected thickness(d) and
3  %Temperature(T).
4
5  %It should be noted that this code is designed to start with
6  %l(2,:) corresponding to air external to the structure.
7  %
8  %Output~[lambda;r;R;t;Tr]
9
10 %%Version changes: In 2013, I (Mark) added funtionality to call
11 %the desired structure by passing a string of the code file
      name.
12 %This saves the user from having to rename the structure
      function
13 %in the matlab active directory each time a new structure is
14 %computed.
15 %%Version changes: In 2015, I (Mark) removed the scalar
16 %d and T inputs to Reflectivity...() and replaced these with a
17 %structure input
18
```

```
19 function DataOut=Reflectivity2015(lambda,DesignParams,
     StructureName)
20 %This code computes the reflectivity/transmissivity of a
21 %microcavity vs wavelength for selected sample conditions
22 %(DesignParams). This calculator assumes TEM (normal incidence)
23     %the sample design is defined in the function
24     %'StructureName', which is a function that generates an
         array
25     %of indices of refraction and thicknesses as a function of
26     %wavelength.  The structure function may be kept in the
27     %working directory, while the rest of the polariton
28     %calculator is to be added to the Matlab Path.
29     %
30     %In principle, DesignParams can be any set of parameters
31     %relevant to the given structure.  Traditionally, we
         worried
32     %about changing thickness of the entire structure (multiply
33     %all layer thicknesses uniformly by some coeficient) or the
34     %temperature (which was used to compute refractive index of
35     %materials).  However, in some cases, we have no need to
36     %change temperature, but may want to vary different
37     %thicknesses independently.  Therefore, I have chosen to
38     %remove these scalar inputs and accept a variable
39     %(DesignParams) which in general might as well be a
         STRUCTURE
40     %variable type.  Just like the "d" and "T" variables of
         yore,
41     %this variable is just passed along to the layer code.  In
42     %the new format, these old variables would be defined like:
43         %DesignParams.d=...
```

```
44          %DesignParams.T=...
45      %This adds flexibility to the user, and generally cleans up
46      %the input to Reflectivity2015...
47      %The exact format of DesignParams is irrelevant to this
           code.
48      %It only has to match the format required by the specific
49      %layer code specified by 'StructureName'
50  %It should be noted that this code is designed to start with
51  %l(2,:) corresponding to air external to the structure.
52  %
53  %
54  %
55  %Output~[lambda;r;R;t;Tr]
56
57  %The following variable is used in conjunction with some of the
58  %structure functions (particularly some index calculations that
59  %use table lookup) for example, the Silver refractive index is
60  %interpolated based on a dataset of n and k vs lambda.  To
61  %execute that interpolation, I must first load the dataset from
       a
62  %.txt file on the hard drive to memory.  Since the structure
63  %function is called iteratively for each wavelength, this .txt
64  %file must be loaded for each wavelength.  The solution is to
65  %load the data into a persistent variable so that it need only
       be
66  %loaded once for a given call to Reflectivity...()
67  global GlobalWaveStartEnd
68
69
70  %Note input from layer_.m code generally has the form
```

```
%l~[thickenss, index]
%!!caution: for historical resions, this layer.m deines
%thicknesses in Angstroms, while we usualy work in nm.  This is
%converted mid code below!
thet = 0*pi/180;%fix Theta, since this is a TEM calculator
% lambda = lambdamin:step:lambdamax;
Ref(1:5,1:length(lambda))=0;

%iteratively compute R for each lambda
for i=1:length(lambda);
    if i==1
        GlobalWaveStartEnd=1;
    elseif i>1&&i<length(lambda)
        GlobalWaveStartEnd=2;
    elseif i==length(lambda)
        GlobalWaveStartEnd=0;
    end
    Tran=[1,0;0,1];%initialize transfer as identity matrix
    %retrieve structure design at given wavelength and with
    %design specifications.  eval() is the best way that I
       found
    %to execute a function .m-file handle passed as a string.
    %There may be better ways.
    l(:,:)=eval([StructureName '(lambda(i),DesignParams)']);
    for j=3:length(l);
        B=Transfer(l(j-1,2),l(j,2),lambda(i),.1*l(j,1),thet);
        Tran=Tran*B;
    end

    r=(Tran(2,1)/Tran(1,1));
```

```
100     R=abs(r)^2;
101     t=1/Tran(1,1);
102     Tr=abs(t)^2;
103     Ref(:,i)=[lambda(i);r;R;t;Tr];
104 end
105 %Now we note that the electric field prior to the sample (END
        of
106 %l(2,:) is [E0,rE0] and after the sample (start of last layer)
        is
107
108 DataOut=Ref;
109 %plot(Ref(1,:),Ref(3,:));
110
111 end
112
113
114
115 function B = Transfer(n1, n2, lambda, d, thet)
116 %This function returns the propagation matrix from one medium
        to
117 %another. By now d should be in nm. d is the second layer
118 %thickness.
119
120 %See section 5.8 in S.L. Chuang:
121 %[E0;rE0]=[B01]*[B12]*[B23]*...*[Bn(n+1)]*[C(n+1);D(n+1)]
122 %        =[b11,b12;b21,b22]*[tE0,0]
123 %where we have set D(n+1)=0 because there is no reflected wave
        at
124 %the end of the structure.
125
```

```
%thus [c1;d1]=B[c2;d2] yields the coefficients at the end of
%layer 1 given the coeffs at the end of layer 2

%ky is conserved across layers.  kx for each layer is computed
%according to the external wavelength and index of that layer.
    ky = 2*pi./lambda*sin(thet);
    k1x = ((2*pi*n1./lambda).^2-ky.^2).^(0.5);
    k2x = ((2*pi*n2./lambda).^2-ky.^2).^(0.5);
    P = k2x./k1x;%Assumes that both relative permeabilities are
        1
    B(1,1,:) = 0.5*(1 + P).*exp(-1i*k2x*d);
    B(1,2,:) = 0.5*(1 - P).*exp(1i*k2x*d);
    B(2,1,:) = 0.5*(1 - P).*exp(-1i*k2x*d);
    B(2,2,:) = 0.5*(1 + P).*exp(1i*k2x*d);
end
```

A.3.3 OBLIQUE INCIDENCE REFLECTIVITY CALCULATOR

This code extends Reflectivity2015.m to compute reflectivity at oblique angles. Typically the following code is stored in the matlab path and called from some other script, such as ReflectivityScript2015.m.

Practically speaking, this code Reflectivity2015.m and ReflectivityTETM2015.m should be combined into one function, since the first is just a special case of the second. The TETM calculator should be modified to produce *either* TE or TM polarization rather than both, which would recover the computational efficiency of the TEM calculator. However, this is not implemented at this moment because these functions are called by the PlotSingleMode functions, so making these modifications will require further testing.

ReflectivityTETM2015.m

```

```

```
2
3  function DataOut=ReflectivityTETM2015(lambda,DesignParams,theta
       ,StructureName)
4  %See notes in Reflectivity2015
5  %   In addition to that function, this allows for the
       calculation
6  %   of R at non-zero angles of incidence.  These must be
7  %   separately calculated for TE and TM polarizations, which is
8  %   why the output matrix is larger.
9  %Theta should be in radians!!
10
11 %l~[thickenss, index]
12 Ref(1:9,1:length(lambda))=0;
13
14
15
16 for i=1:length(lambda);
17     TranTE=[1,0;0,1];
18     TranTM=[1,0;0,1];
19     l(:,:)=eval([StructureName '(lambda(i),DesignParams)']);
20     for j=3:length(l);
21         %Step for TE polarization
22         B=TransferTE(l(j-1,2),l(j,2),lambda(i),.1*l(j,1),theta)
               ;
23         TranTE=TranTE*B;
24         %Step for TM polarization
25         B=TransferTM(l(j-1,2),l(j,2),lambda(i),.1*l(j,1),theta)
               ;
26         TranTM=TranTM*B;
27     end
```

```
28      %Coeffs for TE polarization
29      rTE=(TranTE(2,1)/TranTE(1,1));
30      RTE=abs(rTE)^2;
31      tTE=1/TranTE(1,1);
32      TrTE=abs(tTE)^2;
33      Ref(1:5,i)=[lambda(i);rTE;RTE;tTE;TrTE];
34
35      %Coeffs for TM polarization
36      rTM=(TranTM(2,1)/TranTM(1,1));
37      RTM=abs(rTM)^2;
38      tTM=1/TranTM(1,1);
39      TrTM=abs(tTM)^2;
40      Ref(6:9,i)=[rTM;RTM;tTM;TrTM];
41  end
42  %Now we note that the electric field prior to the sample (END
        of
43  %l(2,:) is [E0,rE0] and after the sample (start of last layer)
44  %is...[tE0,0]?
45
46  DataOut=Ref;
47  %plot(Ref(1,:),Ref(3,:));
48
49  end
50
51
52
53  function B = TransferTE(n1, n2, lambda, d, thet)
54  %This function returns the propagation matrix from one medium
        to
55  %another. By now d should be in nm. d is the second layer
```

```
%thickness.

%See section 5.8 in S.L. Chuang:
%[E0;rE0]=[B01]*[B12]*[B23]*...*[Bn(n+1)]*[C(n+1);D(n+1)]
%        =[b11,b12;b21,b22]*[tE0,0]
%where we have set D(n+1)=0 because there is no reflected wave at
%the end of the structure.

%thus [c1;d1]=B[c2;d2] yields the coefficients at the end of
%layer 1 given the coeffs at the end of layer 2

%ky is conserved across layers.  kx for each layer is computed
%according to the external wavelength and index of that layer.
    ky = 2*pi./lambda*sin(thet);
    k1x = ((2*pi*n1./lambda).^2-ky.^2).^(0.5);
    k2x = ((2*pi*n2./lambda).^2-ky.^2).^(0.5);
    P = k2x./k1x;%Assumes that both relative permeabilities are 1
    B(1,1,:) = 0.5*(1 + P).*exp(-1i*k2x*d);
    B(1,2,:) = 0.5*(1 - P).*exp(1i*k2x*d);
    B(2,1,:) = 0.5*(1 - P).*exp(-1i*k2x*d);
    B(2,2,:) = 0.5*(1 + P).*exp(1i*k2x*d);
end

function B = TransferTM(n1, n2, lambda, d, thet)
%This function returns the propagation matrix from one medium to
%another. By now d should be in nm. d is the second layer
%thickness.
```

```
83
84 %See section 5.8 in S.L. Chuang:
85 %[E0;rE0]=[B01]*[B12]*[B23]*...*[Bn(n+1)]*[C(n+1);D(n+1)]
86 %        =[b11,b12;b21,b22]*[tE0,0]
87 %where we have set D(n+1)=0 because there is no reflected wave
      at
88 %the end of the structure.
89
90 %thus [c1;d1]=B[c2;d2] yields the coefficients at the end of
91 %layer 1 given the coeffs at the end of layer 2
92
93 %ky is conserved across layers.  kx for each layer is computed
94 %according to the external wavelength and index of that layer.
95     ky = 2*pi./lambda*sin(thet);
96     k1x = ((2*pi*n1./lambda).^2-ky.^2).^(0.5);
97     k2x = ((2*pi*n2./lambda).^2-ky.^2).^(0.5);
98     P = (n1^2*k2x)./(n2^2*k1x);
99     %P Assumes that both relative permeabilities are 1
100    B(1,1,:) = 0.5*(1 + P).*exp(-1i*k2x*d);
101    B(1,2,:) = 0.5*(1 - P).*exp(1i*k2x*d);
102    B(2,1,:) = 0.5*(1 - P).*exp(-1i*k2x*d);
103    B(2,2,:) = 0.5*(1 + P).*exp(1i*k2x*d);
104 end
```

A.3.4 REFLECTIVITY CALCULATION SCRIPT

The following script is very simple code demonstrating how to implement the above calculator.

ReflectivityScript2015.m

```
%use this script to calculate and plot the reflectivity of a
%structure

lam=[300:420 421:.2:520 521:700];

%chosen structure
test='layer2015_GenericCavity';
%relevant design handles:
DesignParams.n1=2.2;%index of DBR material 1
DesignParams.n2=1.5;%index of DBR material 2
DesignParams.nc=1.4;%index of cavity material
DesignParams.dc=1;%multiplier of cavity thickness
DesignParams.dt=1;%multiplier of overall thickness
DesignParams.mc=3;
%Cavity Order (ie. m*lambda/2, our samples use m=3)
DesignParams.mF=5;%periods of front DBR
DesignParams.mB=5;%periods of back DBR
DesignParams.lam0=500;%DBR design wavelength

R=Reflectivity2015(lam,DesignParams,test);

% use the following code to calculate reflectivity for non-
   normal
% incidence light
theta=degtorad(40);
R2=ReflectivityTETM2015(lam,DesignParams,theta,test);

figure(1)
% plot(lam,R(3,:),lam,R2(3,:),lam,R2(7,:))
hold all
```

```
% plot(lam,R(3,:))
plot(lam,R(3,:),lam,R2(3,:),lam,R2(7,:))
```

A.3.5 MODE PROFILE CALCULATOR

Typically, the following code is used for a cavity mode or polariton wavelength once such a feature has been identified with the reflectivity calculator, although it can in principle be used at any wavelength. We input a single wavelength and the sample design, and get out the index and E-field vs depth into the structure. Typically we plot the magnitude (or magnitude-squared) of this field, since a local QW will feel the time-averaged field.

We note that this calculator can also be extended to oblique incidence. However we do not include that code here. The currently implemented code is somewhat ambiguous as to what the relevant field are. It is the case that the electric and magnetic fields are out of phase in these structures. Thus, when we plot the TM calculated mode, we need to take care to think of the electric field that is implied for that mode. While these details are readily solvable, they have not been critical enough to work out explicitly.

PlotSingleMode2015.m

```


%Plan: have this function output the amplitude of the electric
%field vs x for a given design (layer2 function) after
   computing
%reflectivity coeffs.
%
%1)use Reflectivity.m to compute reflectivity vs lambda.  Look
   at
%interesting values of lambda
%
%2)Now use this function to plot a wavelength of interest.
```

```
11      %a) the initial coeffs will be determined by making a call
            to
12      %reflectivity.m for the one wavelength in question (This
13      %occurs in photonAB function).  The coefficients are
            computed
14      %for END of each layer and ammended as new columns to the
15      %layer array (LayerCoeffs)
16      %b)This array is passed back to the main function where
            each
17      %layer is expanded (with spacial points determined by the
18      %wavelength and index of refraction such that the mode can
            be
19      %plotted well).  The electric field is propogated backwards
20      %from the end of each interface from the computed A B
            coeffs
21  %
22  %3) The output contains the index of refraction and electric
23  %field vs x. Note that x is a nonuniform measure of depth into
24  %the sample from the front face.  (nonuniform because spacing
        was
25  %chosen by the wavelength of light in each medium.) the output
26  %looks like [x;n;E]
27
28  %%Version changes: In 2013, I (Mark) added funtionality to call
29  %the desired structure by passing a string of the code file
        name.
30  %This saves the user from having to rename the structure
        function
31  %in the matlab active directory each time a new structure is
32  %computed.
```

```
33 %%Version changes: In 2015, I (Mark) removed the scalar
34 %d and T inputs to Reflectivity...() and replaced these with a
35 %structure input
36
37
38 function [Out,layer]=PlotSingleMode2015(lambda,DesignParams,
      StructureName)
39 %PlotSingleMode2013 outputs an array of [distance;index;E-field
      ].
40
41
42 %Immediately, the sample design is ammended with the A and B
43 %coefficients of the field within each region:
44 layer=photonAB(lambda,DesignParams,StructureName);
45 m=size(layer);
46
47 clear Out
48 for i=2:m(1)
49     di=layer(i,1);%layer thicknesses
50     ni=layer(i,2);%layer indexes
51     cforward=layer(i,3);
52     cbackward=layer(i,4);
53 %     if i==2
54 %         ampin=abs(cforward+cbackward);%obsolete?
55 %     end
56
57     %Here we decide the spacing within the current layer.
          Since
58     %we are going to be plotting sinusiodal fields, it is good
          to
```

```
59      %have enough points to smooth the curve.  However, the
            actual
60      %wavelength is going to be material (index) dependent.  We
61      %don't want to use a fixed value, because some layers can
            be
62      %only a few nm (<< lambda), while in principle, some can be
63      %tens or hundreds of lambda
64      NumPoints=max(20,round(100*.1*di*ni/lambda));
65      xi=linspace(0,di,NumPoints);
66      nivect=ones(size(xi))*ni;
67      ki=2*pi*ni/lambda;
68
69      if i==2
70          x=xi-layer(i,1);
71      else
72          x=xi+max(x);
73      end
74
75          phi=ki*(layer(i,1)-xi)*.1;
76
77  %Because the coeffs are computed for the *back* of the layer,
        we
78  %must propogate backwards through the thickness. (thus
79  %phi~(di-xi))
80      Ei=(cforward)*exp(-1i*phi)+(cbackward)*exp(1i*phi);
81      if i==2
82          Out=[x;nivect;Ei];
83      else
84          Out=[Out,[x;nivect;Ei]];
85      end
```

```
86  end
87
88
89  end
90
91
92
93
94
95  function LayerCoeffs=photonAB(lambda,DesignParams,StructureName
        )
96  %this will propagate the photon mode through each layer of the
97  %sample.  It will ammend the layer matrix to have the A and B
98  %coefficients just at the end of each interface.
99
100 % I figured it woud be better to calculate all the A's and B's
101 % using the full forward transfer matrix for each layer, rather
102 % than propagate across boundary, then interface, and keep
103 % passing the complex amplitudes back and forth.  Perhaps it
104 % doesn't make much of a difference.
105 LayerCoeffs=eval([StructureName '(lambda,DesignParams)']);
106 l=LayerCoeffs;
107 Coeffs(1:length(l),1:2)=0;
108 r=Reflectivity2015(lambda,DesignParams,StructureName);
109 C=[1;r(2)]; %These are the mode coeffs at the surface
110
111 %enter the sample:
112 Coeffs(2,1)=C(1);
113 Coeffs(2,2)=C(2);
114
```

```
for j=3:length(l);
    Trans=ForwardTransfer(l(j-1,2),l(j,2),lambda,.1*l(j,1));
    C=Trans*C;
    Coeffs(j,1)=C(1);
    Coeffs(j,2)=C(2);

end
LayerCoeffs(:,3:4)=Coeffs;
end


%This function returns the propagation matrix from one medium
   to
%another
function B = ForwardTransfer(n1, n2, lambda, d)
%This propogates across the interface, then medium 2 Thus Bi
   will
%take the coefficients from just before interface 1->2 to just
%before interface 2->3.

    %ky = 2*pi./lambda*sin(thet);
    %k1x = ((2*pi*n1./lambda).^2-ky.^2).^(0.5);
    %k2x = ((2*pi*n2./lambda).^2-ky.^2).^(0.5);
    k2=2*pi*n2/lambda;
    P = n1/n2;
    B(1,1,:) = 0.5*(1 + P).*exp(1i*k2*d);
    B(1,2,:) = 0.5*(1 - P).*exp(1i*k2*d);
    B(2,1,:) = 0.5*(1 - P).*exp(-1i*k2*d);
    B(2,2,:) = 0.5*(1 + P).*exp(-1i*k2*d);
end
```

A.3.6 MODE PROFILE CALCULATOR SCRIPT

The following script is very simple code demonstrating how to implement the above mode profile calculator.

ModePlotScript2015.m

```
lamP=500;%wavelength mode to plot


%chosen structure
test='layer2015_GenericCavity';
%relevant design handles:
DesignParams.n1=2.2;%index of DBR material 1
DesignParams.n2=1.5;%index of DBR material 2
DesignParams.nc=1.4;%index of cavity material
DesignParams.dc=1;%multiplier of cavity thickness
DesignParams.dt=1;%multiplier of overall thickness
DesignParams.mc=3;
%Cavity Order (ie. m*lambda/2, our samples use m=3)
DesignParams.mF=5;%periods of front DBR
DesignParams.mB=5;%periods of back DBR
DesignParams.lam0=500;%DBR design wavelength

[A,B]=PlotSingleMode2015(lamP,DesignParams,test);


%A~[x;index;E] B~data about the structure including
%forward/backward coeffs in each layer
x=A(1,:);
```

```
24 n=A(2,:);
25 E=A(3,:);
26 En=E/max(abs(E));
27 En2=abs(En).^2;
28 %This is mod(E)^2.  Not properly intensity, since we would have
29 %to add the magnetic field also
30
31 plot(x,n,x,En2+1)
```

APPENDIX B

RAY TRACING FOR PERFECT LENSES

Here I will discuss the method and code used to trace rays

As discussed in chapter 3, ray tracing can be used to design or better understand a complex optical setup. This can be particularly useful to assess field of view and numerical aperture throughput of a design. Commercial products such as OSLO, Code V or ZEMAX can be especially useful for estimating efficiencies and exploring aberrations. However, full versions of these packages come with high costs, and while they may offer a plethora of measurables as output, the learning curve and setup time can be a major trade off.

The benefit of these commercial packages is that they typically track rays across each interface. For a simple lens, a ray will refract at the first surface, propagate through the lens, and then refract at the second surface. Often the programs have the physical structure of many commercial lenses preloaded. Since the optical materials can have dispersive optical density, the code will properly account for chromatic aberrations. Since the actual curvatures of lens surfaces are used, spherical and other aberrations can also be calculated.

However, for most of our applications, these aberrations are only a secondary concern. Often chromatic aberrations are irrelevant since we typically observe only a narrow spectral range. We are more interested in image (real and angle resolved) formation and magnification, matching numerical aperture to detectors, and avoiding clipping of high NA and large field-of-view rays. The most basic understanding of these effects come directly from geometric optics with an assumption of perfect lenses.

To extend the usefulness of geometric optics beyond the thin lens equation $\frac{1}{f} = \frac{1}{d_0} + \frac{1}{d_i}$,

we develop a technique to trace rays through lenses. Considering parameters defined in Fig. B1, one can work out a relationship between the initial and final rays. We know that the final ray will be bent according to its initial angle and the location it strikes the lens. Assessing the action on a few principle rays quickly reveals that the action can be expressed as

$$\tan(\theta_f) = \tan(\theta_i) - \frac{h}{f}. \tag{B.1}$$

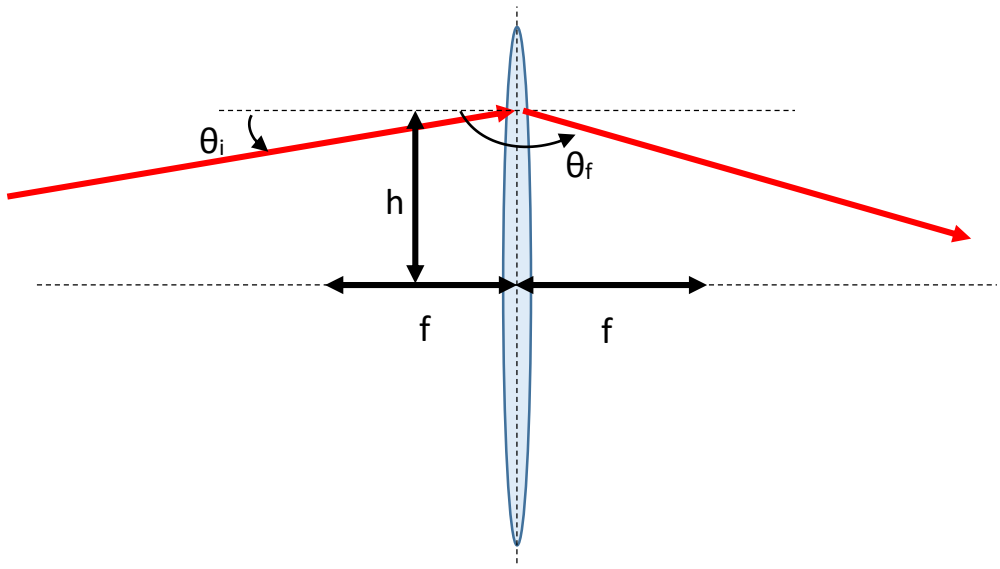

Figure B1: The action of a lens is to change the angle of a ray of light depending on where it strikes the lens.

Practically, propagating a ray of light through a lens system now amounts to finding the intersection of this ray with the lens to calculate h, finding the slope of the ray in the next region, then finding the intersection with the next lens. By keeping track of the apertures of lenses, mirrors and other optics in the system, one can plot the ray and visually inspect it for clipping. Locating image planes amounts to bookkeeping: different NA rays from a single point should be checked for intersections in each region. Locating angle resolved images is the corollary: all the rays of a single NA but emitted from different spatial points are checked

for intersections. The primary benefit of ray tracing on this scale is developing intuition for image formation, particularly angle resolved imaging.

Additionally, since the lens action is so strightforward, we can easily add offsets and tilts to the lens system. An offset of a lens therefore just means that we must calculate h as the distance from the lens center. Tilting the lens effectively changes the angle of rays relative to the lens. Still, we simply look for intersection of the ray with the lens just as before. This is useful to develop intuition for the effect of misaligned optics and the aberrations that can result. Since we regularly translate lenses to image different parts of the field of view, it can be helpful to simulate those shifts to be sure that using lenses off-axis isn't causing aberrations or extra clipping.

B.1 RAY TRACING CODE

The current version of our ray tracing code is functional, but perhaps not user friendly. Most of the calculation is done in a Matlab script such as RayTracer.m. The first section of this code defines the optical setup to be studied. Lens focal lengths, positions, radius, offsets and tilts are entered manually. Special care must be taken to ensure that lenses are entered iteratively along the optical axis: the code will not correct for a lens accidentally placed *prior* to other lenses, but the rays will be traced back to that lens and refracted. Non-lens objects, such as mirrors, beam splitters and irises, can be included in this code just by setting the focal length to a very large value (i.e. orders of magnitude larger than distances in the system). For example, this may be desirable when a beam splitter is inserted in the beam path at an angle, which results in a reduced aperture and increased clipping. In principle these objects should be included and plotted separately, but the computational cost is minimal in applications of this calculator.

The second section of the code selects the initial rays to study. This means that we select spatial points and numerical aperture rays to track. Really any combination of multiple NA rays and multiple spatial points are enough to locate images, but it is best for the user to specify spatial points and NA rays on the scale that they will actually try to observe. Due to

the current structure of the code, these rays must be selected with the same number of NA rays as spatial points. Typically 3 or 5 of each are used. For example, we may use spatial points $y = -1$, 0 and 1 mm and NA $= -0.4$, 0 and 0.4. This gives us a total of 9 rays which are traced through the system. Also in this section are stylistic choices for the traced rays. The code plots different spatial points in different colors, and different NA rays in different styles (eg. dashed or dotted). These color and style selections can be changed here.

The third section is the proper ray tracing code. In addition to storing each ray segment, this section calls the function PerfectLens.m, which takes the ray and lens properties, then outputs the new ray. This code also calls the function LineIntersection.m, which simply determines the point of intersection of the ray with the lens. Back in the main script, this new ray is plotted *if* there is another lens. For this reason, it is generally useful to add a final object beyond the desired image location. Iteratively, each ray is traced through the lens system in this way.

The final two sections search for all the segment intersections to form real space and angle resolved images. Where intersections are found markers are paced on the ray diagram. This code specifically uses the SegmentIntersection.m function to only select intersections within the current region.

RayTracer.m

```
%problems:
%+fitting image surfaces breaks down when number of spatial
%points does not equal number of angle points.
%+need to add convenient readout of image points
%+be nice to have good way to 'autofocus' a lens--least squares
%vary position of a lens to put image plane in proper place.

%propagate light through a few perfect lenses:

%use format of:
%lens=[f,x,yoffset,tilt,radius]
%where f is focal length
```

```
%x is positional along optical axis
%yoffset is translation of lens perpendicualr to optical axis
%tilt is roation of lens relative to optical axis
%radius is that of the lens, for clipping purposes

%My setup looks something like this where ( ) indicates a lens
%and <- -> indicates a distance:
%Object<-10->(f=10)<-210->(f=100)<-250->(f=75)<-150->Image
%which would be entered as thus:
% L=[[10,10,0,0,4.65];...
%     [100,10+210,0,0,12];...
%     [75,10+210+250,0,0,25];...
%     [1000,10+210+250+100,0,0,4]];%this final 'lens' is the
   image plane


%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%% user: enter optics setup

%if desired, shift the origin.  I think?
Z=0;

X1=10;%M.O.principal plane
X2=210;X1_2=X1+X2;%position of imaging lens
X3=100;X1_3=X1_2+X3;%position of real space iris
X3K=200;X1_3K=X1_2+X3K;%position of real space iris
X4=150;X1_4=X1_3+X4;%position of real space imaging lens
X4K=200;X1_4K=X1_3+X4K;%position of k-space imaging lens
X5=300;X1_5=X1_3+X5;%position of final image

```

```
%In this example, the user should slect EITHER the "Real Space
%lens" OR the "2f-2f k-Space lens" OR "k-Space lens dual A/B"
L=[[  10,   X1, 0, 0,4.65];December 9th 2015...%microscope objective
   [ 100, X1_2, 0, 0,  25];...%imaging lens
   [  75, X1_4, 0, 0,  25];...%Real Space lens
%    [  50,X1_4K, 0, 0,  25];...%2f-2f k-Space lens
%    [  50,X1_3K+50, 0, 0,  25];...%k-Space lens dual A
%    [  50,X1_5-50, 0, 0,  25];...%k-Space lens dual B
%    [1E10, X1_5, 0, 0,   4];...%image sensor/slit
   [1E10, X1_5+100,0,0,1]];%extend plot past slit for user
      benefit

%%%%%%%%%%
%%%%%%%%%%
%%%%%%%%%%

%% Initial rays to model:

%initial rays=[xo,yo,theta0,red,green,blue]


%For now, keep length P and length A the same (5 + 5 is a
   little
%busy, but OK

%P is list of spatial points to model rays from
%use 1, 3 or 5 spatial poitns.  Also use the same number of
%angles (in variable A below).  For some reason the code
   crashes
%otherwise...haven't bothered to fix it
```

```
P=[0 -1 1]/1;
% P=[0 -.5 .5 -1 1]/1;

%A is the list of angles to model rays from

%The NA of our objective is 0.42, so I usually base A off of
   that
NA=[0 -1 1]*.4/1;A=asind(NA);
% NA=[0 -1/3 1/3 -2/3 2/3 -1 1]*.42;A=asind(NA);

%C is the color list.  I have each spatial point a different
   color.
C=[0 0 1;1,0,0;0,.6,0;0,.8,.8;.6,.6,0];
% s={'-','--',':','-.'};
% s={'-','--','--','-.','-.',':',':'};%line style for angle
s={'-','--','-.','-.',':',':'};%line style for angle
s2={'o','^','v','s','p'};%marker style for angle
qq=1;
R=zeros(length(P)*length(A),8);
S=cell(length(P)*length(A),1);


%% Ray Tracing:

for j=1:length(P);
    for i=1:length(A)
        R(qq,:)=[Z P(j) A(i) C(j,:) j i];
        S{qq}=s{i};
        qq=qq+1;
    end
```

```
end

qL=size(L);%->get number of lenses
qR=size(R);%->get number of rays
figure
    %the following variable will be used to archive the
    %slopes/intercepts of all rays
    SlopeIntercept=zeros(2,qL(1),length(P),length(A));
    %The alternate way to record rays is to use the end points:
    XY=zeros(2,qL(1)+1,length(P),length(A));


for i=1:qR(1)
    xri=R(i,1);
    yri=R(i,2);
    thetai=R(i,3);
    color=R(i,4:6);
    %the following merely records the starting point/angle
    %iteration for later organization of results
    Pi=R(i,7);
    Ai=R(i,8);

    style=S{i};
    X=zeros(1,qL(1)+1);
    Y=X;
    X(1)=xri;
    Y(1)=yri;

    for j=1:qL(1)
        f=L(j,1);
```

```
127         xl=L(j,2);
128         yl=L(j,3);
129         tilt=L(j,4);
130         rl=L(j,5);
131         if i==1
132 
133             line([xl-rl*sind(tilt),xl+rl*sind(tilt)],...
134                 [yl+rl*cosd(tilt),yl-rl*cosd(tilt)],...
135                 'Color','k');hold all
136         end
137 
138 %           if i==1
139         %hmm...color should be removed from PerfectLens
140         [yri,xri,thetai]=PerfectLens(xri,yri,thetai,...
141             f,xl,yl,tilt,color);
142         X(j+1)=xri;
143         Y(j+1)=yri;
144         %%%To determine intersections, compute slope and
145         %%%intercept for each section and save it just like we
146         %%%are saving X and Y.  Be sure to save for all
              segments
147         m=(yri-Y(j))/(xri-X(j));
148         b=yri-m*xri;
149 
150 
151         %now record X,Y,m and b before moving on to next ray:
152         XY(:,:,Pi,Ai)=[X;Y];
153         SlopeIntercept(:,j,Pi,Ai)=[m;b];
154 %           pause(2)
155     end
```

```
156
157     line(X,Y,'Color',color,'LineStyle',style);hold all
158 %     line(X,Y,'Color',color,'LineStyle',':');hold all
159 %     pause(2)
160 end
161
162
163 %% Find real image surfaces
164
165 %now use the SlopeIntercept archive to generate positive
       crossings
166 %Real images:
167     RealImages=cell(4,qL(1));
168     xmean=NaN(1,length(P));
169     xstd=xmean;
170     ymean=xmean;
171     ystd=xmean;
172 for k=1:qL(1)
173     %angle crossings:
174         Acrossings=length(A)*(length(A)-1)/2;
175     tempfull=NaN(2,Acrossings*length(P));
176     if k==1
177         xi=min(R(:,1));
178     else
179         xi=L(k-1,2);
180     end
181     xf=L(k,2);
182     for j=1:length(P)%find crossing for each spatial point
183         %by intersection of rays@different angles for that
               point
```

```
184         temp=NaN(2,Acrossings);
185         %temp will record all crossings for a single spatial
186         %point
187         qq=1;
188         for i=1:(length(A)-1)
189             m1=SlopeIntercept(1,k,j,i);
190             b1=SlopeIntercept(2,k,j,i);
191             for ii=(i+1):length(A)
192                 m2=SlopeIntercept(1,k,j,ii);
193                 b2=SlopeIntercept(2,k,j,ii);
194 %                 [m1 m2 b1 b2]
195 
196                 [xc,yc]=SegmentIntersection(m1,b1,m2,b2,xi,xf);
197                 temp(:,qq)=[xc;yc];
198                 qq=qq+1;
199             end
200         end
201         %add values in temp (single spatial point) to tempfull
202         %(all spatial points for this region of x)
203         tempfull(:,((j-1)*Acrossings)+(1:Acrossings))=temp;
204 
205         xmean(j)=mean(temp(1,:));
206         xstd(j)=std(temp(1,:));
207         ymean(j)=mean(temp(2,:));
208         ystd(j)=std(temp(2,:));
209 
210 
211         scatter(temp(1,:),temp(2,:),'d','filled',...
212             'SizeData',50,'MarkerEdgeColor','k',...
213             'MarkerFaceColor',C(j,:))
```

```
    end
    %assign tempfull to the cell in the RealImages cell array
    %that corresponds to this region of x
    RealImages{1,k}=tempfull;

    if ~isnan(sum(xmean))&&~isnan(sum(ymean))
        %fit a polynomial to the image surface
        [coeffs]=polyfit(ymean,xmean,2);
        ylim=1.5*max(abs(ymean));
        yspace=linspace(-ylim,ylim,20);
        plot(polyval(coeffs,yspace),yspace,':k')
    else
        coeffs=NaN(1,3);
    end

    RealImages{2,k}=[xmean;ymean];
    RealImages{3,k}=[xstd;ystd];
    RealImages{4,k}=coeffs;
end

%% Find k-image surfaces

%now use the SlopeIntercept archive to generate positive
%crossings k images:
    KImages=cell(4,qL(1));
    xmean=zeros(1,length(A));
    xstd=xmean;
    ymean=xmean;
    ystd=xmean;
for k=1:qL(1)
```

```
    %angle crossings:
        Pcrossings=length(P)*(length(P)-1)/2;
    tempfull=NaN(2,Pcrossings*length(A));
    if k==1
        xi=0;
    else
        xi=L(k-1,2);
    end
    xf=L(k,2);
    for i=1:length(A)%find crossing for each angle of emission
        %by intersection of rays from different points for that
        %angle
        temp=NaN(2,Pcrossings);
        %temp will record all crossings for a single emission
        %angle
        qq=1;
        for j=1:(length(P)-1)
            m1=SlopeIntercept(1,k,j,i);
            b1=SlopeIntercept(2,k,j,i);
            for jj=(j+1):length(P)
                m2=SlopeIntercept(1,k,jj,i);
                b2=SlopeIntercept(2,k,jj,i);
%                 [m1 m2 b1 b2]

                [xc,yc]=SegmentIntersection(m1,b1,m2,b2,xi,xf);
                temp(:,qq)=[xc;yc];
                qq=qq+1;
            end
        end
        %add values in temp (single spatial point) to tempfull
```

```
274         %(all spatial points for this region of x)
275         tempfull(:,((i-1)*Pcrossings)+(1:Pcrossings))=temp;
276 
277         xmean(i)=mean(temp(1,:));
278         xstd(i)=std(temp(1,:));
279         ymean(i)=mean(temp(2,:));
280         ystd(i)=std(temp(2,:));
281 
282 
283         scatter(temp(1,:),temp(2,:),'^','filled',...
284             'SizeData',40,'MarkerEdgeColor','k',...
285             'MarkerFaceColor','r')
286     end
287     %assign tempfull to the cell in the KImages cell array that
288     %corresponds to this region of x
289     KImages{1,k}=tempfull;
290 
291     if ~isnan(sum(xmean))&&~isnan(sum(ymean))
292         %fit a polynomial to the image surface
293         [coeffs]=polyfit(ymean,xmean,2);
294         ylim=1.5*max(abs(ymean));
295         yspace=linspace(-ylim,ylim,20);
296         plot(polyval(coeffs,yspace),yspace,':k')
297     else
298         coeffs=NaN(1,3);
299     end
300 
301     KImages{2,k}=[xmean;ymean];
302     KImages{3,k}=[xstd;ystd];
303     KImages{4,k}=coeffs;
```

```
end

%% report results
% Ideas for the future: take real space image data from k=0
   rays.
% take k-space image data from y=0 rays.

%% Figure formatting
%Optional figure formatting:
% ylabel('$x$, Distance from optical axis, mm','Interpreter','
   latex')
% xlabel('$z$, Distance along optical axis, mm','Interpreter','
   latex')
% set(gcf,'Units','inches','Position',[.1,.5,5,4],...
%     'PaperSize',[5,4])
% set(gca,'XLim',[0, 700],'YLim',[-28,28])
```

PerfectLens.m

```
%a perfect lens should transform a ray at
%[yi,thetai]->[yf,thetaf] where yf=yi and
%tan(thetaf)=tan(thetai)-h/f where h=yi-ylens and f is the
   focal
%length

function [yr,xr,thetaf]=PerfectLens(xro,yro,thetai,f,xl,yl,tilt
   ,raycolor)
%this function will take a ray from the previous surface,
   propagate it
%through air until it strikes the lens, then refract the ray.

```

```
[xr,yr]=LineIntersection(xro,yro,thetai,xl,yl,90+tilt);
%Adjust for offset of the lens, if there is any
h=(yr-yl)/cosd(tilt);
%the relevant angle of the light is with respect to the lens, not
%the optical axis
thetaf=atand(tand(thetai)-h/f);
end
```

LineIntersection.m

```
function [xc,yc] = LineIntersection( x1,y1,theta1,x2,y2,theta2 )
%LINEINTERSECTION find point of intersection between two lines
%   x/y are points on lines 1 and 2. Theta is the inclination of
%   each line with respect to the x-axis.  Note that simple
%   algorithms will yield a problem if either line is vertical.
%   Another potential problem is if the two lines are parallel

%generally, m is slope, b is intercept, and xc and yc are the
%critical points.
if theta1==theta2
    %there is no intersection if the two lines are parallel
    %(otherwise, they overlap)
    xc=[];yc=[];
elseif theta1==90
    %if one ray is vertical, we must treat the case specially
    m2=tand(theta2);
    b2=y2-m2*x2;
    xc=x1;yc=m2*x1+b2;
```

```
elseif theta2==90
    m1=tand(theta1);
    b1=y1-m1*x1;
    xc=x2;yc=m1*x2+b1;
else
    m1=tand(theta1);
    b1=y1-m1*x1;
    m2=tand(theta2);
    b2=y2-m2*x2;
    xc=-(b1-b2)/(m1-m2);
    yc=b1+m1*xc;
end
end
```

SegmentIntersection.m

```
function [xc,yc]=SegmentIntersection(m1,b1,m2,b2,xi,xf)
%Find intersection between two non-vertical line segments in
   the
%region of xi<=x<=xf
%   define a simple intersection function to be used later.  At
%   this point we shouldn't need to worry about vertical lines,
%   but parallel lines are possible.  Also, we must filter out
%   intersections that occur outside of the line segments:
%Find intersection
if m1==m2
    xc=NaN;yc=NaN;
else
    xc=-(b2-b1)/(m2-m1);
end
%Filter for only intersections within current region
```

```
15 if xc>=xi && xc<=xf
16     yc=b1+m1*xc;
17 else
18     xc=NaN;yc=NaN;
19 end
20 end
```

APPENDIX C

GAUSSIAN-EXPONENTIAL DECAY CONVOLUTION AND FITTING

The form of the time-resolved polariton PL can be understood best as the result of hot excitations relaxing into the polariton states. The rise time indicates a multiple-path relaxation from the hot excitations to the LP state, so the complicated dynamics become difficult to model. Since we cannot measure the intermediate or high energy populations, the uncertainty in the parameters governing the relaxation becomes very large. Because of this, we use simple functions to parametrize the data.

A convolution of a Gaussian with an exponential decay was chosen as a natural function to fit the observed time-resolved PL data with a minimum number of fit parameters. As shown in Fig. 4.3, the data clearly exhibit a long decay time which suggests fitting the data with an exponential decay, and the rise time fits a Gaussian broadening reasonably well; the broadening can be understood as due to the multiple paths for polariton generation from the initial incoherent hot carriers created by the pump laser. The central time of the Gaussian peak gives a convenient parameter to measure the arrival time of the polariton population. Including the overall intensity of the data and background, this means that each curve is fit with 5 parameters.

The Gaussian-exponential convolution (GEC) is calculated according to

$$n(t) = \int_0^\infty \left(\left[\frac{A}{\sigma\sqrt{2\pi}} e^{\left(\frac{-(t-x-t_0)^2}{2\sigma^2}\right)} \right] \left[\frac{1}{\tau} e^{\left(-x/\tau\right)} \right] \right) dx. \tag{C.1}$$

The five parameters of the model are σ, the broadening of the Gaussian; t_0, the peak time of the unconvolved Gaussian; τ, the exponential lifetime; A, the time-integrated intensity,

and ultimately a possible background. Performing the convolution leads to the form

$$n(t) = \frac{A}{2\tau} e^{\left(\frac{\sigma^2 - 2t\tau + 2t_0\tau}{2\tau^2}\right)} \operatorname{erfc}\left(\frac{\sigma^2 + t_0\tau - t\tau}{\sqrt{2}\sigma\tau}\right) \tag{C.2}$$

where $\operatorname{erfc}(t)$ is the complimentary error function.

Since the GEC model is not derived from a theoretical basis of the relaxation of excitations to the LP states, it is dangerous to interpret too much from the parameters of the fit. For example, the decay time τ is not simply the lifetime of the LP population; it includes the effect of the mean lifetime of the reservoir particles to scatter into the LP state. If the excited states, that is, hot free carriers and excitons, take a long time to relax but have no other means to decay quickly, then it is possible to measure a long lifetime for this decay parameter even if the final polariton decay process is fast [158]. However, we note that the rise time to populate the polariton states is on the order of 80 ps, which is not substantially longer than the decay time itself, and the range of decay times measured from these fits are on the same order as the other lifetime estimates, so these values are still in agreement with our assessment that the polaritons themselves have a lifetime on the order of 100-200 ps.

While several parameters of this fit do not directly give information about the polariton dynamics, the t_0 parameter is useful and indicative of the time of arrival of the polaritons at the location being observed. Other methods of assigning this time, such as the peak time, the 10% and 50% turn on times were investigated as well. While all of these data clearly have different offsets, the overall trends fall within their respective uncertainties. These assessments were included in assigning the uncertainty of the time-of-flight data, for example in Figs. 4.4 and 4.5.

APPENDIX D

TRANSPORT LIFETIME MEASUREMENT SUPPLEMENTAL NOTES

D.1 SAMPLE DETAILS

This sample is the same as used in Refs [34,35]. A $3\lambda/2$ microcavity contains three sets of four GaAs quantum wells located at the antinodes of the cavity mode. The QWs are nominally 70 Å pure GaAs embedded in pure AlAs barriers nominally 30 Å thick. The optical mode is confined between distributed Bragg reflectors made of AlAs/$Al_{0.2}Ga_{.8}As$ with 32 pairs on the top surface and 40 pairs on the bottom surface. Molecular beam epitaxial growth of the sample leads to an inherent wedge to the cavity thickness, resulting in a gradient of the cavity as well as exciton energies. At 5 K, the polariton exhibits a Rabi coupling of 6 meV, which corresponds to an upper polariton-lower polariton splitting of 12 meV; the cavity mode gradient is 13 meV/mm and the exciton gradient is 1.5 meV/mm.

D.2 METHODS

The sample was held in a cold-finger cryostat at 5 K for all experiments.

Emission was collected using a N.A.=0.42 microscope objective. A preliminary imaging lens permitted spatial filtering of the real space image data, and a subsequent iris in the Fourier image plane permitted filtering of the emission angle. Secondary lenses could be exchanged to image either the real-space or angle-resolved emission. Luminescence was

imaged through a spectrometer onto either a standard CCD or onto a Hamamatsu streak camera.

Polaritons were resonantly injected at $\lambda = 778$ nm with a picosecond laser far on the photonic side of the sample. The injected polaritons were created at a detuning of approximately -2 meV corresponding to the high k value of $k_{\|} = 5.5 \times 10^4$ cm^{-1}. This state was created by an external pump angle of roughly 42°. The $k_{\|} = 0$ detuning at this position is -30 meV.

The sample and pump laser were arranged such that the polaritons were moving antiparallel to the cavity gradient, which was aligned with the time slit. The angle of incidence was larger than the collection angle of the optics, so the reflected beam was not collected. Additionally, the pump spot was spatially outside the field of view such that scattered light was not collected. At a distance of approximately 2 mm from the injection point, emission entered the collection range of the optics. At the turn-around point, the 778 nm polaritons are more photonic with a detuning of -7 meV which corresponds to a photon fraction of 75%.

Figure D1 demonstrates the polariton dispersion relation at three points along its trajectory in the data of Fig. 3(a). Dispersions in Fig. D1(a), (b) and (c) correspond respectively to $x = 0, 1.75$ and 2.25 mm from the injection point. Thus, frame (a) indicates the polariton dispersion at the point of injection, frame (b) corresponds to the point at which the polaritons enter the imaging field of view, and frame (c) shows the dispersion just at the turn around point. Black dashed lines indicate the bare photon dispersion at each point, while the black dashed-dotted lines indicate the bare exciton dispersion. Frame (d) overlays the three polariton dispersions for easier comparison. The injection momentum condition is shown by the red circle at ≈ 5 μm^{-1}. The cavity gradient in the $-x$-direction leads to an acceleration to negative momentum, as indicated by the red arrow. The polaritons propagate at constant energy. The black brackets indicate the numerical aperture of our collection optics, so we can only observe polaritons with wavevector $\|k\| < 3 \times 10^4$ cm^{-1}.

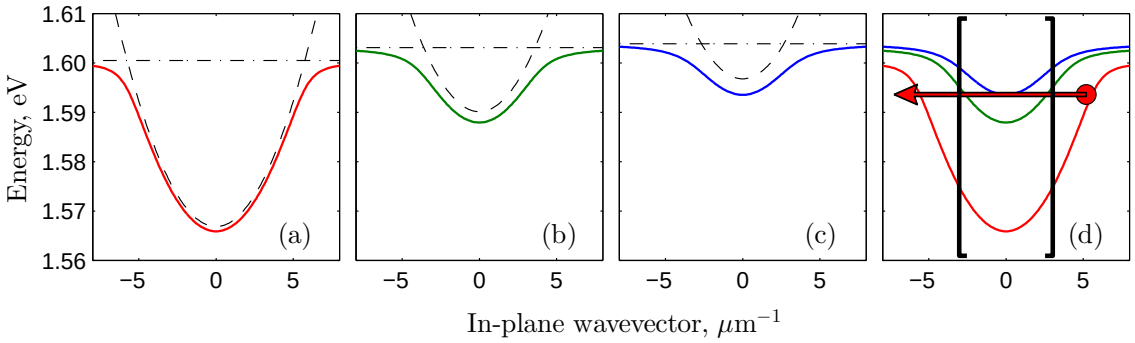

Figure D1: Calculated polariton dispersion relations for three spatial points relevant to the data presented in Fig. 5.3(a). Frames (a-c) show the dispersions at the injection point, where the polaritons enter the field of view, and the turn around point respectively. Black dashed and dashed-dotted lines show the bare photon and exciton dispersions respectively. The polariton dispersions are overlayed in (d) and the injection conditions are shown by a red circle. The collection angle of the optics is shown by the black brackets. The red arrow indicates the trajectory of the polaritons in momentum space as they are accelerated by the cavity wedge.

D.3 DISTANCE CALIBRATION

The x-distance from the injection point for the 40° injection cases was estimated as follows: the x-distance in Figs. 5.2, and 5.3(a), was determined from the fit of the polariton x vs t trajectory presented in Fig. 5.3(a). Extrapolation of this fit back to time $t = 0$ (as determined by locating scattered laser light) determines the initial position of excitation. This initial excitation position is consistent with the sample parameters and injection conditions. This method assumes that the acceleration of the polaritons is strictly constant from creation to turn around. Variation in the acceleration due to a non-constant energy gradient in addition to the changing mass of the polariton implies uncertainty on the overall offset of this axis, but the spatial magnification was measured directly.

APPENDIX E

ZNSE MBE SAMPLE DESIGNS

The following designs were planned for MBE growth at NREL to be subsequently etched and made into microcavities. The thickness of GaAs is just a placeholder for a (typically thick) buffer layer used in MBE growth to smooth the substrate in preparation for the subsequent growth. The QW design is based on the sample 26_0063 discussed in Section 7.2.1.

These ZnSe-based MBE growths are grown on GaAs substrates because, due to a fortunate coincidence, ZnSe has the same lattice structure and almost exactly the same lattice constant as GaAs. Since GaAs technology is well developed, substrates made from GaAs are relatively inexpensive and high quality. Clearly, the mismatch of the materials at the interface will give rise to interface states and potentially disorder at that boundary. However, the bulk of the structure should be relatively crystalline and reasonably defect free.

The following designs assume the concentration of Cd in the $Cd_xZn_{1-x}Se$ is $x = .28$, just as sample 26_0063. The 'rotate' column indicates whether the sample is to be rotated or not. Due to the anisotropy of the MBE growth, if the sample is not rotated, the final sample will have a thickness gradient. This is desirable in the spacer layers for us, since it will impart a gradient to the cavity energy, giving us a range of detunings possible on a single sample.

The first sample design (Table E1) uses a single QW at each antinode of a 2λ cavity. The second design (Table E2) uses three QWs at each antinode, potentially increasing the Rabi splitting by a factor of $\sqrt{3}$. While increasing the number of QWs is ideal for enhancing the coupling, it may also increase the disorder in the exciton mode.

Table E1: Single QW/antinode cavity

Layer	Thickness (nm)	Material	Notes	Rotate
0	1000.00	GaAs	Substrate	yes
1	5.00	AlAs		
2	20.00	GaAs		
3	95.62	ZnSe	Spacer	no
4	5.00	QW	Antinode	yes
5	93.12	ZnSe	Spacer	no
6	5.00	QW	Antinode	yes
7	93.12	ZnSe	Spacer	no
8	5.00	QW	Antinode	yes
9	95.62	ZnSe	Spacer	no

Table E2: Multiple QW/antinode cavity

Layer	Thickness (nm)	Material	Notes	Rotate
0	1000.00	GaAs	Substrate	yes
1	5.00	AlAs	Substrate	yes
2	20.00	GaAs	Substrate	yes
3	80.62	ZnSe	Spacer	no
4	5.00	QW	Antinode	yes
5	10.00	ZnSe		
6	5.00	QW		
7	10.00	ZnSe		
8	5.00	QW		
9	63.12	ZnSe	Spacer	no
10	5.00	QW	Antinode	yes
11	10.00	ZnSe		
12	5.00	QW		
13	10.00	ZnSe		
14	5.00	QW		
15	63.12	ZnSe	Spacer	no
16	5.00	QW	Antinode	yes
17	10.00	ZnSe		
18	5.00	QW		
19	10.00	ZnSe		
20	5.00	QW		
21	80.62	ZnSe	Spacer	no

BIBLIOGRAPHY

[1] J. Kasprzak, M. Richard, S. Kundermann, A. Baas, P. Jeambrun, J. M. J. Keeling, F. M. Marchetti, M. H. Szymaska, R. André, J. L. Staehli, V. Savona, P. B. Littlewood, B. Deveaud, and L. S. Dang, "BoseEinstein condensation of exciton polaritons," *Nature*, vol. 443, pp. 409–14, sep 2006.

[2] R. Balili, V. Hartwell, D. Snoke, L. Pfeiffer, and K. West, "Bose-Einstein condensation of microcavity polaritons in a trap," *Science*, vol. 316, pp. 1007–1010, may 2007.

[3] B. Nelsen, R. Balili, D. W. Snoke, L. Pfeiffer, and K. West, "Lasing and polariton condensation: Two distinct transitions in GaAs microcavities with stress traps," *Journal of Applied Physics*, vol. 105, no. 12, p. 122414, 2009.

[4] H. Deng, G. Weihs, D. Snoke, J. Bloch, and Y. Yamamoto, "Polariton lasing vs. photon lasing in a semiconductor microcavity," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 100, no. 26, pp. 15318–15323, 2003.

[5] D. Bajoni, P. Senellart, E. Wertz, I. Sagnes, A. Miard, A. Lemaître, and J. Bloch, "Polariton Laser Using Single Micropillar GaAs-GaAlAs Semiconductor Cavities," *Physical Review Letters*, vol. 100, p. 047401, jan 2008.

[6] P. Bhattacharya, B. Xiao, A. Das, S. Bhowmick, and J. Heo, "Solid state electrically injected exciton-polariton laser," *Physical Review Letters*, vol. 110, no. 20, pp. 1–5, 2013.

[7] D. Snoke, "Polariton Condensation and Lasing," in *Exciton Polaritons in Microcavities: New Frontiers* (V. Timofeev and D. Sanvitto, eds.), Springer Series in Solid State Sciences, ch. 12, pp. 307–327, Springer Berlin Heidelberg, 2012.

[8] S. Christopoulos, G. B. H. von Högersthal, A. J. D. Grundy, P. G. Lagoudakis, A. V. Kavokin, J. J. Baumberg, G. Christmann, R. Butté, E. Feltin, J.-F. Carlin, and N. Grandjean, "Room-Temperature Polariton Lasing in Semiconductor Microcavities," *Physical Review Letters*, vol. 98, p. 126405, mar 2007.

[9] K. S. Daskalakis, P. S. Eldridge, G. Christmann, E. Trichas, R. Murray, E. Iliopoulos, E. Monroy, N. T. Pelekanos, J. J. Baumberg, and P. G. Savvidis, "All-dielectric

GaN microcavity: Strong coupling and lasing at room temperature," *Applied Physics Letters*, vol. 102, no. 10, pp. 2013–2016, 2013.

[10] S. Kéna-Cohen and S. Forrest, "Room-temperature polariton lasing in an organic single-crystal microcavity," *Nature Photonics*, vol. 4, pp. 371–375, 2010.

[11] T. Klein, S. Klembt, E. Durupt, C. Kruse, D. Hommel, and M. Richard, "Polariton lasing in high-quality Selenide-based micropillars in the strong coupling regime," *ArXiv*, pp. 1–5, 2015.

[12] A. V. Yulin, O. A. Egorov, F. Lederer, and D. V. Skryabin, "Dark polariton solitons in semiconductor microcavities," *Physical Review A*, vol. 78, p. 061801, dec 2008.

[13] O. Egorov, D. Skryabin, A. Yulin, and F. Lederer, "Bright Cavity Polariton Solitons," *Physical Review Letters*, vol. 102, p. 153904, apr 2009.

[14] A. Amo, S. Pigeon, D. Sanvitto, V. G. Sala, R. Hivet, I. Carusotto, F. Pisanello, G. Lemenager, R. Houdre, E. Giacobino, C. Ciuti, and A. Bramati, "Polariton Superfluids Reveal Quantum Hydrodynamic Solitons," *Science*, vol. 332, pp. 1167–1170, jun 2011.

[15] M. Sich, D. N. Krizhanovskii, M. S. Skolnick, A. V. Gorbach, R. Hartley, D. V. Skryabin, E. A. Cerda-Méndez, K. Biermann, R. Hey, and P. V. Santos, "Observation of bright polariton solitons in a semiconductor microcavity," *Nature Photonics*, vol. 6, pp. 50–55, nov 2011.

[16] G. Grosso, G. Nardin, F. Morier-Genoud, Y. Léger, and B. Deveaud-Plédran, "Dynamics of dark-soliton formation in a polariton quantum fluid," *Physical Review B*, vol. 86, p. 020509, jul 2012.

[17] K. G. Lagoudakis, M. Wouters, M. Richard, A. Baas, I. Carusotto, R. André, L. S. Dang, and B. Deveaud-Plédran, "Quantized vortices in an excitonpolariton condensate," *Nature Physics*, vol. 4, pp. 706–710, aug 2008.

[18] K. G. Lagoudakis, T. Ostatnický, A. V. Kavokin, Y. G. Rubo, R. André, and B. Deveaud-Plédran, "Observation of half-quantum vortices in an exciton-polariton condensate.," *Science*, vol. 326, no. 5955, pp. 974–976, 2009.

[19] K. G. Lagoudakis, F. Manni, B. Pietka, M. Wouters, T. C. H. Liew, V. Savona, A. V. Kavokin, R. André, and B. Deveaud-Plédran, "Probing the Dynamics of Spontaneous Quantum Vortices in Polariton Superfluids," *Physical Review Letters*, vol. 106, p. 115301, mar 2011.

[20] G. Liu, D. W. Snoke, A. Daley, L. N. Pfeiffer, and K. West, "A new type of half-quantum circulation in a macroscopic polariton spinor ring condensate," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 112, pp. 2676–81, mar 2015.

[21] K. G. Lagoudakis, B. Pietka, M. Wouters, R. André, and B. Deveaud-Plédran, "Coherent oscillations in an exciton-polariton Josephson junction," *Physical Review Letters*, vol. 105, no. 12, pp. 1–4, 2010.

[22] H. Deng, G. Weihs, C. Santori, J. Bloch, and Y. Yamamoto, "Condensation of semiconductor microcavity exciton polaritons," *Science*, vol. 298, no. 5591, pp. 199–202, 2002.

[23] M. Richard, J. Kasprzak, R. Andre, R. Romestain, L. S. Dang, G. Malpuech, and A. Kavokin, "Experimental evidence for nonequilibrium Bose condensation of exciton polaritons," *Physical Review B*, vol. 72, no. 20, pp. 201301–1–4, 2005.

[24] H. Deng, *Dynamic Condensation of Semiconductor Microcavity Polaritons*. Ph.d., Stanford University, 2006.

[25] O. L. Berman, Y. E. Lozovik, and D. W. Snoke, "Theory of Bose-Einstein condensation and superfluidity of two-dimensional polaritons in an in-plane harmonic potential," *Physical Review B*, vol. 77, p. 155317, apr 2008.

[26] R. Balili, B. Nelsen, D. Snoke, L. Pfeiffer, and K. West, "Role of the stress trap in the polariton quasiequilibrium condensation in GaAs microcavities," *Physical Review B*, vol. 79, p. 075319, feb 2009.

[27] J. D. Plumhof, T. Stöferle, L. Mai, U. Scherf, and R. F. Mahrt, "Room-temperature Bose-Einstein condensation of cavity exciton-polaritons in a polymer," *Nature Materials*, vol. 13, no. 3, pp. 247–52, 2014.

[28] M. Nakayama, K. Murakami, Y. Furukawa, and D. Kim, "Polariton-condensation effects on photoluminescence properties in a CuBr microcavity," *Journal of Physics: Conference Series*, vol. 619, p. 012015, 2015.

[29] D. Snoke and P. Littlewood, "Polariton condensates," *Physics Today*, vol. 63, no. 8, pp. 42–47, 2010.

[30] H. Deng, H. Haug, and Y. Yamamoto, "Exciton-polariton Bose-Einstein condensation," *Reviews of Modern Physics*, vol. 82, pp. 1489–1537, may 2010.

[31] J. Kasprzak, *Condensation of exciton polaritons*. Ph.d. thesis, Joseph Fourier University-Grenoble, 2006.

[32] R. B. Balili, *Bose-Einstein condensation of microcavity polaritons*. Ph.d. thesis, University of Pittsburgh, 2009.

[33] B. Nelsen, *Polariton condensates in a trap and photon lasing in two-dimensional semiconductor microcavities*. PhD thesis, University of Pittsburgh, 2012.

[34] M. Steger, G. Liu, B. Nelsen, C. Gautham, D. W. Snoke, R. Balili, L. Pfeiffer, and K. West, "Long-range ballistic motion and coherent flow of long-lifetime polaritons," *Physical Review B*, vol. 88, p. 235314, dec 2013.

[35] B. Nelsen, G. Liu, M. Steger, D. W. Snoke, R. Balili, K. West, and L. Pfeiffer, "Dissipationless Flow and Sharp Threshold of a Polariton Condensate with Long Lifetime," *Physical Review X*, vol. 3, p. 041015, nov 2013.

[36] M. Steger, C. Gautham, D. Snoke, L. Pfeiffer, and K. West, "Slow reflection and two-photon generation of microcavity excitonpolaritons," *Optica*, vol. 2, no. 1, pp. 1–5, 2015.

[37] M. Steger, C. Gautham, B. Nelsen, D. Snoke, L. Pfeiffer, and K. West, "Single-wavelength, all-optical switching based on exciton-polaritons," *Applied Physics Letters*, vol. 101, no. 13, p. 131104, 2012.

[38] D. W. Snoke, *Solid State Physics; Essential Concepts*. Addison-Wesley, 2009.

[39] N. W. Ashcroft and N. D. Mermin, *Solid State Physics*. Thomson Learning, 1976.

[40] C. Kittel, *Introduction to Solid State Physics*. John Wiley and Sons, Inc., 8th ed., 2005.

[41] J. J. Sakurai, *Modern Quantum Mechanics*. Addison-Wesley, revised ed., 2011.

[42] D. W. Snoke, W. W. Rühle, K. Köhler, and K. Ploog, "Spin flip of excitons in GaAs quantum wells," *Physical Review B*, vol. 55, no. 20, pp. 13789–13794, 1997.

[43] C. Gautham, D. W. Snoke, A. Rastelli, and O. G. Schmidt, "Time-resolved two-photon excitation of dark states in quantum dots," *Applied Physics Letters*, vol. 104, no. 14, p. 143114, 2014.

[44] S. L. Chuang, *Physics of Photonic Devices*. Hoboken, NJ: Wiley, 2nd ed., 2009.

[45] C. G. Van De Walle and R. M. Martin, "Theoretical calculations of heterojunction discontinuities," *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, vol. 4, p. 1055, 1986.

[46] C. G. Van de Walle, K. Shahzad, and D. J. Olego, "Strained-layer interfaces between II–VI compound semiconductors," *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, vol. 6, no. 4, p. 1350, 1988.

[47] C. G. Van De Walle, "Band Lineups and Deformation Potentials in the Model-Solid Theory," *Physical Review B*, vol. 39, no. 3, pp. 1871–1883, 1989.

[48] G. Bastard, E. E. Mendez, L. L. Chang, and L. Esaki, "Exciton binding energy in quantum wells," *Physical Review B*, vol. 26, no. 4, pp. 1974–1979, 1982.

[49] H. Mathieu, P. Lefebvre, and P. Christol, "Excitons in semiconductor quantum wells : A straightforward analytical calculation," *Journal of Applied Physics*, vol. 72, no. 1, pp. 300–302, 1992.

[50] C. Gautham, M. Steger, D. Snoke, K. West, and L. Pfeiffer, "Time-resolved two-photon excitation of long-lifetime polaritons," *Optica-Under Review*, 2015.

[51] J. K. Wuenschell, *TRAPPING AND TRANSPORT OF INDIRECT EXCITONS IN COUPLED QUANTUM WELLS*. Ph.d. dissertation, University of Pittsburgh, 2014.

[52] R. Balili, B. Nelsen, D. W. Snoke, R. H. Reid, L. Pfeiffer, and K. West, "Huge splitting of polariton states in microcavities under stress," *Physical Review B - Condensed Matter and Materials Physics*, vol. 81, p. 125311, mar 2010.

[53] N. W. Sinclair, J. K. Wuenschell, Z. Vörös, B. Nelsen, D. W. Snoke, M. H. Szymanska, a. Chin, J. Keeling, L. N. Pfeiffer, and K. W. West, "Strain-induced darkening of trapped excitons in coupled quantum wells at low temperature," *Physical Review B*, vol. 83, p. 245304, jun 2011.

[54] A. P. Arya, *Introduction to Classical Mechanics*. Upper Saddle River, NJ: Prentice-Hall, Inc, 2nd ed., 1998.

[55] M. E. Purcell, "Spontaneous Emission Probabilities at Radio Frequencies," *Physical Review*, vol. 69, no. June, p. 681, 1946.

[56] J. J. Hopfield, "Theory of the Contribution of Excitons to the Complex Dielectric Constant of Crystals," *Physical Review*, vol. 112, no. 5, pp. 1555–67, 1958.

[57] P. Savvidis, C. Ciuti, J. Baumberg, D. Whittaker, M. Skolnick, and J. Roberts, "Off-branch polaritons and multiple scattering in semiconductor microcavities," *Physical Review B*, vol. 64, p. 075311, jul 2001.

[58] J. Keeling, P. Eastham, M. Szymanska, and P. Littlewood, "BCS-BEC crossover in a system of microcavity polaritons," *Physical Review B*, vol. 72, p. 115320, sep 2005.

[59] J. Keeling, F. M. Marchetti, M. H. Szymanska, and P. B. Littlewood, "Collective coherence in planar semiconductor microcavities," *Semiconductor Science and Technology*, vol. 22, no. 5, p. 46, 2007.

[60] F. M. Marchetti, M. H. Szymaska, J. M. J. Keeling, J. Kasprzak, R. André, P. B. Littlewood, and L. Si Dang, "Phase diagram for condensation of microcavity polaritons: From theory to practice," *Physical Review B*, vol. 77, no. 23, p. 235313, 2008.

[61] A. Kavokin and G. Malpuech, *Cavity Polaritons*. Thin Films and Nanostructures, San Diego, Ca: Elsevier Inc, 2003.

[62] D. M. Whittaker, P. Kinsler, T. A. Fisher, M. S. Skolnick, A. Armitage, A. M. Afshar, M. D. Sturge, and J. S. Roberts, "Motional narrowing in semiconductor microcavities," *Physical Review Letters*, vol. 77, no. 23, pp. 4792–4795, 1996.

[63] D. Whittaker, "What Determines Inhomogeneous Line Widths in Semiconductor Microcavities?," *Physical Review Letters*, vol. 80, no. 21, pp. 4791–4794, 1998.

[64] A. V. Kavokin, J. J. Baumberg, G. Malpuech, and F. P. Laussy, *Microcavities*. Series on Semiconductor Science and Technology, New York: Oxford University Press, 2007.

[65] V. E. Hartwell, *Dynamics of trapped polaritons in stressed GaAs quantum well-microcavity structures: experiments and numerical simulations*. Ph.d. thesis, University of Pittsburgh, 2008.

[66] V. E. Hartwell and D. W. Snoke, "Numerical simulations of the polariton kinetic energy distribution in GaAs quantum-well microcavity structures," *Physical Review B*, vol. 82, p. 075307, aug 2010.

[67] G. Tosi, G. Christmann, N. G. Berloff, P. Tsotsis, T. Gao, Z. Hatzopoulos, P. G. Savvidis, and J. J. Baumberg, "Sculpting oscillators with light within a nonlinear quantum fluid," *Nature Physics*, vol. 8, pp. 190–194, jan 2012.

[68] G. Christmann, G. Tosi, N. G. Berloff, P. Tsotsis, P. S. Eldridge, Z. Hatzopoulos, P. G. Savvidis, and J. J. Baumberg, "Polariton ring condensates and sunflower ripples in an expanding quantum liquid," *Physical Review B*, vol. 85, p. 235303, jun 2012.

[69] C. Anton, T. Liew, G. Tosi, M. Martin, T. Gao, Z. Hatzopoulos, P. Eldridge, P. Savvidis, and L. Viña, "Dynamics of a polariton condensate transistor switch," *Applied Physics Letters*, vol. 101, no. 26, pp. 261116–1–5, 2012.

[70] P. Cristofolini, A. Dreismann, G. Christmann, G. Franchetti, N. G. Berloff, P. Tsotsis, Z. Hatzopoulos, P. G. Savvidis, and J. J. Baumberg, "Optical Superfluid Phase Transitions and Trapping of Polariton Condensates," *Physical Review Letters*, vol. 110, p. 186403, may 2013.

[71] A. Hayat, C. Lange, L. A. Rozema, R. Chang, S. Potnis, H. M. van Driel, A. M. Steinberg, M. Steger, D. W. Snoke, L. N. Pfeiffer, and K. W. West, "Macroscopic coherence between quantum condensates formed at different times," *Optics Express*, vol. 22, no. 25, pp. 1–24, 2014.

[72] R. Stevenson, V. Astratov, M. Skolnick, D. Whittaker, M. Emam-Ismail, A. Tartakovskii, P. Savvidis, J. Baumberg, and J. Roberts, "Continuous wave observation of massive polariton redistribution by stimulated scattering in semiconductor microcavities," *Physical Review Letters*, vol. 85, pp. 3680–3, oct 2000.

[73] M. Saba, C. Ciuti, J. Bloch, V. Thierry-Mieg, R. André, L. S. Dang, S. Kundermann, a. Mura, G. Bongiovanni, J. L. Staehli, and B. Deveaud, "High-temperature ultrafast

polariton parametric amplification in semiconductor microcavities.," *Nature*, vol. 414, no. 6865, pp. 731–735, 2001.

[74] S. Kundermann, M. Saba, C. Ciuti, T. Guillet, U. Oesterle, J. L. Staehli, and B. Deveaud, "Coherent control of polariton parametric scattering in semiconductor microcavities.," *Physical Review Letters*, vol. 91, no. 10, p. 107402, 2003.

[75] W. Langbein, "Polariton correlation in microcavities produced by parametric scattering," *Physica Status Solidi (B)*, vol. 242, pp. 2260–2270, sep 2005.

[76] C. Palmer, *Diffraction Grating Handbook*. Rochester, NY: Newport Corporation, sixth ed. ed., 2005.

[77] E. Hecht, *Optics*. San Francisco: Pearson Addison Wesley, 4th ed. ed., 2002.

[78] H. Goldstein, C. Poole, and J. Safko, *Classical Mechanics*. San Francisco: Addison-Wesley, third ed., 2002.

[79] P. Wen, Y. Sun, K. A. Nelson, B. Nelsen, G. Liu, M. Steger, D. W. Snoke, L. N. Pfeiffer, and K. West, "Phase diagram of Bose condensation of long-lifetime polaritons in equilibrium," *In Preparation*, 2013.

[80] M. Holzmann, G. Baym, J.-P. Blaizot, and F. Laloë, "Superfluid transition of homogeneous and trapped two-dimensional Bose gases.," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 104, pp. 1476–81, jan 2007.

[81] M. Davis, S. Gardiner, T. Hanna, N. Nygaard, N. Proukakis, and M. Szymanska, "Fluctuations and Degeneracy in Low Dimensions," in *Cold Atoms Vol. 1: Quantum Gases: Finite Temperature and Non-Equilibrium Dynamics* (N. Proukakis, S. Gardiner, M. Davis, and M. Szymanska, eds.), pp. 72–76, London: Imperial College Press, 2013.

[82] M. Wouters, I. Carusotto, and C. Ciuti, "Spatial and spectral shape of inhomogeneous nonequilibrium exciton-polariton condensates," *Physical Review B*, vol. 77, p. 115340, mar 2008.

[83] M. Richard, J. Kasprzak, R. Romestain, R. André, and L. S. Dang, "Spontaneous Coherent Phase Transition of Polaritons in CdTe Microcavities," *Physical Review Letters*, vol. 94, p. 187401, may 2005.

[84] E. Wertz, L. Ferrier, D. Solnyshkov, R. Johne, D. Sanvitto, A. Lemaître, I. Sagnes, R. Grousson, A. V. Kavokin, P. Senellart, G. Malpuech, and J. Bloch, "Spontaneous formation and optical manipulation of extended polariton condensates," *Nature Physics*, vol. 6, pp. 860–864, aug 2010.

[85] C. Weisbuch, M. Nishioka, A. Ishikawa, and Y. Arakawa, "Observation of the coupled exciton-photon mode splitting in a semiconductor quantum microcavity," *Physical Review Letters*, vol. 69, no. 23, pp. 3314–3317, 1992.

[86] F. Tassone and Y. Yamamoto, "Lasing and squeezing of composite bosons in a semiconductor microcavity," *Physical Review A*, vol. 62, p. 063809, nov 2000.

[87] D. Ballarini, M. De Giorgi, E. Cancellieri, R. Houdré, E. Giacobino, R. Cingolani, A. Bramati, G. Gigli, and D. Sanvitto, "All-optical polariton transistor," *Nature Communications*, vol. 4, p. 1778, jan 2013.

[88] J. E. Heebner and R. W. Boyd, "Enhanced all-optical switching by use of a nonlinear fiber ring resonator.," *Optics Letters*, vol. 24, no. 12, pp. 847–849, 1999.

[89] T. Pertsch, U. Peschel, and F. Lederer, "All-optical switching in quadratically nonlinear waveguide arrays," *Optics Letters*, vol. 28, no. 2, pp. 102–104, 2003.

[90] V. R. Almeida, C. A. Barrios, R. R. Panepucci, M. Lipson, M. A. Foster, D. G. Ouzounov, and A. L. Gaeta, "All-optical switching on a silicon chip," *Optics Letters*, vol. 29, no. 24, pp. 2867–2869, 2004.

[91] G. Ctistis, E. Yuce, A. Hartsuiker, J. Claudon, M. Bazin, J.-M. Gerard, and W. L. Vos, "Ultimate fast optical switching of a planar microcavity in the telecom wavelength range," *Applied Physics Letters*, vol. 98, no. 16, p. 161114, 2011.

[92] A. M. C. Dawes, L. Illing, S. M. Clark, and D. J. Gauthier, "All-optical switching in rubidium vapor," *Science (New York, N.Y.)*, vol. 308, no. 5722, pp. 672–674, 2005.

[93] X. Wei, J. Zhang, and Y. Zhu, "All-optical switching in a coupled cavity-atom system," *Physical Review A*, vol. 82, p. 033808, sep 2010.

[94] J. Sheng, X. Yang, U. Khadka, and M. Xiao, "All-optical switching in an N-type foul-level atom-cavity system," *Optics Express*, vol. 19, no. 18, pp. 17059–64, 2011.

[95] R. Houdré, J. Gibernon, and P. Pellandini, "Saturation of the strong-coupling regime in a semiconductor microcavity: Free-carrier bleaching of cavity polaritons," *Physical Review B*, vol. 52, no. 11, pp. 7810–7813, 1995.

[96] D. Snoke, "Polariton condensates: a feature rather than a bug," *Nature Physics*, vol. 4, no. September, pp. 674–675, 2008.

[97] T. Gao, P. Eldridge, T. Liew, S. Tsintzos, G. Stavrinidis, G. Deligeorgis, Z. Hatzopoulos, and P. Savvidis, "Polariton condensate transistor switch," *Physical Review B*, vol. 85, p. 235102, jun 2012.

[98] A. Amo, T. C. H. Liew, C. Adrados, R. Houdré, E. Giacobino, A. V. Kavokin, and A. Bramati, "Excitonpolariton spin switches," *Nature Photonics*, vol. 4, pp. 361–366, apr 2010.

[99] T. C. H. Liew, A. V. Kavokin, T. Ostatnický, M. Kaliteevski, I. A. Shelykh, and R. A. Abram, "Exciton-polariton integrated circuits," *Physical Review B*, vol. 82, p. 033302, jul 2010.

[100] C. Symonds, A. Lematre, E. Homeyer, J. C. Plenet, and J. Bellessa, "Emission of Tamm plasmon/exciton polaritons," *Applied Physics Letters*, vol. 95, no. 15, p. 151114, 2009.

[101] M. Kaliteevski, S. Brand, R. A. Abram, I. Iorsh, A. V. Kavokin, and I. A. Shelykh, "Hybrid states of Tamm plasmons and exciton polaritons," *Applied Physics Letters*, vol. 95, no. 25, p. 251108, 2009.

[102] M. Wouters, T. K. Paraïso, Y. Léger, R. Cerna, F. Morier-Genoud, M. T. Portella-Oberli, and B. Deveaud-Plédran, "Influence of a nonradiative reservoir on polariton spin multistability," *Physical Review B*, vol. 87, p. 045303, jan 2013.

[103] R. Cerna, Y. Léger, T. K. Paraïso, M. Wouters, F. Morier-Genoud, M. T. Portella-Oberli, and B. Deveaud, "Ultrafast tristable spin memory of a coherent polariton gas," *Nature Communications*, vol. 4, no. May, p. 2008, 2013.

[104] G. Grosso, S. Trebaol, M. Wouters, F. Morier-Genoud, M. T. Portella-Oberli, and B. Deveaud, "Nonlinear relaxation and selective polychromatic lasing of confined polaritons," *Physical Review B*, vol. 90, p. 045307, jul 2014.

[105] W. L. Zhang, F. Wang, Y. J. Rao, R. Ma, and X. M. Wu, "Spin-Dependent Polaritonic Flip-Flop Operation in Semiconductor Microcavity," *Journal of Lightwave Technology*, vol. 33, no. 18, pp. 3933–3937, 2015.

[106] G. Christmann, R. Butt, E. Feltin, J. F. Carlin, and N. Grandjean, "Room temperature polariton lasing in a GaNAlGaN multiple quantum well microcavity," *Applied Physics Letters*, vol. 93, no. 5, pp. 1–4, 2008.

[107] D. G. Lidzey, D. D. C. Bradley, M. S. Skolnick, T. Virgili, S. Walker, and D. M. Whittaker, "Strong exciton-photon coupling in an organic semiconductor microcavity," *Nature*, vol. 395, no. 6697, pp. 53–55, 1998.

[108] A. Pawlis, A. Khartchenko, O. Husberg, D. As, K. Lischka, and D. Schikora, "Large room temperature Rabi-splitting in a ZnSe/(Zn,Cd)Se semiconductor microcavity structure," *Solid State Communications*, vol. 123, pp. 235–238, aug 2002.

[109] R. Butté, G. Christmann, E. Feltin, J. F. Carlin, M. Mosca, M. Ilegems, and N. Grandjean, "Room-temperature polariton luminescence from a bulk GaN microcavity," *Physical Review B - Condensed Matter and Materials Physics*, vol. 73, no. 3, pp. 1–4, 2006.

[110] X. Liu, T. Galfsky, Z. Sun, F. Xia, E.-C. Lin, Y.-H. Lee, S. Kéna-Cohen, and V. M. Menon, "Strong light-matter coupling in two-dimensional atomic crystals," *Nature Photonics*, vol. 9, pp. 30–34, dec 2014.

[111] I. Akasaki, H. Amano, A. Nobel, and W. Akasaki, "Blue LEDs - Filling the world with new light," *The Royal Swedish Academy of Sciences-The Nobel Prize in Physics*, 2014.

[112] N. Antoine-Vincent, F. Natali, D. Byrne, A. Vasson, P. Disseix, J. Leymarie, M. Leroux, F. Semond, and J. Massies, "Observation of Rabi splitting in a bulk GaN microcavity grown on silicon," *Physical Review B*, vol. 68, no. 15, pp. 1–4, 2003.

[113] G. Christmann, D. Simeonov, R. Butté, E. Feltin, J. F. Carlin, and N. Grandjean, "Impact of disorder on high quality factor III-V nitride microcavities," *Applied Physics Letters*, vol. 89, no. 26, pp. 1–4, 2006.

[114] J. F. Carlin, C. Zellweger, J. Dorsaz, S. Nicolay, G. Christmann, E. Feltin, R. Butté, and N. Grandjean, "Progresses in III-nitride distributed Bragg reflectors and microcavities using AlInN/GaN materials," *Physica Status Solidi (B) Basic Research*, vol. 242, no. 11, pp. 2326–2344, 2005.

[115] P. Kelkar, V. Kozlov, H. Jeon, A. Nurmikko, C.-C. Chu, D. Grillo, J. Han, C. Hua, and R. Gunshor, "Excitons in a II-VI semiconductor microcavity in the strong-coupling regime," *Physical Review B*, vol. 52, no. 8, pp. 5491–5494, 1995.

[116] A. Pawlis, *Development and investigation of II-VI semiconductor microcavity structures*. PhD thesis, Universität Paderborn, 2003.

[117] A. Balocchi, A. Curran, T. C. M. Graham, C. Bradford, K. A. Prior, and R. J. Warburton, "Epitaxial liftoff of ZnSe-based heterostructures using a II-VI release layer," *Applied Physics Letters*, vol. 86, no. 1, p. 011915, 2005.

[118] A. Rajan, I. A. Davidson, R. T. Moug, and K. A. Prior, "Epitaxial lift-off of II-VI semiconductors from III-V substrates using a MgS release layer," *Journal of Applied Physics*, vol. 114, p. 243510, dec 2013.

[119] A. Curran, J. K. Morrod, K. A. Prior, A. K. Kar, and R. J. Warburton, "Exciton-photon coupling in a ZnSe-based microcavity fabricated using epitaxial liftoff," *Semiconductor Science and Technology*, vol. 22, pp. 1189–1192, nov 2007.

[120] A. Curran, *Exciton-photon hybridisation in ZnSe based microcavities*. PhD thesis, Heriot-Watt University, 2008.

[121] K. Sebald, M. Seyfried, S. Klembt, S. Bley, A. Rosenauer, D. Hommel, and C. Kruse, "Strong coupling in monolithic microcavities with ZnSe quantum wells," *Applied Physics Letters*, vol. 100, no. 16, p. 161104, 2012.

[122] D. Lidzey, D. Bradley, T. Virgili, A. Armitage, M. Skolnick, and S. Walker, "Room Temperature Polariton Emission from Strongly Coupled Organic Semiconductor Microcavities," *Physical Review Letters*, vol. 82, no. 16, pp. 3316–3319, 1999.

[123] P. Schouwink, H. Von Berlepsch, L. Dähne, and R. F. Mahrt, "Observation of strong exciton-photon coupling in an organic microcavity," *Chemical Physics Letters*, vol. 344, no. 3-4, pp. 352–356, 2001.

[124] P. Schouwink, H. Von Berlepsch, L. Dähne, and R. F. Mahrt, "Observation of strong exciton-photon coupling in an organic microcavity in transmission and photoluminescence," *Journal of Luminescence*, vol. 94-95, pp. 821–826, 2001.

[125] P. Schouwink, H. Berlepsch, L. Dahne, and R. F. Mahrt, "Dependence of Rabi-splitting on the spatial position of the optically active layer in organic microcavities in the strong coupling regime," *Chemical Physics*, vol. 285, no. 1, pp. 113–120, 2002.

[126] P. Schouwink, J. Lupton, H. von Berlepsch, L. Dähne, and R. Mahrt, "Nonequilibrium polariton dynamics in organic microcavities," *Physical Review B*, vol. 66, no. 8, pp. 2–5, 2002.

[127] N. Takada, T. Kamata, and D. D. C. Bradley, "Polariton emission from polysilane-based organic microcavities," *Applied Physics Letters*, vol. 82, no. 12, pp. 1812–1814, 2003.

[128] R. J. Holmes and S. R. Forrest, "Strong exciton-photon coupling and exciton hybridization in a thermally evaporated polycrystalline film of an organic small molecule," *Physical Review Letters*, vol. 93, no. 18, p. 186404, 2004.

[129] J. H. Song, Y. He, A. V. Nurmikko, J. Tischler, and V. Bulovic, "Exciton-polariton dynamics in a transparent organic semiconductor microcavity," *Physical Review B - Condensed Matter and Materials Physics*, vol. 69, no. 23, pp. 1–5, 2004.

[130] J. Wenus, R. Parashkov, S. Ceccarelli, A. Brehier, J. S. Lauret, M. S. Skolnick, E. Deleporte, and D. G. Lidzey, "Hybrid organic-inorganic exciton-polaritons in a strongly coupled microcavity," *Physical Review B*, vol. 74, no. 23, p. 235212, 2006.

[131] R. J. Holmes and S. R. Forrest, "Strong exciton-photon coupling in organic materials," *Organic Electronics*, vol. 8, pp. 77–93, 2007.

[132] S. Kéna-Cohen, M. Davanço, and S. Forrest, "Strong Exciton-Photon Coupling in an Organic Single Crystal Microcavity," *Physical Review Letters*, vol. 101, no. 11, pp. 1–4, 2008.

[133] P. Michetti and G. C. La Rocca, "Polariton states in disordered organic microcavities," *Physical Review B*, vol. 71, no. 11, p. 115320, 2005.

[134] P. Michetti and G. C. La Rocca, "Polariton dynamics in disordered microcavities," *Physica E*, vol. 40, no. 6, pp. 1926–1929, 2008.

[135] M. Litinskaya, P. Reineker, and V. M. Agranovich, "Exciton-polaritons in organic microcavities," *Journal of Luminescence*, vol. 119, pp. 277–282, 2006.

[136] G. M. Akselrod, Y. R. Tischler, E. R. Young, D. G. Nocera, and V. Bulovic, "Exciton-exciton annihilation in organic polariton microcavities," *Physical Review B*, vol. 82, p. 113106, sep 2010.

[137] A. Splendiani, L. Sun, Y. Zhang, T. Li, J. Kim, C. Y. Chim, G. Galli, and F. Wang, "Emerging photoluminescence in monolayer MoS2," *Nano Letters*, vol. 10, no. 4, pp. 1271–1275, 2010.

[138] K. F. Mak, C. Lee, J. Hone, J. Shan, and T. F. Heinz, "Atomically thin MoS_2: A new direct-gap semiconductor," *Physical Review Letters*, vol. 105, no. 13, pp. 2–5, 2010.

[139] S. Schwarz, S. Dufferwiel, P. M. Walker, F. Withers, A. A. P. Trichet, M. Sich, F. Li, E. A. Chekhovich, D. N. Borisenko, N. N. Kolesnikov, K. S. Novoselov, M. S. Skolnick, J. M. Smith, D. N. Krizhanovskii, and A. I. Tartakovskii, "Two-dimensional metal-chalcogenide films in tunable optical microcavities," *ACS Nano*, vol. 14, pp. 7003–7008, 2014.

[140] S. Dufferwiel, F. Li, E. Cancellieri, L. Giriunas, A. A. P. Trichet, D. M. Whittaker, M. Walker, F. Fras, E. Clarke, J. M. Smith, M. S. Skolnick, and D. N. Krizhanovskii, "Spin Textures of Polariton Condensates in a Tunable Microcavity with Strong Spin-Orbit Interaction," *ArXiv*, 2015.

[141] M. I. Vasilevskiy, D. G. Santiago-Pérez, C. Trallero-Giner, N. M. R. Peres, and A. Kavokin, "Exciton-polaritons in 2D dichalcogenide layers placed in a planar microcavity: tuneable interaction between two Bose-Einstein condensates," *ArXiv*, vol. 4, no. 1, pp. 1–15, 2015.

[142] W. Liu, B. Lee, C. H. Naylor, H.-S. Ee, J. Park, and A. T. Charlie, "Strong exciton-plasmon coupling in MoS_2 coupled with plasmonic lattice," *arXiv*, p. 1511.03750, 2015.

[143] K. F. Mak, K. He, C. Lee, G. H. Lee, J. Hone, T. F. Heinz, and J. Shan, "Tightly bound trions in monolayer MoS_2," *Nature Materials*, vol. 12, no. 3, pp. 207–11, 2013.

[144] A. Das, J. Heo, M. Jankowski, W. Guo, L. Zhang, H. Deng, and P. Bhattacharya, "Room temperature ultralow threshold GaN nanowire polariton laser," *Physical Review Letters*, vol. 107, no. 6, pp. 1–5, 2011.

[145] J. Heo, S. Jahangir, B. Xiao, and P. K. Bhattacharya, "Room temperature polariton lasing from GaN nanowire array clad by dielectric microcavity," *Nano Letters*, vol. 13, pp. 2376–2380, 2013.

[146] A. Das, J. Heo, A. Bayraktaroglu, W. Guo, T.-K. Ng, J. Phillips, B. S. Ooi, and P. Bhattacharya, "Room temperature strong coupling effects from single ZnO nanowire microcavity," *Optics Express*, vol. 20, no. 11, p. 11830, 2012.

[147] L. K. van Vugt, S. Rühle, P. Ravindran, H. C. Gerritsen, L. Kuipers, and D. Vanmaekelbergh, "Exciton Polaritons Confined in a ZnO Nanowire Cavity," *Physical Review Letters*, vol. 97, no. 14, p. 147401, 2006.

[148] D. Xu, W. Xie, W. Liu, J. Wang, L. Zhang, Y. Wang, S. Zhang, L. Sun, X. Shen, and Z. Chen, "Polariton lasing in a ZnO microwire above 450K," *Applied Physics Letters*, vol. 104, p. 082101, 2014.

[149] S.-H. Gong, S.-M. Ko, M.-H. Jang, and Y.-H. Cho, "Giant Rabi Splitting of Whispering Gallery Polaritons in GaN/InGaN CoreShell Wire," *Nano Letters*, p. 150610090401002, 2015.

[150] T.-C. Lu, Y.-Y. Lai, Y.-P. Lan, S.-W. Huang, J.-R. Chen, Y.-C. Wu, W.-F. Hsieh, and H. Deng, "Room temperature polariton lasing vs photon lasing in a ZnO-based hybrid microcavity," *Optics Express*, vol. 20, no. 5, p. 5530, 2012.

[151] F. Li, L. Orosz, O. Kamoun, S. Bouchoule, C. Brimont, P. Disseix, T. Guillet, X. Lafosse, M. Leroux, J. Leymarie, M. Mexis, M. Mihailovic, G. Patriarche, F. Réveret, D. Solnyshkov, J. Zuniga-Perez, and G. Malpuech, "From Excitonic to Photonic Polariton Condensate in a ZnO-Based Microcavity," *Physical Review Letters*, vol. 110, no. 19, p. 196406, 2013.

[152] J.-H. Jiang and S. John, "Photonic Crystal Architecture for Room-Temperature Equilibrium Bose-Einstein Condensation of Exciton Polaritons," *Physical Review X*, vol. 4, no. 3, p. 031025, 2014.

[153] J.-H. Jiang and S. John, "Photonic Architectures for Equilibrium High-Temperature Bose-Einstein Condensation in Dichalcogenide Monolayers," *Scientific Reports*, vol. 4, p. 7432, 2014.

[154] S. Ogawa, M. Imada, S. Yoshimoto, M. Okano, and S. Noda, "Control of light emission by 3D photonic crystals.," *Science (New York, N.Y.)*, vol. 305, no. 5681, pp. 227–229, 2004.

[155] A. Rubio-Ponce, D. Olguín, and I. Hernández-Calderón, "Calculation of the effective masses of II-VI semiconductor compounds," *Superficies y Vacio*, vol. 16, no. 2, pp. 26–28, 2003.

[156] Y. Sun, Y. Yoon, M. Steger, G. Liu, L. N. Pfeiffer, K. West, D. W. Snoke, and K. A. Nelson, "Polaritons are Not Weakly Interacting: Direct Measurement of the Polariton-Polariton Interaction Strength," *ArXiv*, pp. 1–27, 2015.

[157] S. Gehrsitz, F. Reinhart, C. Gourgon, N. Herred, V. A, and H. Sigg, "The refractive index of AlGaAs below the band gap: Accurate determination and empirical modeling," *Journal of Applied Physics*, vol. 87, no. 11, pp. 7825–7836, 2000.

[158] G. Nardin, K. G. Lagoudakis, M. Wouters, M. Richard, A. Baas, R. André, L. S. Dang, B. Pietka, and B. Deveaud-Plédran, "Dynamics of Long-Range Ordering in an Exciton-Polariton Condensate," *Physical Review Letters*, vol. 103, p. 256402, dec 2009.