Characterization, Modeling, And Design Of Microcavity Exciton-Polariton Samples

by

Jonathan C. Beaumariage

B.S.,Physics, Rochester Institute of Technology, 2013

M.S.,Physics, University of Pittsburgh, 2022

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

2024

UNIVERSITY OF PITTSBURGH

KENNETH P. DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Jonathan C. Beaumariage

It was defended on

August 1st, 2024

and approved by

David W. Snoke, Department of Physics and Astronomy

Gurudev Dutt, Department of Physics and Astronomy

Arthur Kosowsky, Department of Physics and Astronomy

David Pekker, Department of Physics and Astronomy

Sean Garrett-Roe, Department of Physics and Astronomy

Characterization, Modeling, And Design Of Microcavity Exciton-Polariton Samples

Jonathan C. Beaumariage, PhD

University of Pittsburgh, 2024

Exciton-polaritons are particles combining the aspects of solid matter with light. They interact and scatter off one another, while also having a very light effective mass. They thermalize well, but are also bosons, meaning they can condense into Bose-Einstein condensates at much higher temperatures than many other systems. They are typically generated in semiconductor structures grown in extremely precise molecular beam epitaxy machines. Very high quality samples are commonplace, and long polariton lifetimes over 100 ps are observed.

However, when the samples are of such high quality, the polaritons may no longer be easily observed in reflectivity measurements, because they have surpassed the resolution of normal laboratory spectrometers. Also, one of the two polariton states, known as the upper polariton, is no longer visible in standard photoluminescence measurements, because its decay is dominated by down-scattering into lower polariton states. This has been a consistent problem for calculating the exciton fraction, which is an essential parameter in many-body theories that include the polariton-polariton interactions.

In this work, I present my work on accurate calculation of the exciton fraction of polaritons in semiconductor microcavities. This work includes redesign of several experiments, creation of a photoluminescence excitation measurement setup, which allows detection of of the upper polariton, and simulation of the electromagnetic properties of polariton microcavity structures using the transfer-matrix method. These simulations show great success, and may be used for designing future samples. Additionally, I present work on cavity structures with transition-metal dichalcogenide monolayers, as a means towards room temperature polariton physics.

Table of Contents

List of Tables

List of Figures

Preface

First, I would like to thank my advisor Dr. David Snoke. Without his tireless help and support I could never have completed this work. He shows an unending commitment towards scientific understanding, while also volunteering countless hours of his time towards sharing that understanding. The University of Pittsburgh is very fortunate to have such a professor.

Additionally, I would like to thank the other four members of my committee; Drs. Gurudev Dutt, Sean Garrett-Roe, Arthur Kosowsky, and David Pekker. Specifically I wish to thank them for reviewing my work, and providing guidance towards strengthening it.

Furthermore, I would like to thank the National Science Foundation for their funding through Grant No. DMR-2306977 as well as U.S. Army Research Office for their funding under MURI award W911NF1710312.

Next, I would like to thank my parents John and Mary Beaumariage. Every trip to Darien Lake, every time you drove me home from bowling, every time you helped me with my schooling; I remember it. You gave me everything I could ever ask for in life, I could never fully express my gratitude.

I would also like to thank my grandmother Matilda "Tilly" Guerriero, whose house was always a home away from home. The various machines and tools I used on her property inspired me, and steered me towards experimental physics.

Thank you to my aunt Irene Guerriero and uncle Fred Guerriero. The bonds that hold a family together are not free, they require constant maintenance. The work that Fred and Irene do helps keeps our family woven together. Additionally, I would like to thank my uncle Vincent Guerriero, and my cousins and respective cousin-in-laws Kent Guerriero, Manessa Zamingir, Jim Guerriero, Meghan Brown, Kate Guerriero, Kyle Van Pelt, Beth Guerriero, and Matt Potts. I would also like to thank my extended family, in particular Thomas E. and Paula L. Johnson. My life is improved because of each one of you.

I would also like to thank all my friends who have supported me throughout this work. Particularly Benjamin Liu who helped me with countless MATLAB questions, Joshua Kahn who was happy to discuss physics with me, and Héctor Martínez Rodríguez who fought through the core classes with me. Additionally I would like to thank Emma Anderson, Matt Anderson, Laurel Buchanan and all my other friends. Each one of you brightens my life, and I'm very grateful to have met you all.

My labmates have also been important in this work. Thank you to Mark Steger, Chitra Gautham, and Gangqiang Liu for teaching me the basics of work in our lab. Thank you also to Ryan Balili who was very helpful in his visits. Thank you to David Myers, Shouvik Mukherjee, Qi Yao, Adam Grusky, Hassan Alnatah, Benjamin Ash, Qiaochu Wan, and Daniel Edward Vaz; who have helped me with countless things in the lab. I've also enjoy our many conversations. I would also like to thank our two postdocs Burcu Ozden and Zheng Sun for their help, especially Zheng Sun, who was very helpful in collecting the data.

Furthermore I would like to thank all the staff at Pitt who have helped me along my way. In particular I would like to mention the great people who work in the machine shop. Countless pieces of equipment in our lab don't work without their help. In particular I've worked a lot with Jeff Tomaszewski, Shawn Artman, Thomas Gasmire and Jeffrey Sicher.

Additionally, countless teachers have inspired my love of science and mathematics and taught me priceless lessons, who I would like to thank. Specifically I would like to thank Mr. and Mrs. Hill and coach Mike McFadden from Hamburg High School, as well as many others. I would also like to thank my many great professors at the Rochester Institute of Technology, especially Dr. David Ross and Dr. John Andersen.

Finally, I wish to thank all who have ever embodied the spirit of scientific discovery. Through our collective efforts, we are able to draw understanding from the chaos.

1.0 Introduction

In countless science fiction media we've seen depictions of light behaving as a solid object. We've seen sparks fly as swords of light clash with each other, cars drive on bridges made of light, and light used as a prison to restrain our would-be heroes. But back here on Earth, two beams of light don't bounce off one another. However, in the world of condensed matter physics, every solid object can be seen as simulating its own little universe. Inside one crystal, we can see the chemistry of a universe where the proton and electron have the same mass. Inside another, we can see a universe where electrons attract one another, and pair up; and inside some special materials we can even see a universe where light flows like water.

In this Chapter, we will give an introduction to some of the various topics of condensed matter physics, and describe the ways we can blur the lines between solid matter and light.

As with any field, ours has plenty of jargon and abbreviations. Table 1 lists common abbreviations used in this thesis.

1.1 Electron Bands

Typically, we are interested in crystals which are highly ordered periodic structures. These crystals are of course simply made up of atoms. The nuclei of those atoms provide the fixed points of a lattice, while the electrons are much more free to move. For this reason, the electrons are where we focus our efforts. In fact, even most of the electrons are not free to move. Rather, the outermost shell of electrons are the ones ultimately responsible for most of the phenomenon we are concerned with. For a more thorough explanation of solid state physics see Refs. 73 and 49.

First we wish to explain where the phenomenon of electron bands originates. To begin, we consider a simple model known as *linear combination of atomic orbitals* (LCAO), but the results we derive are broadly applicable to many systems. Consider an atom with some number of non-degenerate eigenstates with energies E_n, whose Hamiltonian is written simply

Table 1: Table of common abbreviations

Abbreviation	Meaning
CCD	charge-coupled device (camera)
CW	Continuous Wave
DBR	Distributed Bragg Reflector
FWHM	Full Width at Half Maximum
GUI	Graphical User Interface
LP	Lower Polariton
MBE	Molecular Beam Epitaxy
NA	Numerical Aperture
PECVD	Plasma Enhanced Chemical Vapor Deposition
PCI	Peripheral Component Interconnect (card slot on motherboards)
PL	Photoluminescence
RHEED	Reflection High Energy Electron Diffraction
TE	Transverse Electric
TMD	Transition Metal Dichalcogenide (Monolayer)
TM	Transverse Magnetic
TMM	Transfer Matrix Method
UP	Upper Polariton

as

$$H_1 = \begin{bmatrix} E_1 & 0 & 0 & 0 & \\ 0 & E_2 & 0 & 0 & \\ 0 & 0 & E_3 & 0 & \\ 0 & 0 & 0 & \ddots & \\ & & & & \end{bmatrix}. \tag{1.1}$$

Now we consider bringing in a second identical atom and allowing weak coupling with the first atom. This interaction shows up as a coupling strength between the various eigenstates.

In general these coupling strengths are all independent, and come from the specifics of our physical system. For simplicity, we will limit our written equations to just three energy levels going forward, but later when we simulate our system we are free to add more. The Hamiltonian of the two-atom system is now

$$H_2 = \begin{bmatrix} E_1 & 0 & 0 & g_{1,1} & g_{1,2} & g_{1,3} \\ 0 & E_2 & 0 & g_{2,1} & g_{2,2} & g_{2,3} \\ 0 & 0 & E_3 & g_{3,1} & g_{3,2} & g_{3,3} \\ g_{1,1} & g_{1,2} & g_{1,3} & E_1 & 0 & 0 \\ g_{2,1} & g_{2,2} & g_{2,3} & 0 & E_2 & 0 \\ g_{3,1} & g_{3,2} & g_{3,3} & 0 & 0 & E_3 \end{bmatrix} = \begin{bmatrix} H_1 & G \\ G & H_1 \end{bmatrix}, \tag{1.2}$$

where we have rewritten it in the block matrix notation, defining G accordingly. Next, we wish to bring in many more atoms. However, we restrict ourselves so that atoms are only interacting with their nearest neighbor. Our Hamiltonian now has the form

$$H_n = \begin{bmatrix} H_1 & G & 0 & 0 & 0 & 0 & & \\ G & H_1 & G & 0 & 0 & 0 & & \\ 0 & G & H_1 & G & 0 & 0 & & \\ 0 & 0 & G & H_1 & G & 0 & & \\ 0 & 0 & 0 & G & H_1 & G & & \\ & & & & & \ddots & & \\ & & & & & & & G \\ & & & & & & G & H_1 \end{bmatrix}. \tag{1.3}$$

If we had allowed coupling between every pair of atoms, the block zero matrices would have been replaced by block G matrices.

We now wish to find the eigenvalues of the system's Hamiltonian. We need specifics to proceed further, but we find the specifics we choose are not so important. As long as we choose the values of G to be small relative to the separation between the original energies E_n, we find that the energies of our new many-atom system bunch together in distinct bands.

For example, in Figure 1 we have chosen E_n to be the energies of the electron in the Bohr model of the hydrogen atom, given as

$$E_n = \frac{-13.6\ eV}{n^2}. \tag{1.4}$$

For Figure 1A we chose a constant matrix for G, with $g_{i,j} = 0.1$ eV, and we show the resulting energies for 1, 2, 3, and 100 atoms. We limited ourselves to just the first four energies of the model. We see that at 100 atoms, essentially have a continuum of single-electron energies, forming bands near where the original single-electron energy states were.

However, our particular choices were not too important. This phenomenon of bands arises often whenever a large number of identical systems are coupled together. To demonstrate this, we introduce some randomness into our model. In Figure 1B, we have chosen G to be a random matrix, with each element having a separate random value between 0 and 0.2 eV. In Figure 1C, we let the random values range from 0 to 0.4 eV, and in 1D the range is 0 to 0.8 eV.

We could have chosen from many different physical systems for E_n, but as long as the values of the coupling matrix are small relative to the separation, we see these distinct bands form. However, as Figure 1D shows, the bands will begin to mix together if the coupling is strong enough. Notice that due to the nature of the Bohr model's energy states, the higher energy bands are closer together than the lower energy bands. In practice, this often means that lower energy bands are distinctly separated with gaps between them, and higher energy bands mix. This leads to the high energy continuum that many objects exhibit, where above a minimum energy all energies have viable states.

Now that we have some understanding of what states are available, we can begin filling those states with electrons. In order to proceed, I must now discuss many-particle states. Everything so far has been calculating the states available to a single electron. Of course, electrons experience a Coulomb repulsion, and so a thorough proper treatment would involve including that term into the Hamiltonian. However, as discussed in many textbooks (e.g. Ref. [73], most of the Coulomb interaction of the electrons and atomic nuclei can be accounted for as an effective potential energy, in what is called density functional theory. This electron-density-dependent potential leads to corrections to the single-particle band

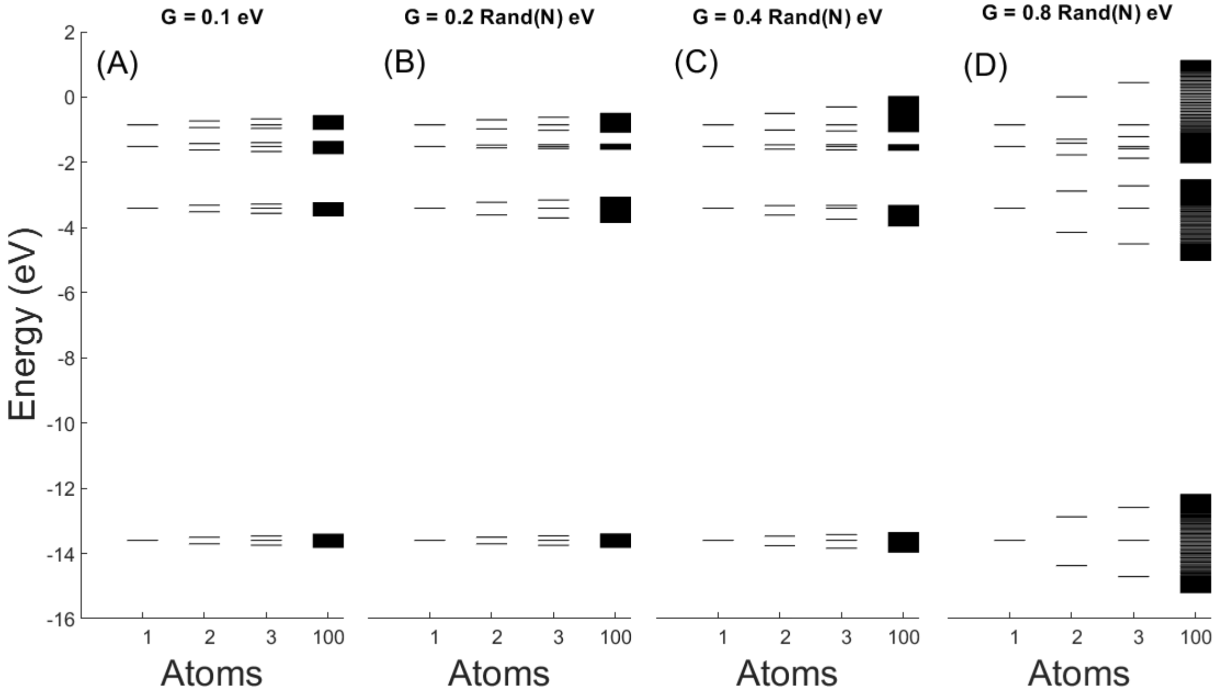

Figure 1: Energies of the single-electron eigenstates for an exact solution of the LCAO model. We see that when large numbers of identical atoms couple to one another, energy bands form. In A, we chose the elements of our coupling matrix G to all be 0.1 eV. In B, C, and D, all elements of G are randomly chosen between 0 and 0.2 eV, 0.4 eV, and 0.8 eV, respectively.

energies, after which the electrons can be treated as non-interacting. Only when there is an excitation out of the ground state (e.g., the creation of an exciton, discussed I will discuss) do we need to account for the Coulomb interactions.

Although our simulation only went up to 100 atoms, we are now going to consider a much larger system, nominally 10^{23} atoms. Although accurate calculation of electron bands is far more complicated, our analysis above does give us an approximate magnitude for what to expect. The energy scale of band gaps are typically set by the Rydberg energy [73], and so band gaps are often a few electron-Volts. This is is quite large compared to $k_B T$, which at room temperatures is about 25 meV. This means all of the electrons are heavily favored to fall down into the lowest band.

To put N electrons into the empty bands, we simply add in the electrons with a probability given by the appropriate statistical distribution. For electrons at low density, the relative occupation probability of two states with energies E_1 and E_2 is [64]

$$\frac{P_2}{P_1} = e^{\frac{(E_1-E_2)}{k_B T}}. \tag{1.5}$$

However, we have our Pauli exclusion principle. which prevents them from all falling into the lowest band. Combining equation (1.5) with our Pauli exclusion principle leads to the Fermi-Dirac distribution [71],

$$\langle n_E \rangle = \frac{1}{e^{(E-E_f)/k_B T} + 1}. \tag{1.6}$$

The distribution is characterized by a special energy level which unfortunately has two common names; the "Fermi level" E_f, which is also commonly called the "chemical potential" μ.

The bands which arise directly relate to the orbitals of the atoms, and the number of states in the bands directly corresponds to the number of electron states in the orbital shells. Therefore, if the orbital was full before the atoms combined to form our solid, the corresponding band will be filled as well. Similarly, the higher energy orbitals which are empty will correspond to empty higher energy bands; with the occasional thermal fluctuation filling a small number of states in these higher energy bands.

At zero temperature, every electron state above the Fermi level is empty, and all states below it are filled. At finite temperature, an electron state at the Fermi level has a 50% chance of being occupied at thermal equilibrium. However, as discussed above, the scale of bands and band gaps is typically much larger than $k_B T$, and so even at room temperature we will often speak as though all electrons are below the Fermi level. Going forward we may colloquially think of the Fermi level as being smeared out over the range $E_f \pm k_B T$.

The specifics of where a material's Fermi level lies relative to its bands structure is one of its most critical defining characteristics. This is because it tells us how the electrons inside the material may absorb energy. There are two main cases which may occur: the Fermi energy may lie within a continuum of energy states, or it may lie with a gap of states.

A Fermi level may lie within a continuum of states by being within a single band, or by being between two bands which overlap in energy due to band smearing, as seen, e.g., in Figure 1D. Either way, if some infinitesimally small impulse of energy were delivered to the material, the electrons just below the Fermi level could absorb it as there is a continuum of unfilled states directly above for them to transition into. This is the hallmark of a conductor; electrons flow even with very small applied voltages. A simple cartoon showing this situation is shown in Figure 2. Conductors are not of interest to our work, but we mention them here for completeness.

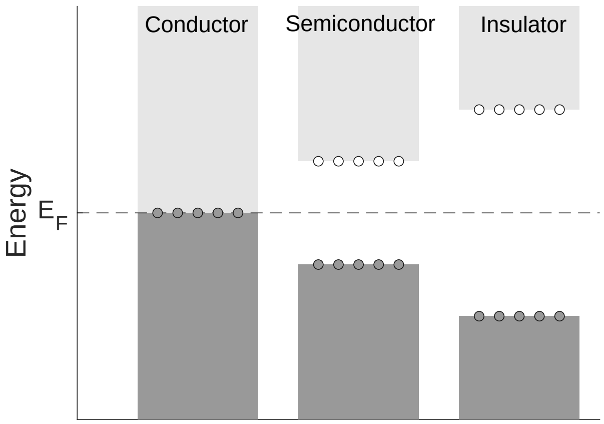

Figure 2: A simple illustration of the Fermi level relative to the bands of a conductor, semiconductor, and insulator. The Fermi level is given by the dashed line. Electron states below it are filled, as indicated by the darker color. A few filled states are represented as filled in circles, and a few empty states are represented as unfilled circles.

The second case, in which the Fermi level lies within a gap, is where most of our work will be conducted. This situation is further broken down into two subcategories, insulators and semiconductors. The main difference is the size of the band gap. Nominally a semiconductor has a band gap between 0.3 eV and 3 eV, while an insulator has a band gap greater than 3 eV [73]. Simple illustrations of this situation are shown in Figure 2.

Commonly, these materials are made by a specific combination of atoms that results in the outermost orbital shells being completely full, which then corresponds to the Fermi level being in a gap. As such, it is most useful to use the Chemical Abstracts Service (CAS) naming scheme for periodic table groups. The CAS naming convention allows us to quickly see that if the numbers of the groups add up to eight, we will fill our shell. However, as is common in the field, we drop the letters from our group names as we are working exclusively with elements in the P block. So for example, we may combine gallium (Ga) which is in the III group with Arsenic (As) in the V group to create the III-V semiconductor gallium arsenide (GaAs). The classification of GaAs as a semiconductor rather than an insulator comes from experimental measurements of the band gap, which we've measured to be around 1.425 eV at room temperature and 1.5 eV at cryogenic temperatures. The elements we combine to form our semiconductors and insulators do not have to be in the same period either. We may combine aluminum and arsenic to create another III-V semiconductor aluminum arsenide (AlAs). Generally AlAs is less chemically stable than GaAs, and more at risk of oxidation when exposed to atmosphere [61]. We say, "risk" here because in our work oxidation is generally bad, but depending on the goals of the project it can also be a beneficial property.

The band directly below the Fermi level is typically called the valence band, and the band directly above it is typically called the conduction band. The simple illustration in Figure 2 shows a semiconductor at absolute zero temperature. At a finite temperature, a small number of electrons will be excited out of the valence band into the conduction band. These conduction band electrons will generally have a continuum of states around them which they can move into if an external voltage is applied, and so a small amount of current may be induced even with a low voltage.

There are other ways to get electrons into the conduction band of a semiconductor, such as doping and gating. The ability to modify the conduction behavior of these materials is the origin of the name, "semiconductor."

1.2 Renormalization, Holes, And The Exciton

In our prior discussion of conduction inside of semiconductors, we see that the electrons in energy states closest to the Fermi level are very important for that discussion. But all the other lower energy electrons are effectively a background term, rather than contributing directly to any of the interesting physics. This motivates us to adopt the renormalized vacuum point of view. In this picture, rather than trying to create a model of physics that accounts for all $\sim 10^{23}$ protons and electrons, we ignore most of them, and pay attention only to the select few that are most relevant.

Of course, the ignored particles still play a role, and they are not truly ignored. Instead, their influence shows up as modifications of our standard physics equations. For example, an electron in the conduction band may behave as though it has a smaller mass than the mass we measure for electrons in vacuum; or the screening effects may modify the effective value of ϵ inside our Maxwell's equations.

One of the most important concepts in this renormalized vacuum picture is the concept of holes. We begin by viewing the semiconductor with a completely filled valence band and a completely empty conduction band as our vacuum state. We may now consider excitations out of the vacuum state. Of particular interest to us will be situations in which an impulse of energy is absorbed by our system, driving an electron out of the valence band into the conduction band. Of particular interest to us will be the absorption of a photon of light, with an energy close to the band gap energy. A simple illustration of this situation is shown in Figure 3.

Relative to the default negative charges in the valence band in the vacuum state, the lack of a negative charge in the valence band in the excited state acts like a positive charge. This is similar to the method of image charges employed near the surface of a conductor [42], where the presence of a macroscopic number of charges mimics the effects of a single particle. This positive charge is referred to as a "hole." We also still have the negatively charged electron in the conduction band. The electron has a continuum of nearby states available to it, and so it may freely move and for this reason is called a "free electron" to distinguish it from the sea of electrons filling the valence band. The hole also may easily move about by having any nearby electron in the valence band fill its current state, thus opening up a new state for the hole to occupy.

These particles are routinely called “quasiparticles”. A deep philosophical question comes up here regarding these quasiparticles [51]: are they real particles? Furthermore, what is it that makes the isolated electrons we experience outside of solid materials real? Can we be so sure that we ourselves are not already in some other renormalized vacuum? Fortunately, we do not need to answer such deep questions as the mathematics governing our system are laid out and not dependent on those answers.

With our renormalized vacuum picture in place, we now see the excited state in Figure 3 as a pair production of oppositely charged particles near one another, and so we expect some Coulomb attraction between them. Because of our renormalized vacuum, we expect physical constants from our Maxwell’s equations to possibly be modified as well as things such as the mass of our particles. These are not bare electrons and positrons, but they will behave similarly.

If a very large amount of energy was responsible for this pair production, then the two particles will have enough kinetic energy to fly infinitely far away from each other. This is analogous to a positively ionized hydrogen atom (except for differences in the relative masses which we will discuss). However, with a smaller amount of energy the hole and free electron may bind to and orbit one another, more akin to the standard hydrogen atom. This two particle complex is called the “exciton”, and it is one of the fundamental excitations of solid state systems, and commonly found in many materials.

In order to better understand the exciton, we need to delve further into the specifics of electron bands and the dispersion curves of the electrons. We’ve shown that when a large number of identical systems are coupled together, bands are formed. However, we have not discussed the specifics of the energy states within the bands. Similar to electrons in a vacuum, there is a dispersion relationship between the energy E and the momentum $\vec{k}$ (technically this is the wave number, but in solid state physics and optics it is commonly referred to as momentum). However, there is a crucial difference. In a vacuum there is continuous rotational symmetry in all directions, and so a simple dispersion relationship exists in which only the magnitude of the momentum vector is relevant:

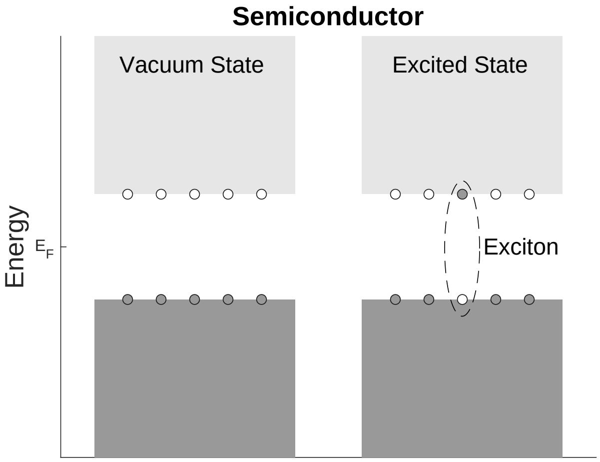

Figure 3: On the left side we have our newly defined renormalized vacuum state. On the right, we have the result after the system has absorbed some energy, causing a valence electron to jump into the conduction band, and creating a hole in the valence band. The electron and hole may bind together to form an exciton.

$$E = \frac{\hbar^2 \vec{k} \cdot \vec{k}}{2m} = \frac{\hbar^2 k^2}{2m} \tag{1.7}$$

We are most interested in GaAs, so we will now narrow our focus. Inside GaAs the atoms are arranged in a highly ordered manner. Specifically, our nuclei form the fixed points of a lattice. This lattice is made by periodic repetition of a unit cell. So, rather than continuous rotational symmetry we have discrete angles about which we have rotational symmetry, as a result the dispersion relation is far more complicated. The symmetries and periodicity in real space leads to critical points in the momentum space, which physicists typically denote with a series of Greek letters. Our work is entirely near the zone center, which is denoted as Γ and corresponds to zero momentum.

GaAs is a direct gap semiconductor, which means that the valence band has a local maximum and the conduction band has a local minimum at the Γ point. This is one of the reasons we are especially interested in GaAs. For our purposes, the presence of these extrema allow us to greatly simplify things, because any local minimum or maximum may be approximated as a parabola, and we are most interested in working around $\vec{k} = 0$, we may write the dispersion relations of our valence and conduction bands as

$$E_{c,v} = E_o \pm \frac{\hbar^2 k^2}{2m_{e,h}}. \tag{1.8}$$

Where the dispersion relationship of the conduction band utilizes the mass of the electron, and the valence band uses the mass of the hole, and E_o is different for the two bands. Because the free electrons and holes occupy different bands, they do not necessarily have the same mass. Here we are utilizing the definition of the effective mass of our particles, given as

$$m_{\text{eff}} = \frac{\hbar^2}{2\frac{d^2E}{dk^2}}. \tag{1.9}$$

So then, the procedure is now clear. First, electronic bands are calculated to obtain the true dispersion relationship $E(\vec{k})$, then a second derivative is evaluated at $\vec{k} = 0$ in order to calculate the effective mass of particles in the bands. Finally, a parabolic approximation is utilized near $\vec{k} = 0$. This analysis is just a standard Taylor series expansion about a local extremum, formatted so that our equations remind us of vacuum electron behavior. We should note that the mass of a free electron in the conduction band is not simply the mass of a vacuum electron. For examples of detailed band calculations see Reference 16. A simple qualitative illustration of their results are shown in the left image of Figure 4.

We see that actually in bulk GaAs the story is slightly more complicated, as two different bands have the same energy at the Γ point. From our definition for the effective mass, we see that the flatter band has holes with a heavier mass, and the curvier band has holes with a lighter mass. These bands are appropriately named the, “heavy-hole band” and the, “light-hole band.”

Rather than bulk material, we are going to be specifically interested in GaAs inside of quantum wells. A quantum well is created by sandwiching a thin slice of a lower-band-gap semiconductor between two slabs of larger-band-gap semiconductors. This creates a situation similar to the classic quantum mechanics problem of a particle in a one-dimensional box. In our case, we will be working with GaAs in between two slices of AlAs. This causes the heavy-hole and light-hole bands to separate [19], as shown in the right image of Figure 4. In

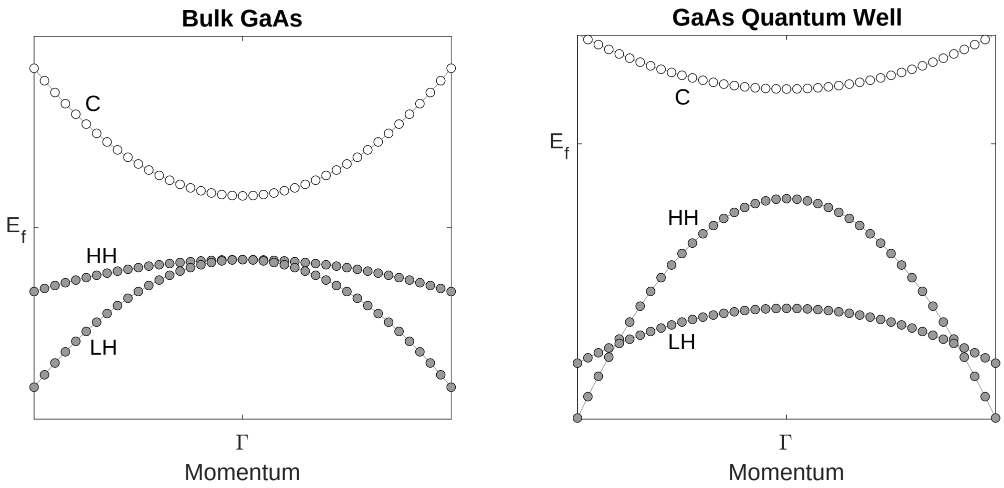

Figure 4: Qualitative illustration of parabolic approximation of GaAs conduction and valence bands. We see that making GaAs into a quantum well removes the degeneracy between the heavy-hole and light-hole bands. Ironically, the heavy-hole exciton now has the lighter mass in the in-plane direction [27].

fact, it is much more complex than the simple illustration we have shown, as multiple bands split out of the original heavy-hole and light-hole bands. But we will only be interested in the first two bands, as shown. The holes also have a directional dependent mass, which is modified when placed in a quantum well. Ironically, the in-plane mass of the heavy-holes are actually lighter than the in-plane mass of the light-holes [27], and so the energy of the holes may cross in this simple parabolic approximation.

Just as in Figure 3, either of our two lower bands in Figure 4 may absorb an appropriate amount of energy (typically in the form of a photon), and promote an electron into the conduction band while leaving behind a hole, and once again the conduction band electron and hole may coulomb bind to one another to create an exciton. So, we have two species of excitons available to us; "heavy-hole excitons" and "light-hole excitons." The mass of the exciton m_{ex} is simply the sum of the masses of its two constituent particles, as both excitons utilize the same conduction band electron, the heavy-hole exciton is heavier than the light-hole exciton. For GaAs the excitons are around 10 % of the vacuum electron mass.

The energy of either of our excitons may be approximately calculated using the Bohr model [50]. Ironically a great deal of complicated quantum mechanics is utilized in renormalizing our vacuum so that we may justify the use of simpler physics. We also account for center-of-mass motion of our exciton using the center of mass momentum k_{COM}, and so the energy of our particle is [73]

$$E_{ex} = \frac{-Ry_{ex}}{n^2} + \frac{\hbar^2 k_{COM}^2}{2m_{ex}}. \tag{1.10}$$

It should be noted that the energy associated with the creation of the electron and hole is not included. Here Ry_{ex} is the renormalized Rydberg energy, given as

$$Ry_{ex} = \frac{e^2}{8\pi\epsilon a_{ex}} = \frac{\epsilon_o}{\epsilon}\frac{a_o}{a_{ex}}\frac{e^2}{8\pi\epsilon_o a_o} = \frac{\epsilon_o}{\epsilon}\frac{a_o}{a_{ex}}\ 13.6\ \text{eV}. \tag{1.11}$$

The excitonic Rydberg energy here is also called the binding energy, as it tells us how much energy is needed to ionize our exciton in the lowest energy state. Here, the typical vacuum permittivity ϵ_o is replaced by the low frequency material specific permittivity ϵ due to the screening of all the surrounding particles. The Bohr radius $a_o = 0.529$ Å is also replaced with the renormalized exciton Bohr radius

$$a_{ex} = \frac{4\pi\hbar^2\epsilon}{e^2 m_r}. \tag{1.12}$$

Often in dealing with the hydrogen atom, the proton mass is much higher than the electron mass, and so the reduced mass m_r is effectively the electron mass. But in our case, the hole mass and the conduction band electron mass are comparable, and so we must utilize the true reduced mass

$$m_r = \frac{m_h m_e}{m_h + m_e}. \tag{1.13}$$

These equations apply to both the heavy-hole exciton and light-hole exciton, with slightly different values for m_h. As denoted by the n in equation 1.11, there are many energy states for the exciton. However, we are going to be most interested in the lowest energy state $n = 1$, as that is the one we encounter most often. It is worth mentioning that excitons are usually characterized into two types of limiting behavior based on the exciton Bohr radius. When the radius is small so that the exciton is effectively localized to a single lattice site we are in the Frenkel regime. When the radius is large so that we are spanning multiple lattice sites, we are in the Wannier-Mott regime. A simple illustration of these two regimes is shown in figure 5.

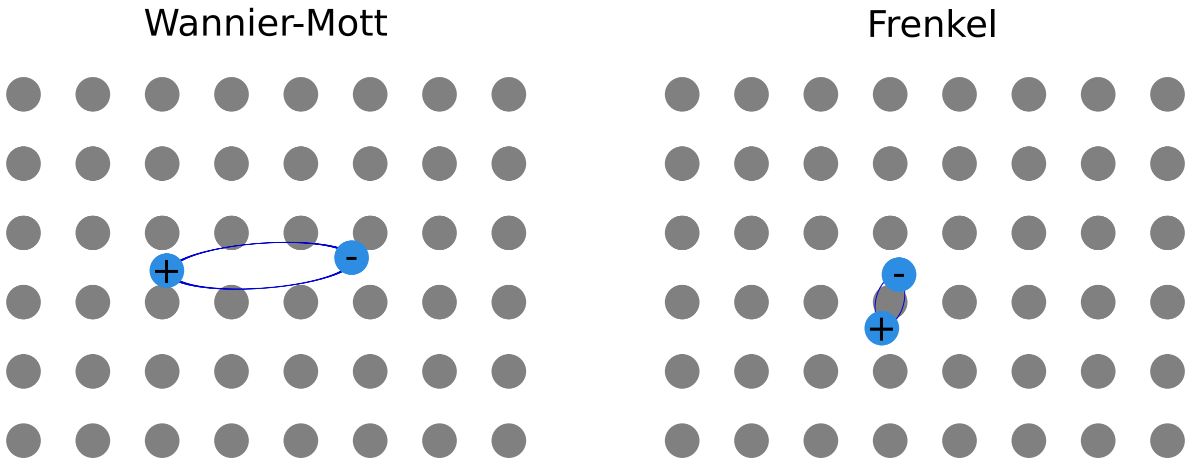

Figure 5: A simple illustration of the two exciton regimes. The background circles are lattice sites, and the excitons are shown as bound charges.

Bulk GaAs has a lattice spacing of about 5.65 Å [2] and its excitons have a radius of 150 Å, and so bulk GaAs has Wannier-Mott excitons. Inside our GaAs/AlAs quantum wells,the physics is effectively two-dimensional, instead of three-dimensional, and so the hydrogenic Bohr atom model must be solved in two dimensions. This leads to a correction of the binding energy, making it bigger by about a factor of two [10], which corresponds to a Bohr radius about half as big, around 75 Å.

The three-dimensional solution for bulk GaAs, using $\epsilon/\epsilon_0 = 13.1$ [73], gives us a binding energy of 3.6 meV; for our quantum wells the binding energy is about 8 meV [10, 14]. For our work, a larger binding energy is better. Because $k_B T$ is about 25 meV at room temperature, our excitons will be torn apart by thermal fluctuations at room temperature, and so we will often be working at cryogenic temperatures where $k_B T$ is around 1 meV. A higher binding energy would allow us to work at room temperature which would be much more practical for various devices we might wish to create. (It is not very convenient to carry a liquid helium canister to cool your cellphone down.) Later on, I will discuss our work in transition metal dichalcogenide monolayers which are similar, but with a much higher binding energy.

As discussed above, excitons may be created when an impulse of energy promotes an electron from the valence band to the conduction band. In theory, this could happen in various ways such as a phonon absorption or even something truly exotic like an alpha particle from a radioactive emitter. In our lab it is in fact quite common to pump a sample with a laser with much higher energy than the band gap. This creates a large number of very hot free electrons and holes, which are initially out of equilibrium. Over time, these particles thermalize through phonon scattering with the underlying lattice, and once cooled, they will begin to coulomb bind to form excitons. Alternatively, we may resonantly pump our samples with a laser tuned to the specific energy required to make excitons, although typically we reserve resonant pumping for exciton-polariton systems, which will be introduced later.

Thus, we have multiple processes to create excitons. However, excitons may also be destroyed. In principle, there are a few avenues for this to occur. For example, a thermal fluctuation could occur with particularly energetic phonon that ionizes the exciton. However, in a system devoid of some sort of engineering to facilitate such processes, exciton decay is dominated by electron-hole recombination. In this recombination, a photon is released with an energy E_γ given as

$$E_\gamma = E_{bg} - \frac{Ry_{ex}}{n^2} + \frac{\hbar^2 k_{COM}^2}{2m_{ex}}, \tag{1.14}$$

where E_{bg} is the band gap energy of the semiconductor, about 1.5 eV for GaAs at cryogenic temperatures. This is also the energy required to resonantly pump the sample to create excitons. Therefore we may actually pump our samples with a laser slightly below the band gap energy and still create excitons. The photons released by recombination exit the material, and may be imaged using a calibrated spectrometer to measure their energy, these photons are called photoluminescence (PL). Its worth mentioning that similar to how the free electron in the conduction band is not the same as a vacuum electron outside of the material, photons inside the material are fundamentally different than vacuum photons. Later, we will more thoroughly introduce the proper terminology for photons inside a material; the polariton.

Figure 6 a photoluminscence (PL) spectrum measured from excitons in GaAs multiple-quantum-well structure. The sample consisted of 12 quantum wells on top of a distributed Bragg reflector (DBR). We will more thoroughly discuss DBRs later, but they are essentially mirrors with very high reflectivity. The 12 different quantum wells are responsible for the multiple peaks in the spectrum. Essentially we have 12 different populations of excitons, with slightly different energies due to the slightly different quantum well thicknesses, which give different zero-point confinement energies. In order to compare their relative intensities, we would need to determine relative density of states and radiative lifetimes, which is not trivial. For this reason, we may not simply fit Equation (1.5) to the entire data set in the left image of Figure 6. However, we may fit the high energy tail of the highest energy peak relatively well. The highest-energy data must be discarded because our signal-to-noise ratio is diminished there. This best fit is shown in the figure on the right. We see a good fit indicating thermalization, and find a temperature of about 15 K. This experiment was conducted in a liquid helium cryostat, measured to be at 4 K. This small discrepancy between our cryostat's temperature reading and our Maxwell-Boltzmann fits is typical. We find that populations of our various quasiparticles are usually warmer than our underlying lattice (which should be accurately measured by our cryostat's temperature sensor).

The lifetime of the excitons is inversely proportional to the rate of their recombination. This recombination is directly related to the overlap of the electron and hole wave functions. As such, we would expect Frenkel excitons to have a shorter radiative lifetime than Wannier-Mott excitons. In GaAs, the excitons have a lifetime of about 1000 ps, but a thermalization time of only 10 ps [22], and so we would expect our excitons to be thermalized when we measure them. This is in good agreement with our Maxwell-Boltzmann fit above.

One final crucial detail of excitons must be mentioned. The free electron is a spin one-half particle and the hole is a spin three-halves particle (at least in our materials). When they combine, the resulting exciton is therefore an integer spin particle, i.e., a boson. As such, the Pauli exclusion principle no longer applies. Two excitons should be able to exist right on top of each other (although the electrons and holes are never at exactly the same place). In fact, many more than two should be possible, i.e., a Bose condensate, which is a concept we will

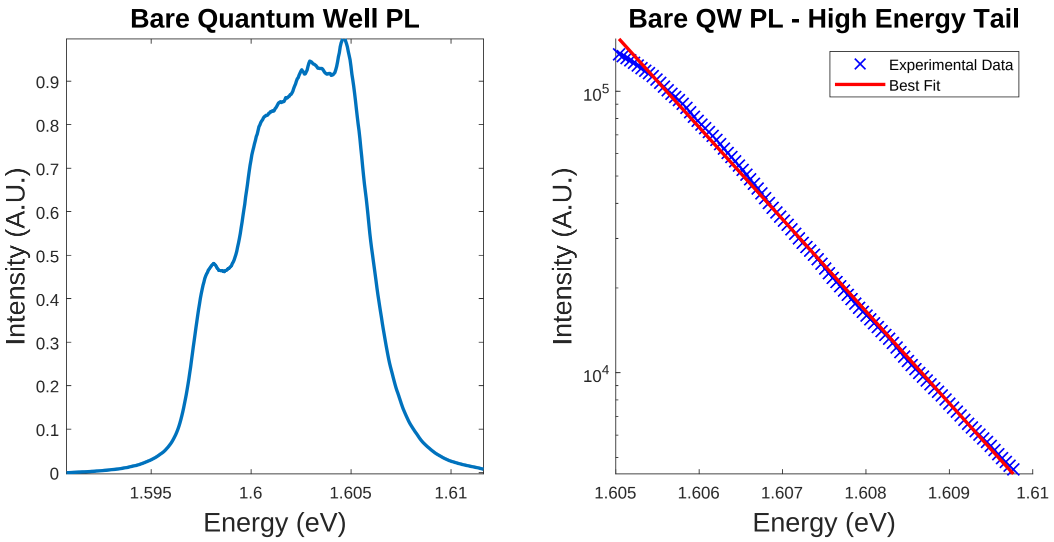

Figure 6: Photoluminescence measured from Pfieffer sample 9-3-14.1 which is a 20 period bottom DBR with 12 exposed quantum wells grown on top. The fit is to the function $I = I_o \exp(-E/K_b T)$

more thoroughly introduce in Section 1.3. If we step back from our renormalized vacuum picture, and remember that the exciton is essentially a free electron orbiting the empty space it used to occupy in the lattice, this is somewhat surprising. Nevertheless, this does seem to be an accurate picture of physics, with significant evidence supporting the renormalized vacuum picture as valid [74] [86].

1.3 Condensation and Critical Temperature

In Section 1.2, we introduced the exciton and concluded that it is a chargeless boson. One of the key features of bosons is their ability to condense into a Bose-Einstein condensate (BEC), colloquially referred to simply as a condensate. Here we will review what that means and derive several of the critical equations which ultimately motivate us to move on from excitons towards polaritons. For a more thorough analysis see Ref. 64 and 73. The dimensionality of the system plays a critical role. Bulk excitons live in a three-dimensional space, while quantum well excitons live in a two-dimensional space. We will first analyse the three-dimensional problem, then briefly discuss the two-dimensional problem.

We begin by considering the Bose-Einstein distribution for bosons, the average occupation number for a state with energy E is

$$\langle n_E \rangle = \frac{1}{e^{(E-\mu)/k_B T} - 1} = \frac{1}{e^{\beta(E-\mu)} - 1}, \tag{1.15}$$

where μ is the chemical potential, and we have introduced the notation $\beta = 1/k_B T$ for convenience. We see that there is a singularity if the energy is equal to the chemical potential, thus the chemical potential must be lower than the lowest energy state in the system. For many applications we have a density of states which we approximate as going to zero at the lowest energy state. As long as the lowest energy state does not have too many particles in it, this is a reasonably good approximation. One such example is blackbody radiation [71]. However, we see that as the temperature decreases, more and more particles will occupy the lowest energy state, and so special care must be taken.

To examine this behavior, we look at the total number of bosons. We have

$$N = \sum_E \frac{1}{e^{\beta(E-\mu)} - 1}. \tag{1.16}$$

Now we wish to convert this sum into an integral. To do so, we must calculate the density of states $D(E)$ given as

$$D(E)dE = \frac{V}{(2\pi)^3} 4\pi \frac{k^2(E)}{|\nabla_{\vec{k}} E|} dE. \tag{1.17}$$

This equation only applies for three dimensional space. The differential dE is included because the density of states only has meaning when multiplied by a differential, and a different density of states is utilized when integrating over other variables such as wavelength. Now we need to insert our parabolic approximation for the dispersion relation of the exciton. However, our excitons have a qualitatively different dispersion relationship, given as

$$E_x = E_{bg} - \frac{Ry_{ex}}{n^2} + \frac{\hbar^2 k^2}{2m_{ex}}, \tag{1.18}$$

where we have dropped the COM subscript from the momentum of the exciton, as we understand we are working in the exciton picture. Our excitons have a minimum energy associated with creating them, unlike photons, which are the first bosons typically taught when learning Bose-Einstein statistics. Electrons, on the other hand, do have a minimum energy associated with creating them, some of which is the rest mass energy; however typically we have number conservation when working with electrons, and so we do not need to worry about the energy of creating and destroying electrons. The exciton has characteristics reminiscent of both electrons and photons. Because we are going to ultimately be interested in low-temperature behavior, we may restrict ourselves to only having excitons in the $n = 1$ Rydberg state. This is valid because the separation between the $n = 1$ and $n = 2$ excitons is large compared to cryogenic k_BT. Additionally, to make our math cleaner, we may define the $E = 0$ point to be at $E_{bg} + Ry_{ex}$. This is valid because at the heart of all statistical physics derivations is maxwell-boltzmann exponential term to weight probabilities, which ultimately only depends on relative energy differences not absolute values. So then, our exciton energy may be simply written as

$$E_x = \frac{\hbar^2 k^2}{2m_{ex}}. \tag{1.19}$$

Utilizing equation 1.17 we have

$$|\vec{\nabla}_{\vec{k}} E| = |\frac{\hbar^2}{2m_{ex}}(2k_x\hat{k_x} + 2k_y\hat{k_y} + 2k_z\hat{k_z})| = \frac{\hbar^2 k}{m_{ex}} \tag{1.20}$$

and

$$D(E)dE = \frac{V}{(2\pi)^3} 4\pi \frac{m_{ex}k(E)}{\hbar^2} dE. \tag{1.21}$$

The $k(E)$ term may be obtained by rearranging equation 1.18. We may now use this density of states to convert our discreet sum into an integral. However, we see that our density-of-states approximation assigns a weight of zero to the state at $E = 0$, so we manually add this state's contribution outside of the integral:

$$N = \frac{1}{e^{-\beta\mu} - 1} + \int_{E=0}^{\infty} \frac{V}{(2\pi)^3} 4\pi \frac{m_{ex}}{\hbar^2} \sqrt{E} \sqrt{\frac{2m_{ex}}{\hbar^2}} \frac{1}{e^{\beta(E-\mu)} - 1} dE. \tag{1.22}$$

The lower bound of the integral is not $-\infty$ because no exciton can exist below $E_{bg} - Ry_{ex}$ which we've defined as the zero point for energy. The term before the integral is the number of particles in the ground state, which we denote as N_o. We rearrange the equation as follows:

$$\frac{N}{V} = \frac{N_o}{V} + \frac{2\pi}{h^3}(2m)^{3/2} \int_{E=0}^{\infty} \frac{\sqrt{E}}{e^{\beta(E-\mu)} - 1} dE, \tag{1.23}$$

i.e.,

$$\frac{N - N_o}{V} = \frac{2\pi}{h^3}(2m)^{3/2} \int_{E=0}^{\infty} \frac{\sqrt{E}}{e^{\beta(E-\mu)} - 1} dE. \tag{1.24}$$

At this point it is prudent to define the fugacity $z = e^{\mu/k_B T}$, and use one final U-substitution to remove the β from our integral, leaving us with

$$\frac{N - N_o}{V} = \frac{2\pi}{h^3}(2mk_B T)^{3/2} \int_{x=0}^{\infty} \frac{\sqrt{x}}{z^{-1}e^{x} - 1} dx. \tag{1.25}$$

The integral is the well known Bose-Einstein integral. For a thorough analysis of it see Appendix D of Ref. 64. The generic form of the integral is

$$g_\nu(z) = \frac{1}{\Gamma(\nu)} \int_0^{\infty} \frac{x^{\nu-1}}{z^{-1}e^{x} - 1}, \tag{1.26}$$

where Γ is the standard gamma function. An essential property of this integral which we will utilize is that when $z = 1$ there is a simple relationship to the Riemann zeta function; $g_\nu(1) = \zeta(\nu)$. This property is useful as it allows us to quickly calculate some limiting behavior, as we shall see. Substituting this form into our prior result leaves us with

$$\frac{N - N_o}{V} = \frac{2\pi}{h^3}(2mk_B T)^{3/2} g_{3/2}(z). \tag{1.27}$$

We now consider the fugacity defined above. As we stated at the beginning, the chemical potential must be less than the lowest energy state for the excitons. Recalling that we've set our energy's zero point to be that lowest state, we can say that since $\mu < 0$, this implies that $z < 1$. The Bose-Einstein integral is monotonically increasing with z, and so the Riemann zeta function can provide us an upper bound as follows:

$$\frac{N - N_o}{V} = \frac{2\pi}{h^3}(2mk_B T)^{3/2} g_{3/2}(z) < \frac{2\pi}{h^3}(2mk_B T)^{3/2} \zeta(3/2) = n_{e,max} \tag{1.28}$$

The quantity $(N - N_o)/V$ is the density of excitons in excited states. We see that this quantity is bounded, and we name that upper bound $n_{e,max}$. So then, what happens if the density of excitons $n = N/V$ is greater than this bound? The only state available to them is the ground state, and so we can get macroscopic occupation of a single quantum state. This ground state is then named the condensate. Physically, macroscopic occupation of a single quantum state means a macroscopic fraction of the particles have coherence of their wave functions.

So then, the condition for condensation is

$$n > \frac{2\pi}{h^3}(2mk_BT)^{3/2}\zeta(3/2), \tag{1.29}$$

which we may rearrange as

$$\frac{n^{2/3}}{\zeta(3/2)^{2/3}}\frac{h^2}{(2\pi)^{2/3}}\frac{1}{2mk_B} > T. \tag{1.30}$$

The left hand side is the critical temperature, below which condensation will occur. This is the essential result we have been working towards. In particular, we see that the critical temperature is inversely related to the mass. Heavier particles need to be cooler in order to condense.

Finally, what happens if we try to repeat the above derivation for a two dimensional boson gas?. In two dimensions our density of states changes to

$$D(E)dE = \frac{A}{(2\pi)^2}\frac{2\pi k(E)}{|\nabla_{\vec{k}}E|}dE = \frac{A}{(2\pi)^2}\frac{4m_{ex}\pi k(E)}{\hbar^2 k} = \frac{4Am_{ex}\pi}{h^2}. \tag{1.31}$$

We see a crucial difference which is that our density of states does not depend on energy. All states are weighted equally. This means our justification for pulling the ground state out of the integral is gone. Furthermore, if we were to follow the same steps we would eventually arrive at the integral

$$N = \frac{4Am_{ex}\pi k_BT}{h^2}\int_{x=0}^{\infty}\frac{1}{z^{-1}e^x - 1}dx. \tag{1.32}$$

Unlike our previous integral in the three dimensional case, this one is not bounded when $z = 1$. We can see clearly the numerator does not go to zero. In fact, the Bose-Einstein integrals are specifically defined only for $\nu > 1$ in the case that $z = 1$ [64]. Therefore, there is not a macroscopic number of particles in any one quantum state in 2D. However, the number of particles in any one quantum state can be much larger than 1, in a 2D system, which leads, among other things, to a macroscopic but not infinite coherence length [3]. There is a rich history regarding the exact nature of two-dimensional condensates; for an overview please see Ref. 40. We will summarize that the coherence length of a two- dimensional condensate cannot extend to infinity, but it can in many cases be much larger than the size of the experimental system.

1.4 The Cavity Photon

We now turn our attention to light. The equations which describe light are of course Maxwell's equations, and for a thorough treatment on the subject every physicist knows to refer to Ref. 42. Light is of course the fastest thing in the universe, and not only is it fast, but its always moving. So, if we want to be able to work with light in any sort of device application, we need a method of containing it. Our method of choice is to implement a microcavity. A microcavity is essentially two highly reflective mirrors facing one another with a very small distance in between them. Light may bounce back and forth between these two mirrors until it manages to escape the cavity. The photon can escape because the mirrors are not 100 % reflective, and so every bounce there is a small probability of escape either through transmission through the mirror or absorption into it.

The very small thickness between the two mirrors of a microcavity is advantageous because light inside the microcavity will form a standing wave, which means that we can accurately predict precisely where inside the microcavity the amplitude of the electric and magnetic field will be maximal and minimal. Typically, a microcavity is designed with a specific wavelength of light in mind. For example, let's say we wished to designed a microcavity to produce a standing wave of 770 nm light with 3 anti-nodes inside of it. In this case, we mean 770 nm as measured in vacuum. The microcavity could have a vacuum in between the two mirrors, or we could have a material with an optical index. For this example, lets consider a microcavity in which the cavity is filled with glass, which we will take to have an optical index of 1.4 for this example.

There are two ways of viewing the problem. One way is to calculate the wavelength of light in our materials of interest and work with the real path lengths. The other is to work with the vacuum wavelength and calculate optical path lengths. We generally find the method of optical path lengths to be easier to implement when designing more complicated structures.

In order to create our standing wave, we want the optical path length inside the cavity to be equal to $(3/2) * 770$ nm $= 1155$ nm; this is because the mirrors will form the nodes of our standing wave, and so each half-wavelength inside our cavity will contain one antinode. The optical path length is simply the real path length multiplied by the real part of the optical index. For now, we are only considering a purely real optical index. So we would have $1.4L = 1155$ nm, and ultimately end at $L = 825$ nm.

The general equation for designing a microcavity's width is then given by

$$w = \frac{f\lambda_{design}}{n}. \tag{1.33}$$

where f is our design factor, in the prior example $3/2$, λ_{design} is the design wavelength as measured by light in vacuum after it has exited the cavity, and n is the real part of the optical index of the material filling the cavity. It should be noted that in general the optical index is a function of wavelength, and so care must be taken to use the optical index specifically at the design wavelength.

Figure 7 shows a simple illustration of the microcavity we designed above. We see that in the electromagnetic standing wave, both the forwards- and backwards-propagating waves exist. This deliberate design tells us the electric field will be maximized at the three antinodes; this will be beneficial to us later. Typically, we lay these samples flat, so we think of one of the mirrors as the top mirror and the other as the bottom; by convention we tend to draw the bottom mirror on the left. We take the direction of propagation to be the z-axis, also named the perpendicular direction, as denoted by a subscript $\perp$. As we have drawn it, the wave has momentum only in the z-direction. However, we may also have light trapped in our microcavity with momentum in the other two orthogonal directions. Because

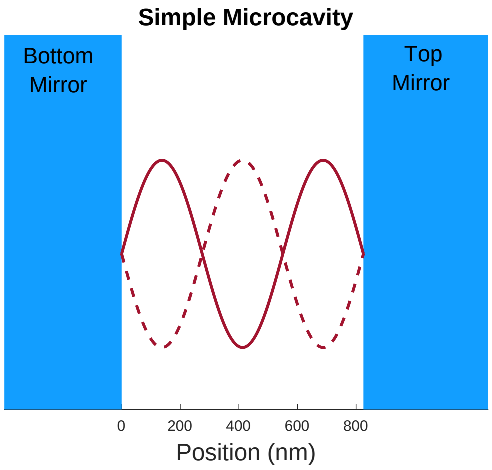

Figure 7: An illustration of the standing wave formed inside of a microcavity.

of the rotational symmetry we tend to combine these two directions as the in-plane direction denoted with the subscript $\|$. So then, our total momentum vector's magnitude will be

$$|\vec{k}| = \sqrt{k_{\|}^2 + k_{\perp}^2}. \tag{1.34}$$

The confinement in the perpendicular direction restricts our $k_{\perp}$ values. There are multiple standing waves possible, we recall that the general form for a standing wave with nodes at both ends is $\lambda_m = 2W/m$, with $m = 1$ being the lowest-energy mode, and the general conversion between wavelength and wave vector is $k = 2\pi/\lambda$. This leaves us with the following possible $k_{\perp}$ values:

$$k_{\perp,m} = \frac{m\pi}{w} = mk_{\perp}, \tag{1.35}$$

where we have defined the first mode's wave number to be simply $k_{\perp}$, which we may now substitute into our total momentum vector's magnitude as

$$|\vec{k}| = \sqrt{m^2 k_{\perp}^2 + k_{\|}^2}. \tag{1.36}$$

We may use this to calculate the energy of the photons using the simple relationship for photons in a material $E = \hbar ck/n$ resulting in

$$E_m = \frac{\hbar c}{n}\sqrt{m^2 k_\perp^2 + k_\parallel^2} = \frac{\hbar c m k_\perp}{n}\sqrt{1 + \left(\frac{k_\parallel^2}{m^2 k_\perp^2}\right)} \tag{1.37}$$

We see now that an important change has occurred by confining our photons to a microcavity. In vacuum, photons may have an arbitrarily small energy, but in a microcavity the minimal energy occurs for the $m = 1$ mode with no in-plane momentum. We denote this minimal energy as $E_o = \hbar c k_\perp / n$, and the minimal energy of the m'th mode of the cavity is then mE_o. Our dispersion relation may now be written more cleanly as

$$E_m = mE_o\sqrt{1 + \frac{\hbar^2 c^2}{m^2 n^2 E_o^2} k_\parallel^2}. \tag{1.38}$$

Next, we recall our definition for the effective mass requires us to calculate the second derivative. However, it is perhaps simpler to apply the binomial approximation to our square root, which we know will successfully approximate our energy as parabolic. We are interested in calculating the effective mass at $k_\parallel = 0$, so the binomial approximation is warranted, giving us

$$E_m = mE_o + \frac{\hbar^2 c^2}{2mn^2 E_o} k_\parallel^2 = mE_o + \frac{\hbar^2 k_\parallel^2}{2m_{cav,m}}. \tag{1.39}$$

Here we have solved for the effective mass of the m'th cavity mode as

$$m_{cav,m} = \frac{mn^2 E_o}{c^2} = \frac{n^2 E_m}{c^2}. \tag{1.40}$$

Notice the division by c^2, which leads us to immediately expect very low effective mass. Let us continue our example from above. The wavelength of 770 nm implies $E_3 = 1.6102 eV$, which leads to a mass of $6x10^{-6} m_e$ where m_e is the vacuum mass of an electron. If instead of the index of glass we had used AlAs we would have ended with a mass of approximately $10^{-5} m_e$. Compare this to the mass of GaAs excitons which are about $0.1 m_e$ [22], and we can see that while confinement to the cavity has imbued our photons with some mass, it is still a very light mass. This can be very beneficial, since as we discussed in Section 1.3, the critical temperature for condensation goes inversely with the mass of the particle. Therefore, we expect cavity photons to more easily condense; however, there are some problems. In general, photons do not interact with the lattice as strongly as excitons, and so thermalization may not occur. Additionally, when slightly imperfect mirrors are used, the lifetime of the photons is very short, on the order of a few picoseconds for great mirrors, and 100 picoseconds for exceptional mirrors [76]. If the photons don't live long enough to thermalize, they won't condense. As we will discuss in Section 1.5, these issues are addressed by introducing exciton coupling.

Next, we want to reformulate our dispersion relationship in terms of angle rather than in-plane momentum. The energy and momentum variables are certainly the best variables for the formulation of the theory of physics. However, the devices and setups in our lab, which we will discuss in depth in Chapter 2, tend to measure angle and wavelength; we also discuss this issue more thoroughly in that chapter. For this reason, a dispersion relationship in the form $E(\theta)$ is most convenient. Rather than attempt to manipulate our above results, its easier to start over from scratch.

Inside of our cavity we consider the lowest energy standing wave with zero in-plane momentum, which will have a momentum vector with magnitude $k_\perp$, and then we consider a wave instead at an angle θ which will have a momentum vector with magnitude $k_\perp$. This second wave isn't exactly a standing wave, as it is traveling along the microcavity. However, the perpendicular component of its momentum vector is still $k_\perp$. This can be visualized in Figure 8.

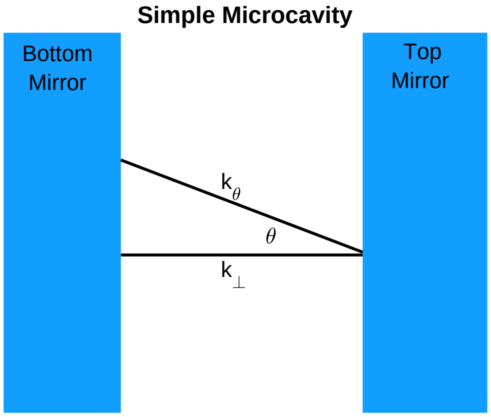

Figure 8: A standing wave inside our microcavity at non-zero angle.

Because the energy of a photon is directly proportional to the magnitude of its momentum vector, we may write

$$\cos\theta = \frac{k_\perp}{k_\theta} = \frac{E(0)}{E(\theta)}. \tag{1.41}$$

We ultimately want to solve this equation for the dispersion curve we will measure through experiments. However, currently it is expressed in terms of the angle of light inside the cavity. Our detector is outside of the cavity, so we must use Snell's law to obtain energy as a function of emission angle $E(\theta_e)$. Taking the index of air to be unity, we have the dispersion curve in terms of the cavity index n_c as

$$E(\theta_e) = \frac{E(0)}{\cos\theta} = \frac{E(0)}{\sqrt{1-\sin^2\theta}} = \frac{E(0)}{\sqrt{1-\left(\frac{\sin\theta_e}{n_c}\right)^2}}. \tag{1.42}$$

So then, the dispersion curve may be written in terms of two parameters, the zero point energy and the cavity index.

As we mentioned, the data we measure comes to us in the form $\lambda(\theta_e)$, but there are times we must insist on converting to energy and in-plane momentum. Converting from wavelength to energy is easy enough for all data points in that curve utilizing $E = hc/\lambda$, as hc is simply a constant 1239.8 eV nm. In order to obtain the in-plane momentum, we need to remember that Snell's law is the statement that at a boundary, refraction preserves the parallel component of momentum. So we may calculate the parallel component of momentum for the light we measure outside of the cavity, and the result will be equal to the in-plane momentum inside the cavity. The calculation is as follows:

$$\sin\theta_e = \frac{k_\parallel}{|\vec{k}|} = \frac{\hbar c k_\parallel}{E} \tag{1.43}$$

$$k_\parallel = \frac{E\sin\theta_e}{\hbar c} = \frac{\sin\theta_e}{2\pi\lambda}. \tag{1.44}$$

So, our measurements of $\lambda(\theta_e)$ may be converted into the typical dispersion curve $E(k_\parallel)$. Next, let us visualize these results. We create the dispersion curves for a cavity with a maximal wavelength of 775 nm and a cavity index of 3; the results are shown in Figure 9. Additionally, we show the parabolic approximation from above using the cavity mass calculation from equation 1.40. We see that the parabolic approximation is very good for the cavity mode. The results are shown for an emission angle of ±90 degrees, which corresponds to an internal angle of ± 17 degrees. However, in practice our optics only collect ±50 degrees, which corresponds to an internal angle of ±5 degrees.

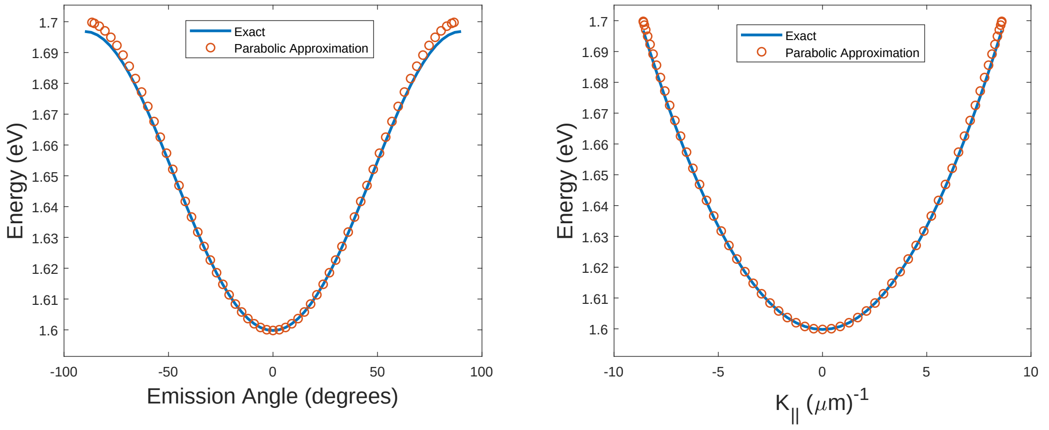

Figure 9: The dispersion relation for a cavity mode, shown in both of our variable systems. The parabolic approximation is also shown for comparison.

An important feature to notice is that both the exact curve and the approximation curve have an inflection point when plotted as a function of emission angle. This is important because it shows that when measured in the lab (emission angle), a perfect parabolic dispersion will still flatten out at high angles. This is an artifact of optical measurements. As we will see in Section 1.5, this flattening out is indicative of exciton-photon coupling when observed in k-space. So, data taken in the lab can be deceiving, where we observe this flattening out, but in k-space no actual flattening is occurring. This can especially be a problem in weak-coupling regimes where our data is not the perfect crisp lines of theory, but rather blurry pixels with noise.

1.5 The Polariton

We've seen that excitons are interacting bosons, which thermalize well because they easily interact with the underlying lattice through phonon exchange. However, condensation of excitons has proven to be troublesome, in part because their mass implies a condensation temperature of 1 K or less at typical exciton densities, which is below typical disorder energies in quantum wells. We've also explored cavity photons, which are bosons imbued with a very light mass. However, cavity photons have very short lifetimes and do not easily thermalize due to poor interaction with the phonons of the underlying lattice. We seek to combine the properties of these two particles, and so we turn our attention now to the polariton.

Broadly speaking, polaritons are a renormalization which comes from the coupling of light with some other excitation. For interesting phenomenon to occur, this coupling needs to be strong enough that a simple quantum mechanical perturbation approach begins to break down, and instead new eigenstates need to be calculated from the whole Hamiltonian; these new eigenstates are the polaritons. There are many different types of polaritons such as surface-plasmon-polaritons [90], phonon-polaritons [9], and Cooper-pair polaritons [15]. For our work, we will be focusing exclusively on exciton-polaritons, specifically inside a microcavity.

Any time a photon is inside a material, it is being renormalized in some way, as there is always some other resonance for it to couple to, such as the nearest electron band. However, without specifically designing a system, the difference between this new polariton state and the regular photon is minimal. In order to achieve a larger change in our photons, we put GaAs quantum wells at the antinode of a standing wave inside a microcavity. A simple illustration of a $\lambda/2$ microcavity with a quantum well at the anti-node is shown in Figure 10.

In this design, the quantum well will be able to readily absorb photons from the microcavity and convert them to excitons, and exciton recombination will result in a photon reentering the cavity. So then, our Hamiltonian needs exchange terms to represent that. Additionally we have the terms representing the energy in the cavity photons and excitons, utilizing the dispersion relationships for the exciton and photon, we have

$$\hat{H} = \sum_{k_\parallel} E_c(k_\parallel)\hat{a}^\dagger_{k_\parallel}\hat{a}_{k_\parallel} + \sum_{k_\parallel} E_{ex}(k_\parallel)\hat{b}^\dagger_{k_\parallel}\hat{b}_{k_\parallel} + \sum_{k_\parallel} g(\hat{a}^\dagger_{k_\parallel}\hat{b}_{k_\parallel} + \hat{b}^\dagger_{k_\parallel}\hat{a}_{k_\parallel}), \tag{1.45}$$

where g is the exciton-photon coupling strength. In principle g can depend on angle [46], however in practice this correction is quite small. Furthermore, in Chapter 4 we will show

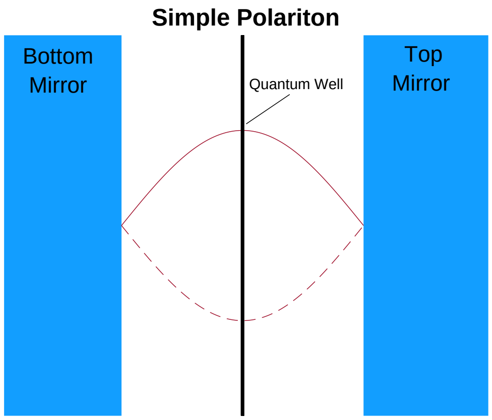

Figure 10: A schematic of a basic microcavity polariton sample.

some of the failings of this simple two-level model. For now, we use it as a simple introduction to polaritons. Because we do not have any coupling between different k-states, each k-state forms a simple two-by-two subspace, which may be diagonalized independent of the other subspaces. So, we could write the Hamiltonian of this subspace as

$$\hat{H}_{k_\parallel} = \begin{pmatrix} E_{ex}(k_\parallel) & g \\ g & E_c(k_\parallel) \end{pmatrix}. \tag{1.46}$$

We diagonalize this matrix to find our new eigen states, written as

$$\hat{L}_{k_\parallel} = X_{k_\parallel}\hat{b}_{k_\parallel} + C_{k_\parallel}\hat{a}_{k_\parallel} \tag{1.47}$$

and

$$\hat{U}_{k_\parallel} = -C_{k_\parallel}\hat{b}_{k_\parallel} + X_{k_\parallel}\hat{a}_{k_\parallel}. \tag{1.48}$$

The X and C are the Hopfield coefficients [22], solved as

$$|X_{k_\parallel}|^2 = \frac{1}{2}\left(1 - \frac{E_{ex}(k_\parallel) - E_c(k_\parallel)}{\sqrt{4g^2 + \left(E_{ex}(k_\parallel) - E_c(k_\parallel)\right)^2}}\right) \tag{1.49}$$

and

$$|C_{k_\parallel}|^2 = \frac{1}{2}\left(1 + \frac{E_{ex}(k_\parallel) - E_c(k_\parallel)}{\sqrt{4g^2 + \left(E_{ex}(k_\parallel) - E_c(k_\parallel)\right)^2}}\right). \tag{1.50}$$

The corresponding eigenenergies are

$$E_{lp}(k_\parallel) = \frac{1}{2}\left(E_{ex}(k_\parallel) + E_c(k_\parallel) - \sqrt{4g^2 + \big(E_{ex}(k_\parallel) - E_c(k_\parallel)\big)^2} \right) \tag{1.51}$$

and

$$E_{up}(k_\parallel) = \frac{1}{2}\left(E_{ex}(k_\parallel) + E_c(k_\parallel) + \sqrt{4g^2 + \big(E_{ex}(k_\parallel) - E_c(k_\parallel)\big)^2} \right). \tag{1.52}$$

Lastly we may write our now diagonalized Hamiltonian in terms of these results as:

$$\hat{H} = \sum_{k_\parallel} E_{lp}(k_\parallel) \hat{L}^\dagger_{k_\parallel} \hat{L}_{k_\parallel} + \sum_{k_\parallel} E_{up}(k_\parallel) \hat{U}^\dagger_{k_\parallel} \hat{U}_{k_\parallel}. \tag{1.53}$$

The energy $E_{lp}(k_\parallel)$ and the operator $\hat{L}$ correspond to the lower energy state, named the lower polariton, and the other state is named the upper polariton. The squared magnitude of the Hopfield coefficients give us the exciton and cavity fractions. This describes what percentage exciton and cavity photon the new particles are. As written, the exciton fraction of the lower polariton at in-plane momentum $k_\parallel$ is $|X_{k_\parallel}|^2$, while the exciton fraction of the upper polariton is $|C_{k_\parallel}|^2$. Conversely, the cavity fraction of the lower polariton at in-plane momentum $k_\parallel$ is $|C_{k_\parallel}|^2$, while for the upper polariton the cavity fraction is $|X_{k_\parallel}|^2$. Our new states are normalized, and so we have $|C_{k_\parallel}|^2 + |X_{k_\parallel}|^2 = 1$.

An interesting result to notice is that the exciton fraction of the lower polariton may be calculated by taking the derivative of the lower polariton's dispersion curve with respect to the energy of the exciton:

$$|X_{k_\parallel}|^2 = \left.\frac{dE_{lp}(k_\parallel)}{dE_{ex}}\right|_{k_\parallel} \tag{1.54}$$

In chapter 4 we will take this as a reasonable generalization of the exciton fraction for more complicated models.

The quantity $\Delta = E_{ex}(k_\parallel) - E_c(k_\parallel)$ is named the detuning; the value at $k_\parallel = 0$ is also commonly used to characterize a microcavity structure. In the literature, if one sees a reported detuning value, one should assume it is at $k_\parallel = 0$ unless otherwise stated. Going forward, we will use the name detuning exclusively to refer to the $k_\parallel = 0$ value. Similarly, the exciton fraction is also dependent on $k_\parallel$, but colloquially the name refers specifically to the $k_\parallel = 0$ value $|X_o|^2$.

By varying the detuning, we may greatly modify the behavior of the upper and lower polariton. For demonstration, we will approximate the mass of the exciton as infinite, i.e., a flat dispersion curve. This is a very reasonable approximation because the mass of the photon is four to five orders of magnitude lower than the mass of the exciton. The dispersion curves for three different detunings are shown in Figure 11, along with the exciton fractions of the polaritons. In general, we see the lower polariton has the highest cavity-photon character (that is, is the most "photonic") at $k_{\|} = 0$ and is most exciton-like (excitonic) at higher angles.

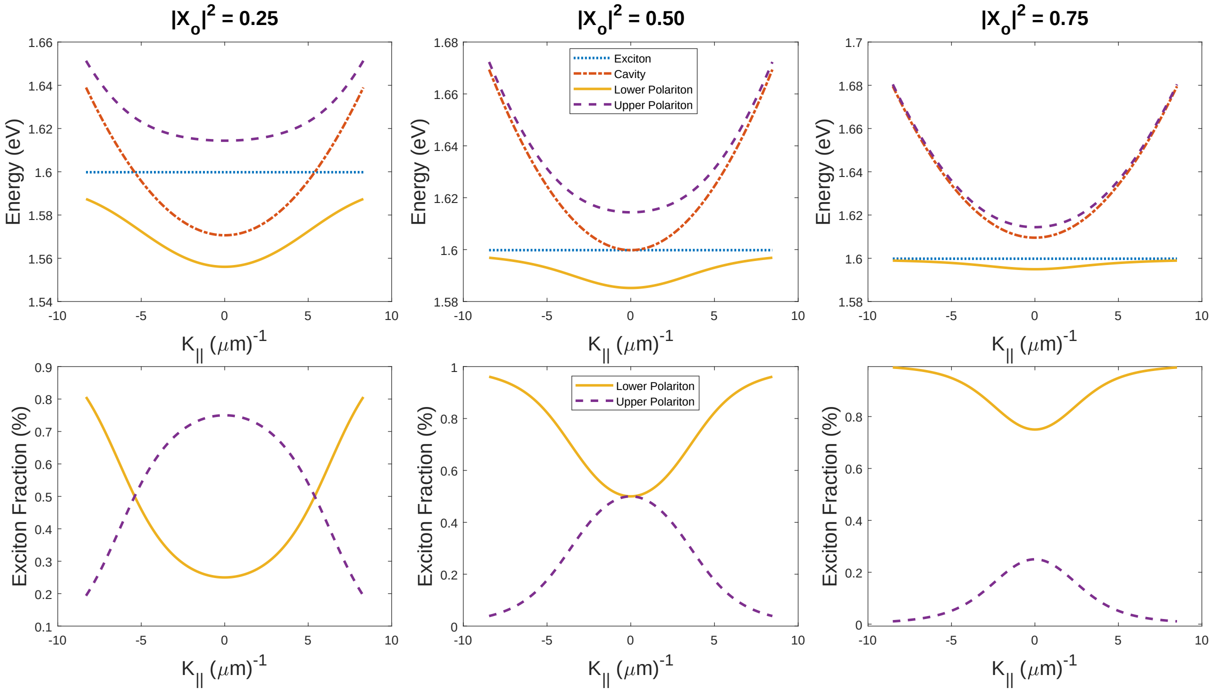

Figure 11: The dispersion curves of polaritons at three different detunings.

Recalling our definition for the effective mass,

$$m_{\text{eff}} = \frac{\hbar^2}{2\frac{d^2E}{dk^2}}, \tag{1.55}$$

it may be shown that the effective mass of the lower polariton at $k_{\|} = 0$ in terms of the mass of the cavity photon and exciton is given by

$$\frac{1}{m_l p} = \frac{|X_o|^2}{m_{ex}} + \frac{|C_o|^2}{m_{cav}}. \tag{1.56}$$

This equation is strictly for determining the effective mass of the lower polaritons at $k_{\|} = 0$. For other momentum values, Equation (1.55) must be used. Looking at our dispersion curves in Figure 11, we notice something surprising. The lower polariton has an inflection point, which means our effective mass would be undefined, and it isn't as though this is a single point that we could easily remove or ignore, as we get closer to the inflection point the mass will still approach gigantic values. Indeed, this is a real phenomenon; the vicinity around the infliction point is known as the bottleneck region , and polaritons can become trapped in this region when thermalizing due to their inability to emit phonons efficiently [84][58].

Examining Equation (1.56), and recalling that the mass of an exciton is about $10^{-1}m_e$ while the cavity mass is $10^{-5}m_e$, we see that polaritons which are close to evenly mixed will have very low effective mass, predominantly determined by the cavity photon mass. This would hint that we should be able to condense them at relatively high temperatures.

The lifetime of the polaritons may also be predicted in a similar manner, but it requires us to add complexity to our model. We must add the lifetime of the exciton and cavity to their energy in the form of an imaginary self-energy. Then, the resulting dispersion curves for the upper and lower polaritons will be complex-valued; however they will have the same exact functional form (with the understanding all the variables are complex valued).

The lifetime of the polaritons τ_{lp} may then be expressed in terms of the exciton lifetime τ_{ex} and the cavity lifetime τ_{cav} as

$$\frac{1}{\tau_{lp}} = \frac{|X_o|^2}{\tau_{ex}} + \frac{|C_o|^2}{\tau_{cav}} \tag{1.57}$$

Recalling that the lifetime of excitons are around 1000 ps and the cavity is significantly shorter, we see that the lifetime of our polaritons will be predominantly determined by the lifetime of photons escaping the cavity. As we will discuss in Chapter 4, our collaborators are able to grow very high quality samples for us, where the cavity lifetime is around 120 ps [76]. The thermalization time of excitons is around a few picoseconds. So, we may hope that our polaritons inherit a thermalization time shorter than their lifetime, and indeed they often do.

The lifetime is inversely related to the energy linewidth ΔE. Polariton data is also commonly measured in wavelength λ, and so there is also a corresponding wavelength linewidth

$\Delta\lambda$. There is additionally a quantity termed the Q-factor, which is the linewidth defined by the central value. In some literature it is defined in terms of wavelength as $\Delta\lambda/\lambda_c$ [22] and in others it is defined as $\Delta E/E_c$ [47]. Ultimately these two definitions are not too different. The Q-factor is how many round trips the photon completes on average before escaping the cavity. So, long-lifetime samples are high-Q samples, and short-lifetime samples are low-Q samples.

Figure 12 shows some PL data measured from one of our long-lifetime samples. At each angle we extract the intensity along the main dispersion curve to get intensity as a function of energy. Then, in the right image we fit that data to confirm that our polaritons are thermalized. The fit returns a temperature of 10 K, fairly close to the cryostat temperature which was nominally 4 K.

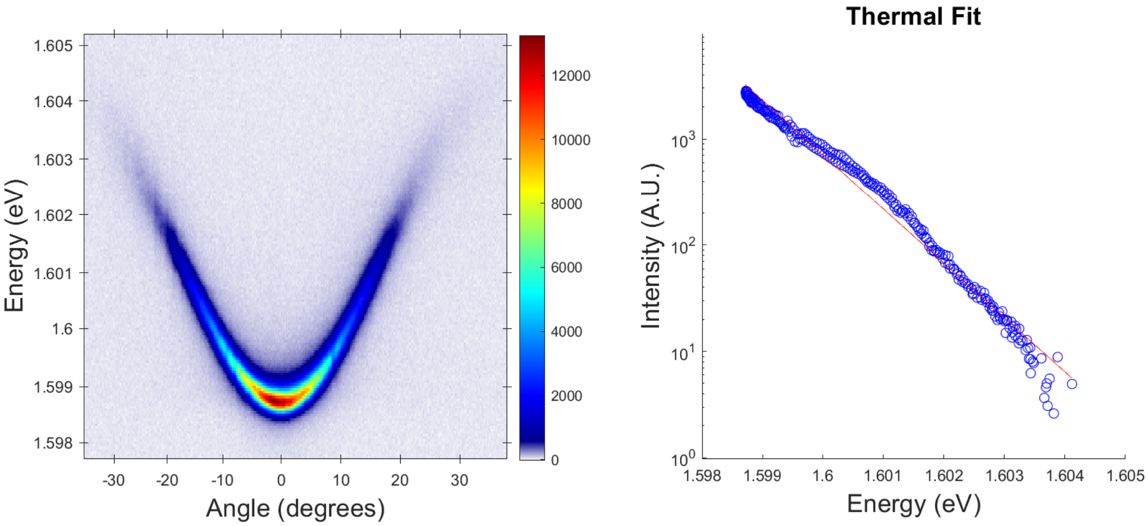

Figure 12: Photoluminescence measured from one of our long-lifetime samples. In the right image we fit the intensity with a Maxwell-Boltzmann distribution to confirm thermalization.

Given these results, it is not too surprising that condensation is regularly observed in microcavity exciton-polariton systems [7, 57, 88], along with many other interesting effects [17, 3, 24, 23, 68, 69, 21, 70, 25, 54, 65, 53].

1.6 Sample Design

We now wish to more concretely describe the creation of our samples in which we observe polaritons. So far, we have described a sample which consists of two high-reflectivity mirrors creating an electromagnetic standing wave in between them, with a quantum well placed at the antinode of the standing wave. There are several ways to realize this, and there are two categories of our work. There are GaAs quantum well samples, and then there are also transition metal dichalcogenide monolayer (TMD) samples. For this section, we will focus solely on the GaAs work, and then the modifications necessary for our TMD work will be discussed in Chapter 6.

Our GaAs samples are grown by our collaborators, the Loren Pfeiffer group at Princeton University and the Zbig Wasilewski group at the University of Waterloo. Both groups utilize a molecular beam epitaxy (MBE) machine to grow these samples. MBE machines are able to grow crystalline structures with atomic precision, including extremely low surface roughness [45, 31], during the growth reflection high-energy electron diffraction (RHEED) measurements are performed to monitor the sample's thickness and roughness [85].

As we will discuss, there are three different materials that ultimately compose our samples, GaAs, AlAs, and a mixed alloy of the two $\mathrm{Al}_x\mathrm{Ga}_{1-x}\mathrm{As}$, where for our work the alloy fraction is $x = 0.2$. In Table 2 essential parameters of these three materials are summarized [1][37]. There are several important features to discuss. The lattice constant of all three materials is very close. This means that MBE growth of one material on the other will work very well, and ultimately samples with low stress at the interfaces will be created. Additionally, AlAs has a band gap significantly larger than GaAs, so it will work well as the sandwiching material to create GaAs quantum wells.

In order to create the necessary mirrors, we create a distributed Bragg reflector (DBR). A DBR is created through alternating layers of high index and low index material, with thicknesses chosen so as to produce destructive interference in transmitted waves, resulting in near perfect reflection of electromagnetic waves. The key to this is designing each layer to have an optical path length of $\lambda_{design}/4$, where the design wavelength is the wavelength for which this mirror will be the most reflective. As we will see, these mirrors will actually work

Table 2: A table of the essential material parameters for our AlGaAs/GaAs structures. All three materials are a cubic zincblende structure. AlGaAs transitions from direct band gap to indirect band gap at x = 0.4.

	$GaAs$	$AlAs$	$Al_{0.2}Ga_{0.8}As$
Lattice Constant (Å)	5.6533	5.661	5.6535
Thermal Expansion Coefficient ($10^{-6}K$)	6.4	5.2	6.16
Band gap Type	Direct	Indirect	Direct
Band Gap (eV) (Room Temp)	1.424	2.168	1.6734
Band Gap (eV) (4 k)	1.518	2.2634	1.7679
Band Gap (nm) (Room Temp)	870	571	741
Band Gap (nm) (4 k)	817	548	700
Optical Index (at 775 nm, 4 k)	3.40	2.90	3.30

very well over a broad range of wavelengths, but the design wavelength will be the center of that range. A simple schematic of a DBR is shown in Figure 13 A.

This figure shows an incident ray of light, and some of the many reflected rays from the interfaces; calculated the resulting phase change for each is shown. Each layer has an optical path length of $\lambda/2$, which adds a total phase of π, as the ray has to travel forward and backward through each layer. Additionally, this structure utilizes the fact that when a ray of light is reflected by a higher index material than it is currently in, it gains a phase of π (and no phase is added when reflected off a lower index material). We have taken some artistic liberties in this schematic diagram. The rays are drawn at an angle to be comprehensible, but we are treating our rays of light as at normal incidence. We see all the rays are in phase with one another, leading to constructive interference in the reflected beam.

There can also be more complicated ray paths, involving multiple reflections inside of the DBR. To truly analyse this system, Maxwell's equations must be solved self consistently. The preferred method of doing so is the transfer-matrix method (TMM), which we thoroughly

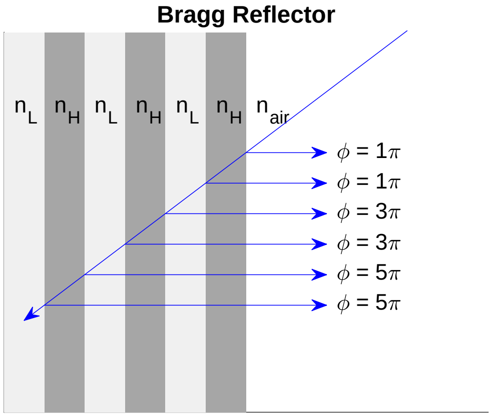

Figure 13: Basic working principle of a distributed Bragg reflector. All reflections are in phase, leading to constructive interference.

discuss in Chapter 3. When characterizing DBRs it is common to talk about the reflectivity. Here, we mean the power reflectivity, calculated as

$$R = \frac{\vec{E_{out}} \cdot \vec{E_{out}^*}}{\vec{E_{in}} \cdot \vec{E_{in}^*}}, \tag{1.58}$$

not to be confused with the reflection coefficient r, which is the complex ratio of the electric fields. The reflectivity curves of three simulated DBRs are shown in Figure 14. Each pair of layers, consisting of one layer of the high index material and one layer of the lower index material, is referred to as a single period. So, a DBR with 60 total layers is made of 30 periods. We show the simulated reflectivity for three different DBRs, all made of the same material with the same design wavelength, but with a different number of periods. Each simulated sample has an extra layer of high index material (hence the extra 0.5 periods), so that both the topmost and bottommost layers are high-index material. In a simulation of just a DBR, the starting and stopping layers can critically affect the final reflectivity curve. In more realistic scenarios, DBRs are typically grown on top of a substrate wafer, which is effectively infinitely thick compared to the DBR. As a result, in real-world DBRs, the starting and stopping layers are not as significant, and the results shown are typical.

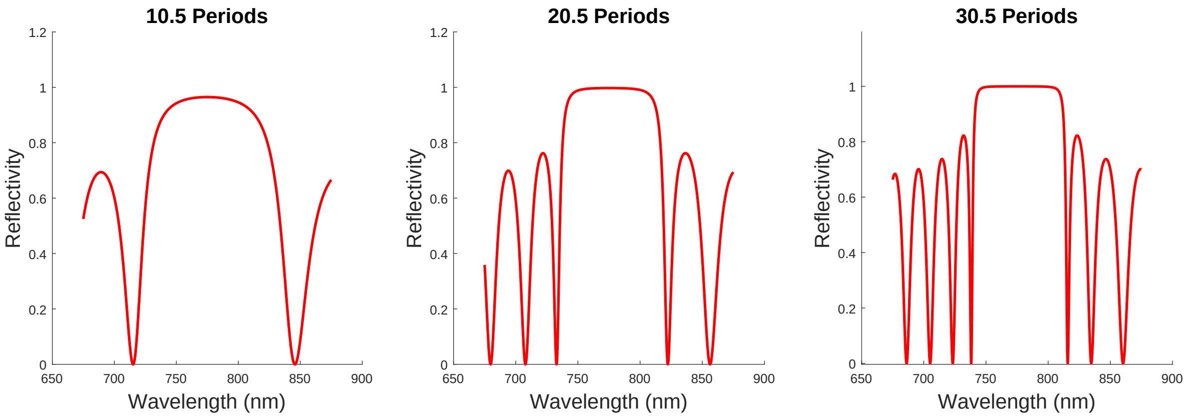

Figure 14: Simulated reflectivity for three different DBRs. Each DBR has a different number of periods.

The large flat region of the reflectivity curve, which is close to a perfect reflector, is called the photonic stopband, or simply stopband. In this case, we designed our DBRs for 775 nm, and we see the stopband is centered on 775 nm. These plots are for normal incidence, but these DBRs are still very good reflectors at larger incident angle, as we will see in chapter 4.

Now, for creating our DBRs we need to decide which of our three materials will be the low-index and high-index layers. We must keep in mind that our goal is ultimately to create a microcavity with quantum wells inside of it. In order to create quantum wells, we need a small band gap material sandwiched between two layers of large band gap material. The greater the difference between these two band gaps, the more trapped our excitons will be. This is desirable because we want radiative decay to be the only decay channel for our excitons, rather than diffusing out into other materials. So then, because GaAs has the smallest band gap of our three materials, and AlAs has the largest, it makes sense to create GaAs quantum wells by sandwiching a thin layer of GaAs in between two layers of AlAs.

As a rough rule, materials absorb photons at energies above their band gap. We do not want the layers of our DBR to absorb light; we only want it to be transmitted or reflected. So, because GaAs quantum wells will release light at energies higher than the band gap of bulk GaAs, GaAs would not be a suitable material for our DBRs. Therefore, we choose AlAs as our low index DBR material and AlGaAs as our high index DBR material.

So, we are ready to put it all together. There are a few more practical considerations however. We find we can create stronger coupling between photons and excitons by putting more than one quantum well at the anti-nodes of our microcavity, so we put 4 quantum wells at each anti-node. Additionally, by utilizing a $3\lambda/2$ cavity we triple the number of quantum wells, giving us a total of 12 quantum wells inside our cavity. We also need a material to fill the space inside the microcavity between the series of quantum wells. In a $\lambda/2$ cavity, the best spacer material is the lower index material. However, in a $3\lambda/2$ cavity, we've actually found it is better to continue the basic pattern of the DBR, so in between the sets of quantum wells are an entire DBR period approximately. Our microcavity design is shown in Figure 15A. Another way of viewing this is that our sample is a $\lambda/2$ microcavity, but we put additional quantum wells outside the cavity. In chapter 3 we will show simulations of the electric field, and how it leaks out of the microcavity, justifying this approach.

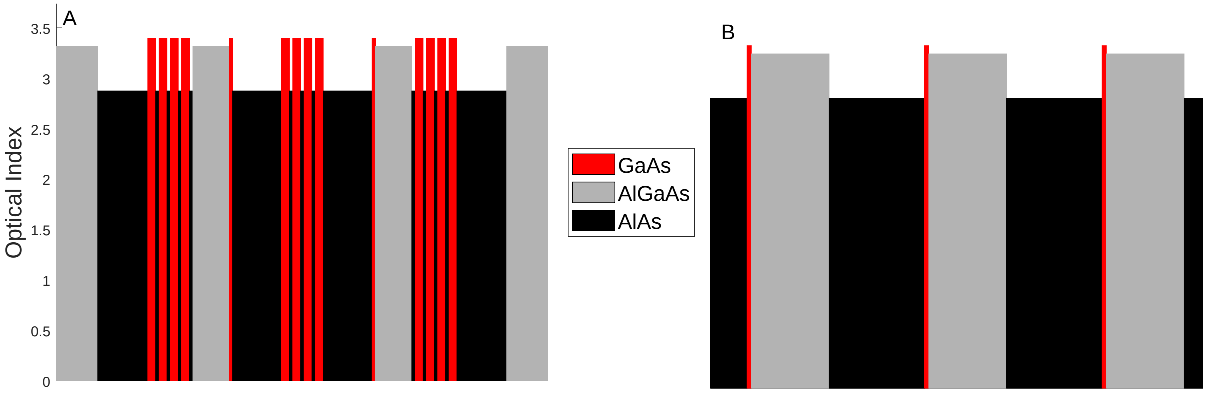

Figure 15: Image A: The design of our microcavity. Image B: The design of our bottom DBR; note the thin wetting layers.

Lastly, in Figure 14 we saw that more periods in our simulated DBR leads to a broader stopband, and overall higher reflectivity. However when growing samples inside an MBE machine, small defects, such as a lattice mismatch at the boundary, propagate upwards, growing and becoming worse with each subsequent layer. We have found that a thin layer of GaAs grown on top of the AlAs leads to better growth of the AlGaAs on top. This is called the wetting layer. A DBR with wetting layers is shown in Figure 15B. Typically, we find that wetting layers are only necessary for the bottom DBR, and not so important in the top DBR. The wetting layers are also present in the microcavity in Figure 15A. Remembering our convention that the left side is the bottom of the sample, and the right side is the top, these samples were grown left to right. One wetting layer is therefore grown on AlGaAs, with AlAs grown on top of it inside the microcavity. This is done to make the microcavity as symmetric as possible.

2.0 Experimental Methods

Two of my flagship achievements in this lab have been the development and refinement of two types of measurements which were previously unavailable to us; photoluminescent excitation (PLE) to measure the upper polariton, and angle resolved reflectivity [13], along with a corresponding simulation [12] covered in Chapter 3. These methods are essential in collecting the data required for long-lifetime sample characterization, as we will discuss in Chapter 4.

The development of these methods required a significant time investment in software development. Our lab had many tools at its disposal, but no one had unified them into a single software package. I integrated nearly all the tools in our lab so that they can now be controlled with a single centralized graphical user interface (GUI). I felt this approach was necessary to facilitate PLE and reflectivity measurements as practical tools. It was important to turn these measurements into practical tools that other students could utilize because the problem of sample characterization is not a one-off occurrence, every sample we receive needs to be characterized. The development of these tools were also useful for a large variety of other experiments conducted in our lab [30, 17, 3, 18, 88, 83, 81, 79, 80, 82, 56, 57, 58, 59, 43].

In the early development of PLE measurements, we struggled to complete three measurements in a single day's experiment. Additionally, because none of it was automated, with many steps, it was prone to human error. However, the most recent data written about in this work, temperature dependent PLE, includes approximately 25 data points measured in one day, which also consisted of significantly higher quality sets of data.

Appendix A serves as a manual for the software that controls the devices in our lab and collects their data, with specific information for how to perform these measurements using that software. In this chapter, I will focus on some of the more generic techniques leading up to PLE, as well as innovations created for our work in transition-metal dichalcogenide monolayer (TMD) systems. Because of its highly specialized nature, specifics of PLE measurements are discussed in Section 4.3 within the broader context of long-lifetime sample characterization. Angle-resolved reflectivity, as a more generic measurement, is discussed in this chapter in Section 2.5.

2.1 Calibration Of Fourier Imaging

One of the essential tools at our disposal is Fourier imaging, which ultimately leads to angle-resolved images, which are the most common type of data taken in our lab. For a thorough discussion of Fourier imaging see Ref. [41]. The key result for us being that through proper lens placement we may create a Fourier plane, where rays of light from our sample are grouped together based on their emission angle rather than their real space position (regular imaging). This Fourier plane can then be imaged through normal imaging optics.

A simple schematic of our basic optics setup is shown in Figure 16. These are free-standing optics on an optical table. Here we are utilizing an infinity-corrected microscope objective; the advantage of this is they tend to give larger numerical apertures, but compressed down into a narrower beam, which is easier to work with downstream. Two important planes form behind Lens 1, the reimaged sample plane and the reimaged Fourier plane. Either of these planes may be reimaged again onto the slit of our spectrometer by appropriate choice of Lens 2. Lens 2 can also be a compound optical system. The slit of the spectrometer may be fully opened or even removed entirely, and the grating angle of the spectrometer set to zero in order to reimage either plane directly onto our CCD camera.

Reimaging the sample plane onto our CCD camera is exactly as it sounds; essentially, a highly magnified image of our sample's surface is formed on our camera for collection. The Fourier plane on the other hand, results in an image showing us the intensity of light as a function of the angle that the light left the sample at. This includes both reflected light and emitted light in the case of PL measurements. In this way, we can measure the dispersion relation of our polaritons by directly imaging them.

However, one issue remains. We must somehow convert the location of the light on our camera from pixel number into angle or in-plane momentum. In order to create this calibration, we place a transmission grating at the sample plane of our optics setup; specifically we used a GT13-03 visible transmission grating (300 Groves/mm) from Thorlabs. Our

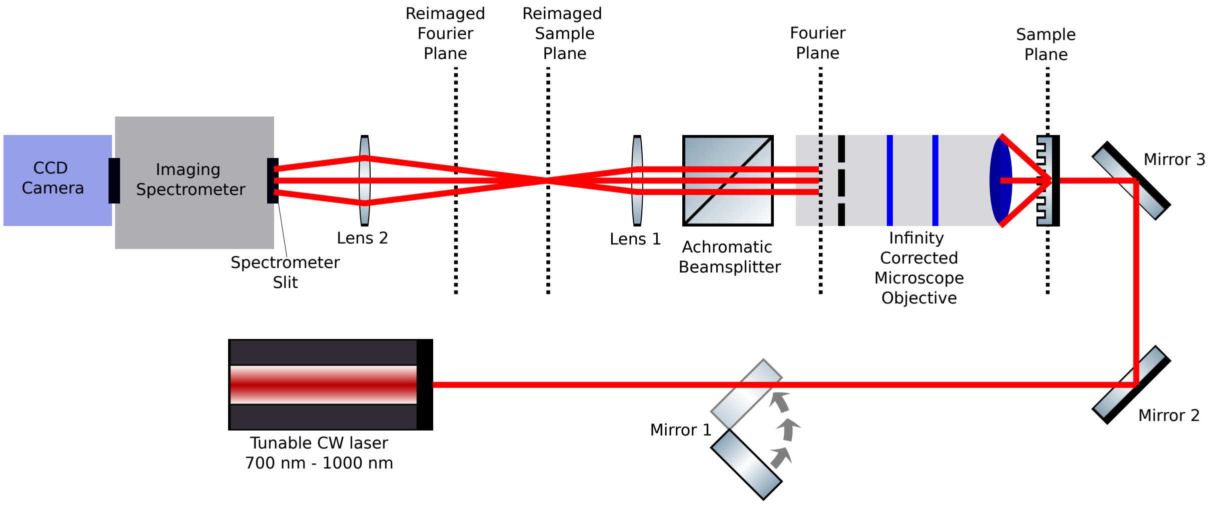

Figure 16: Illustration showing a basic Fourier imaging setup. Mirror 1 may be flipped into the laser path to switch from a transmission geometry to a reflective geometry. Here, a transmission grating is in the sample plane, and three resulting rays are shown propagating through the system. Credit to Alexander Franzen who created the vector graphics library for optical components used in our figures [34].

polariton samples are on top of an opaque substrate, and so all of our data is collected in a reflectivity geometry, so naturally we first tried using a reflective grating. This worked well enough if the microscope objective in our setup was our 20x Mitutoyo Plan Apo NIR infinity- corrected objective, which has a numerical aperture covering approximately ± 25 degrees. However, eventually we upgraded this optical system with a 50x Mitutoyo Plan Apo Infinited Corrected objective, which can collect light at an angle of ± 50 degrees. We like the larger numerical aperture because it allows us to see more of the dispersion curve. Both objectives may be placed into the optical setup with minimal changes because the objectives and their mounts are designed to be interchangeable (as you would see in a nice microscope). The 20x objective can also be used to measure the same range of angles in steps by rotating the cryostat [58], but the 50x objective is significantly easier to work with.

Figure 17A shows the data collected using the transmission grating along with our 50x objective. We see a series of dots form as expected from a grating.

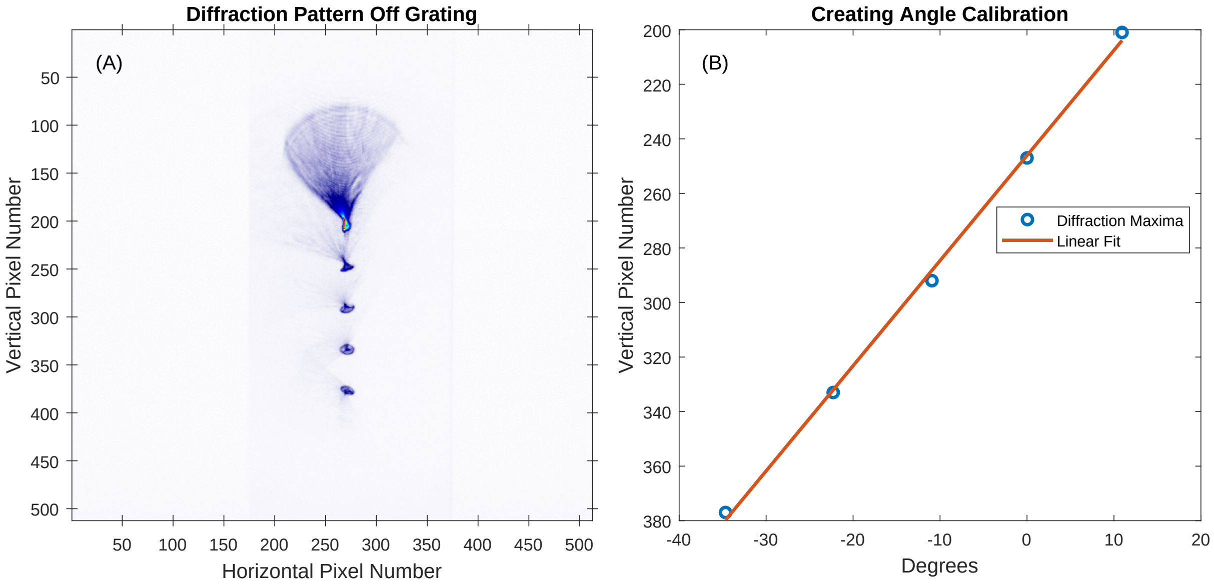

Figure 17: Data used to create the calibration of our angle resolved imaging.

Using the grating equation at normal incidence,

$$m\lambda = d \sin\theta, \tag{2.1}$$

where d is the inverse of the grating density, we may calculate the emission angle of the dots. The grating is blazed so that the brightest dot corresponds to the $m = 1$ diffraction, and no other positive m values show up; instead we see the negative m values as well as the $m = 0$ dot. The pixel number of these dots is extracted, and then emission angle vs pixel number is shown in figure 17B. We see the data is fit very well by a linear fit. This fit will apply to all measurements in this setup as long as we do not change the positions of any of the optics.

2.2 Spectrometer Calibration

Our characterization of high Q polariton samples is ultimately going to combine three sets of data; PL measurements of the lower polariton, PLE measurements of the upper polariton, and reflectivity measurements of the photonic stopband of the entire structure. All of three sets of data will feature wavelength measurements. In the case of PL and reflectivity, our wavelength data comes from our spectrometer. However, as we will discuss in Section 4.3, in PLE the wavelength data comes from the tunable laser, specifically an Ångstrom WS/6 Wavelength Meter from HighFinesse which measures a small pick off beam from the laser in order to report its wavelength. It is essential that these two types of measurements are in agreement with one another. In order to accomplish this, we used the laser to calibrate the spectrometer.

The spectrometer in this case is a SpectraPro-2500, although we've performed this procedure on a few optical setups in our lab. The SpectraPro-2500i has built in software which gives approximately correct results, after being given appropriate grating information. For example, if we set our spectrometer to 750 nm, then a laser beam at 750 nm will approximately land at the center of our camera's sensor. However, it is our job to create a calibration for all the other pixels.

It is critically important to mention that calibrating the spectrometer using the laser will involve a laser going directly into our spectrometer. Extreme caution should be exercised to make sure the laser's power has been cut down enough to not burn the pixels of our camera. This can be easily accomplished by placing a series of neutral density filters in the laser's path. Because the laser is approximately 3 Watts, care must be taken to use reflective filters, absorptive filters will heat up and shatter. Care must also be taken to ensure the throw-away reflected beam from these filters is safely terminated.

The optics setup shown in Figure 17 is slightly modified. Mirror 1 is flipped into the laser's path to convert to a reflective geometry, and a good mirror is placed in the sample plane. The beam coming out of the laser is expanded so that the resulting image is a line of laser rather than a single dot. This isn't strictly necessary, but it makes things easier.

The basic procedure is as follows. We start by setting our spectrometer to a central wavelength λ_{spec}, we then scan our laser in steps ranging from $\lambda_{spec} - \Delta$ to $\lambda_{spec} + \Delta$. Where Δ is chosen to cover the entire field of view of the camera. Determining Δ is done quickly through simple trial and error. An image of the laser is taken at each step, examples of this data are shown in Figure 18 images A, B, and C.

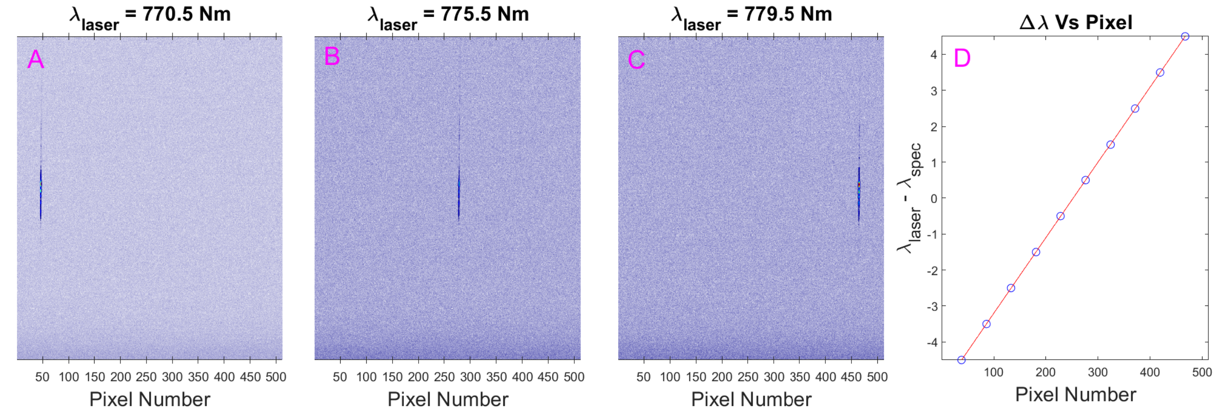

Figure 18: In image A, B, and C we have some of the data used to create our grating calibrations. Image D is the extracted pixel number of our laser, vs the wavelength of the laser relative to the spectrometer's wavelength setting. We see the relationship is very close to linear.

Profiles across the image are then fit with a Lorentzian distribution to extract the center point of the laser as measured by pixel number. The wavelength of the laser is obtained from the laser's attached wavemeter. We express our laser's wavelength relative to the spectrometer's wavelength setting λ_{spec} values, as $\Delta\lambda = \lambda_{laser} - \lambda_{spec}$. The results are shown in figure 18D. We see the data is very linear, we use a quadratic fit to the data given as

$$\Delta\lambda = A + BP + CP^2, \tag{2.2}$$

where P is the pixel number. The possible values of λ_{spec} are essentially continuous. So, in order to calibrate our spectrometer for any of the possible values of it, we fit the coefficients A, B, and C as functions of λ_{spec}. We perform the above fitting at many different λ_{spec} values and plot the values of the coefficients as shown in Figure 19.

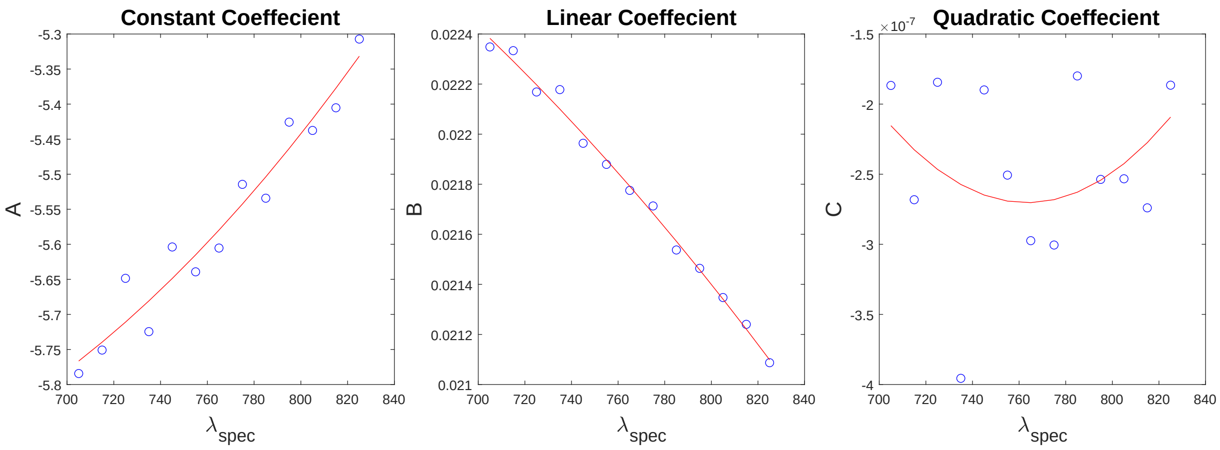

Figure 19: The fits of the coefficients in our grating calibration. We see the quadratic term is not important.

We see that in general the quadratic term contributes minimally to the final fit. The coefficients are now themselves fit as functions of λ_{spec}. We see that fits to quadratic order will be sufficient. We have the following three equations.

$$A = A_0 + A_1\lambda_{spec} + A_2\lambda_{spec}^2 \tag{2.3}$$

$$B = B_0 + B_1\lambda_{spec} + B_2\lambda_{spec}^2 \tag{2.4}$$

$$C = C_0 + C_1\lambda_{spec} + C_2\lambda_{spec}^2. \tag{2.5}$$

So then, our grating calibration can be boiled down to the nine fitting parameters in the above three equations. Those three equations are used to generate the polynomial coefficients for Equation (2.2), finally giving us a calibration between wavelength and pixel number on our screen.

The process of creating this calibration must be performed for every optical setup. For example, in Bay 1 we have two cameras attached to a spectrometer with 3 gratings. Thus, Bay 1 of our lab requires six separate calibrations. This process is streamlined in our software, along with the data-processing algorithms, which are discussed in Appendix A.

The discrepancy between the fits and the data for coefficient A might be a little visually alarming. However, keep in mind that A is essentially a shift in pixels. So we see the discrepancy between our fit and the measured value is of the order of a tenth of a pixel. B is essentially a measurement of nanometers per pixel, around 0.02 nm/pixel, allowing us to estimate a discrepancy of only .002 nm which would be approximately 4 μeV. This is error associated with approximating A as smoothly varying with λ_{spec}; it is not indicative of an absolute error of our system.

Note that because the spectrometer consists of a rotating grating, care is taken to avoid backlash affects. When the spectrometer is set to a wavelength, it always tunes below that wavelength, then approaches the desired wavelength from the negative side. This gives us good repeatability.

2.3 Beamsplitter Calibration

There are many situations in which measuring the power of a laser delivered to our sample is important. This can be done very efficiently with some small changes to the optical setup shown in Figure 16, as well as creating a calibration for our beamsplitter. For details about how to use my software to create that calibration file see Appendix A.2.2. We start by placing a power meter in the sample plane of our imaging setup, which we name the reflective power meter, which measures the power of the beam P_R. Then, we place a second power meter in the path of the transmitted beam of our beamsplitter to measure the power P_T. This optics configuration is shown in Figure 20.

We are typically interested in performing these measurements at many different wavelengths. So, in order to create our calibration, we scan our laser across the region of interest while measuring the power of the laser as measured on both power meters. We assume that no power is lost to absorption, which allows us to calculate the power reflection and transmission ratios of our beamsplitter as

$$R = \frac{P_R}{P_R + P_T} \tag{2.6}$$

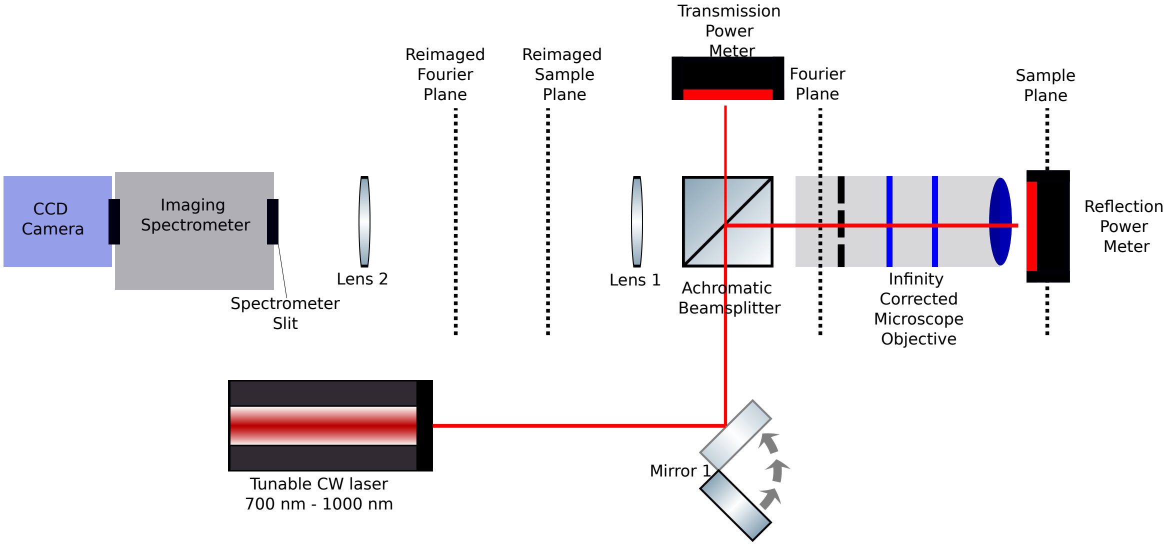

Figure 20: Simple illustration of our optics configuration for calibration of our beamsplitter. Credit to Alexander Franzen who created the vector graphics library for optical components used in our figures [34].

and

$$T = \frac{P_T}{P_R + P_T}. \tag{2.7}$$

The assumption that no power is loss to absorption is not actually important to our final result, as the denominators of these two ratios will cancel one another. R and T are simply nice quantities to plot to understand broadly how our system is performing. Data showing the calibration of one of our setups is shown in Figure 21.

This particular beamsplitter has a high transmission and a low reflection coefficient. This is a design choice when making an optical setup. In our case, for PLE measurements we tend to have more laser power than we need. We actually don't want 100% of our laser to reach our sample because it would significantly warm up our sample and introduce a

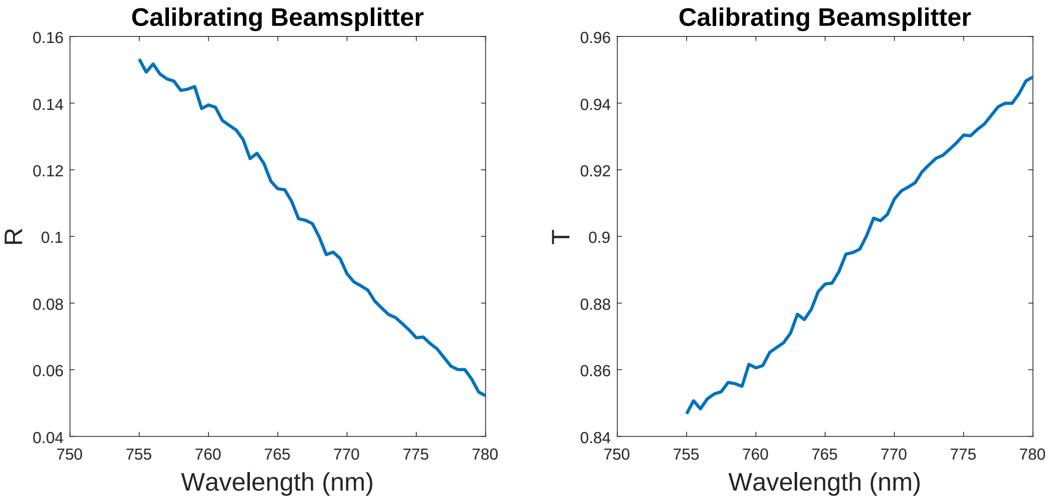

Figure 21: Calibration data for our beamsplitter.

corresponding red shift, and also possibly a density dependent blue shift. For this reason, PLE measurements are best performed at low powers. So, we are actually happy to have more of our laser be transmitted through the beamsplitter, because the upside is that the high transmission rate also applies to the PL we measure from the polaritons. So, we are creating fewer polaritons than we could, but we are successfully collecting more of the light from them.

When we are ready to perform our experiment, we remove the reflection power meter from the setup, and place our sample in the sample plane. Then, when we measure the power on the transmission power meter, the power delivered to the sample may be calculated simply as

$$P_{sample} = \frac{R(\lambda)}{T(\lambda)} P_T. \tag{2.8}$$

We see that the identical denominators in R and T cancel out here. This allows us to easily measure the power of our laser beam as we sweep it in our PLE experiment.

One last thing to mention is that often our power meters have a calibrated voltage response curve for measuring power. This curve is wavelength-dependent. However, our software can be configured to automatically configure our power meters to the wavelength of the laser as we sweep the laser. Some of our larger power meters are thermal power meters, and do not have such wavelength dependence in their calibrations.

2.4 Camera Calibration

In our work we found that the response of our cameras were not scaling linearly with integration time. To test this, we send a weak laser directly into our system, and measured the average intensity of a number of pixels in the resulting image. We swept our integration time from 10 ms to 10000 ms, the minimum and maximum values which the camera is supposed to be able to handle. Longer integration times can be effectively achieved through multiple image being added together. This nonlinear response is very problematic as we frequently wish to compare the intensity of two different images with different integration times. Worse yet, we found that there was a sort of hysteresis in which an image's intensity was affected by the integration time of the prior image. This means even images with the same integration time could not necessarily be compared. Data of this phenomenon is shown in Figure 22.

For comparison, we have plotted a linear function which is fit to the forward sweep data above 2000 ms. We wee that in general the camera response is scaling linearly with time for large integration times, but there is a very clear disparity for short integration times.

We believe we have figured out the problem and a solution. The problem appears to be some sort of read out noise. The value of the readout noise seems to be in some way dependent on the previous image. We found that the solution is to always take a blanking image before taking our real image. The blanking image is collected with an integration time of zero, and then the blanking image is subtracted from the real image. This does not completely fix our error. The data in Figure 22B is collected by repeating our sweeps but with the readout noise subtraction now turned on. We see that the three sweeps have very good agreement above 100 ms or so. Again, we have included a linear fit of the first forward sweep above 2000 ms for comparison.

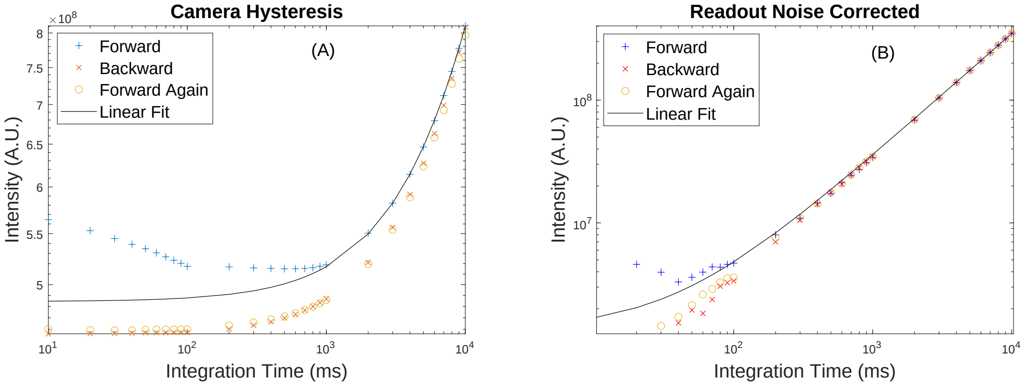

Figure 22: In image A we see the hysteresis exhibited by our camera. In image B, we see that the readout noise subtraction mostly fixes this problem. Both plots are log-log scale.

This is acceptable to us, as we can easily work around these limitations. In general, we try to collect images with an integration time of 1000 ms or more (depending on how bright our signal is). This puts us firmly in the regime of the camera that works well. All data presented in this work has had readout noise subtracted off.

2.5 Angle-Resolved Reflectivity

Now we turn our attention to reflectivity measurements. Reflectivity measurements are certainly not novel to our lab or the world. Indeed, there are many very nice boxed systems one can purchase to perform high quality reflectivity measurements; often reflectivity is just one of many functions an ellipsometer offers. However, before my work, our lab had never done angle-resolved reflectivity measurements.

In the past our lab was focused on short-lifetime samples [7]. In those samples, the polariton branches were visible in the reflectivity spectrum. However, Fourier imaging was not yet a commonly used tool in the field. Indeed, the original angle-resolved PL from that era was collected using about 30 photodiodes mounted along the circular arc of a goniometer. This was not a setup that could practically perform a reflectivity measurement; as a result the reflectivity data from that era was not angle-resolved.

Eventually, Fourier imaging became a well known technique, and so the tools to perform angle-resolved reflectivity were within our grasp. However, around this same time, our lab transitioned towards focusing on long-lifetime samples. In long-lifetime samples, reflectivity was no longer a useful tool to measure polaritons because the linewidth of the polaritons was too narrow to resolve in reflectivity. As a result, our lab never seriously tried to perform angle-resolved reflectivity measurements, sticking purely to reflectivity at normal incidence. Reflectivity became a measurement we would perform one time quickly to confirm our polaritons were indeed in the stopband, as designed. However, as we will discuss in Chapter 4, reflectivity is a critical third measurement we need to characterize our samples.

Furthermore, it is and has been somewhat common in various short-lifetime polariton publications, such as in transition metal dichalcogenide monolayers, to perform an angle-resolved white light measurement which is not truly a reflectivity measurement. This is still colloquially referred to as reflectivity sometimes. In this angle-resolved white light measurement, the sample is illuminated with a broad white light spectrum, and the reflected light is imaged in a Fourier imaging setup. This could be accomplished by flipping mirror 1, as shown in Figure 16,to switch to a reflection geometry, and then replacing the laser with a broadband white light source. The resulting raw image is sufficient for determining the dispersion relationship of the polaritons, as their dispersion curve will show up as a dip in the image. Its a perfectly valid measurement, it just isn't a true reflectivity measurement.

All this is to say that my work here of performing true reflectivity measurements with free-standing optics on a sample in a cryostat is somewhat novel, possibly unique, and certainly a new technique for our lab specifically.

Now let us define what it is we are measuring. We seek to measure the power reflectivity given as

$$R(\vec{\theta}) = \frac{\vec{E_{out}}(\theta) \cdot \vec{E^*_{out}}(\theta)}{\vec{E_{in}}(\theta) \cdot \vec{E^*_{in}}(\theta)}, \tag{2.9}$$

where E is the amplitude of the electric field for the incoming and outgoing wave, with θ defined in Figure 23. In this case, we are ignoring polarization of the light. Essentially we are interested in performing unpolarized reflectivity measurements, which are a 50-50 mixing of the S and P polarization measurements. This is our preferred setup. Polarizers may be added to measure either S or P polarized reflectivity, and then rotated to measure the other polarization, but it takes much more than double the amount of time to measure both S and P polarized reflectivity. Because S- and P-polarized reflectivity are exactly the same at normal incidence, we would be less than doubling our total amount of measured data. Ultimately, for most cases we don't find it makes practical sense to collect polarized reflectivity. But we have done it, and it is possible, and we see the expected behavior for a splitting of the polariton states at higher angles due to TE-TM splitting in the cavity; a high angle the splitting is also observable in unpolarized reflectivity images since the two lines are far apart.

The main issue with performing this measurement is calibrating the incoming electric field. Our incoming electric field is also unknown, so we must measure it. In order to do so, we replace our sample with a very high quality mirror, and use measurements off that mirror as a substitute for E_{in}. Ideally, that mirror would be a perfect reflector; in practice, a very good metal mirror is often good enough, and if needed the manufacturer's stated reflectivity R_M for the mirror may be used to correct our result.

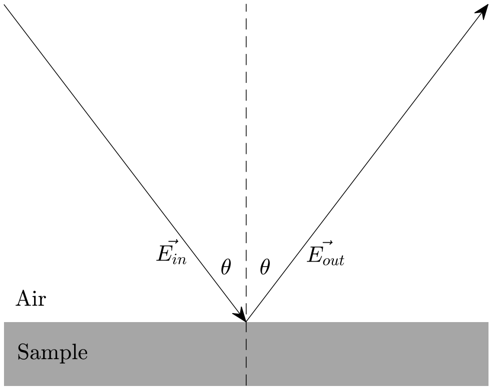

Figure 23: Simple illustration of the light rays in a reflectivity measurement.

We illuminate the mirror with a broad white light source, and take a series of images in order to construct the measured light coming off of it $N(\lambda, \theta)$. Section 2.5.1 discusses the step-and-glue algorithm which is used. In principle N could be obtained with a single image. We have taken to calling N the "normalizer". We will try to refrain from using that name in this text, but we mention it because we have also published a series of videos online detailing how to take these measurements and process them with our software [11], and in those videos that name is used.

We must also take a series of images for the background for N; this is easily accomplished. We may remove the mirror entirely and just let the beam terminate at infinitely, or we may block the mirror with highly absorbing optical tape. Either way, we obtain our background images $B_N(\lambda, \theta)$.

We perform similar measurements with the sample rather than the mirror and obtain the images of the sample $S(\lambda, \theta)$ and their corresponding background $B_S(\lambda, \theta)$. The background for the sample is often obtained in practice by putting highly absorbing optical tape inside our cryostat, next to our mounted sample. This has the advantage of collecting the background caused by reflection of the cryostat mirror, so that it may be properly subtracted from our images.

With our four sets of images, the power reflectivity may now be simply calculated as

$$R(\lambda, \theta) = \frac{S(\lambda, \theta) - B_S(\lambda, \theta)}{N(\lambda, \theta) - B_N(\lambda, \theta)} \frac{1}{R_M(\lambda, \theta)}, \tag{2.10}$$

where $R_M(\lambda, \theta)$ is the manufacturer specifications for the mirror used to collect the N images. Ideally this data would have full angle dependence in it. In practice, often the best we can get is wavelength dependence only, so we have $R_M(\lambda)$ instead. As it turns out, this correction mostly has the affect of uniformly scaling the reflectivity we measure up by a few percent, as in the visible light range, metallic mirrors have a fairy flat reflectivity curve.

For the most part, systematic errors such as collection efficiency will divide out here. As discussed so far, reflectivity looks quite simple. However, the crux of the problem lies in the process of taking data on two separate objects, our sample and the mirror. Ideally, we want everything exactly the same between data sets, specifically the optics alignment. The problem is how to move either our sample or the mirror out of the sample plain, and then move the other into the sample plain, and then align the tilt of it, without changing any of our other optics. Taking high-quality reflectivity data requires careful attention to detail. In the following sections we will discuss some of these finer details.

The reader may wonder why these issues may not be simply circumvented all together through ellipsometry. In ellipsometry, no second reflector would be necessary; instead a polarizer could be placed on the white light source, and rotated to perform two measurements for S and P polarization. Indeed, that is one of the beautiful advantages of ellipsometry: one not need to know the electric field of the source. Ellipsometry gives a slightly different set of data, but for our purposes accurate ellipsometric measurements would work fine.

However, we tried doing ellipsometry and found that this did not work well in practice. We believe the issue came from the beamsplitter. At the time we were using a pellicle beamsplitter. Pellicle beamsplitters are quite nice for several reasons; however they also have a strong polarization dependence. If the issue was caused by the S-P splitting of the beamsplitter, a cubic beamsplitter could mostly fix this. This could be an interesting project for a future student to take up. Our simulator and associated GUI described in Section 3 has options for ellipsometry built into it.

An example set of data is shown in Figure 24A with corresponding profile along $\theta = 0$ shown in image B. The data shown is some of the nicest reflectivity data we have taken. In this case, no processing of the data was required; Equation (2.10) was enough (along with the step and glue algorithm with pixel binning described in Section 2.5.1). However, we have also devised a data-processing method which can partially fix some of the issues which arise from misalignment in the optics between the data sets, discussed in Section 2.5.6.

Section 2.5.2 describes the standard optical alignment procedure to bring a sample into the focal plane of our setup and properly adjust its angle. Then Sections 2.5.4, 2.5.3, and 2.5.5 discuss variants of how to bring the second reflector into the setup with minimal misalignment. We believe the 5-2 routine is the best choice.

In all our reflectivity routines, it is important to take the brightest images first, and so N and B_N are always taken before S and B_S. This is explained more thoroughly in Section 2.5.1.

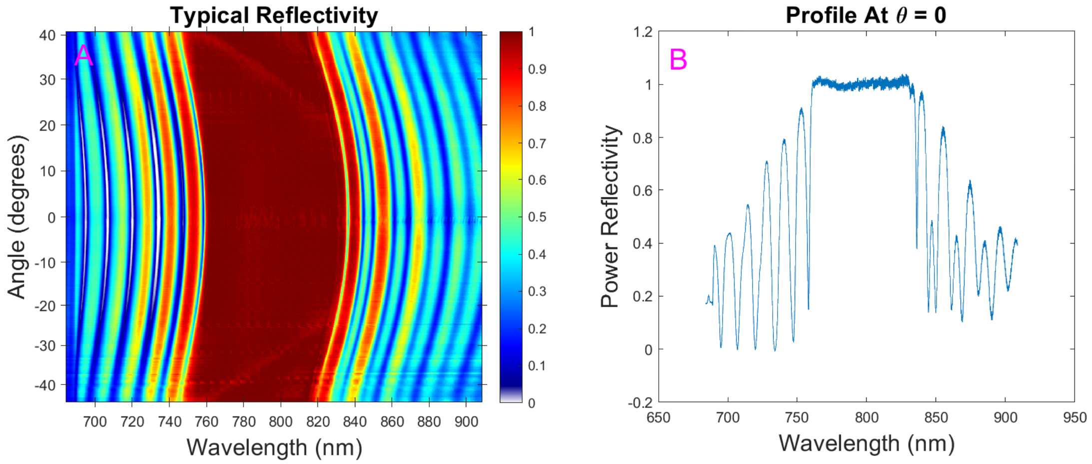

Figure 24: Example reflectivity data measured from one of our long lifetime samples. Image A is the full two dimensional angle resolved reflectivity. Image B is the profile of image A at normal incidence.

2.5.1 The Step-and-Glue Routine

In this section we focus on two algorithms from our software which, when combined with the standard alignment described in Section 2.5.2, make up the bulk of our reflectivity measurements. The step-and-glue algorithm, specifically with the lock-signal-to-noise feature enabled, as well as the copy-prior-scan algorithm.

The step-and-glue algorithm is a method of creating a larger image from smaller images. Its need arises from a mismatch between our spectrometer grating's field of view and our desired wavelength range over which we wish to measure reflectivity. We can see from Figure 18 that our grating covers a range of approximately 10 nm, while Figure 24covers approximately 200 nm of range. It is possible to bypass the need for a step-and-glue algorithm by using a grating in the spectrometer with a lower line density. Indeed, in our spectrometer we have installed both a 50 g/mm and 1200 g/mm grating. The 1200 g/mm grating is the one used with a field of view of 10 nm. The 50 g/mm grating, on the other hand, results in a field of view of about 325 nm. The problem with the 50 g/mm is twofold. Because the field of view is approximately 32 times larger, the resolution is about 32 times lower. At the outset of this project, we did not know with certainty that it would not be possible to measure the polaritons in reflectivity. In order to have a hope of doing so, we needed the best resolution we could obtain. Additionally, the 50 g/mm grating results in a noticeable interference fringes, leading to a stopband spectrum that isn't flat, but rather sinusoidal. There are plenty of situations where the significant time savings of using the 50 g/mm grating are warranted, but for publication quality reflectivity curves, we prefer to stick to the 1200 g/mm grating.

The step-and-glue algorithm sweeps the spectrometer wavelength λ_{spec} in order to cover the entire wavelength range of interest. The step size we use is approximately half the wavelength range displayed on our screen. So for example, the first step may have the wavelength centered on 700 nm, with the screen capturing data from 695 nm to 705 nm. The next step will be centered on 705 nm capturing data from 700 to 710 nm. This step size could be larger, but we do it for very good reason. This way if we do manage to find the polariton in our reflectivity, this will hopefully result in a non-stitched image in which the polariton is positioned nicely on our screen.

One major issue we encounter is that as we step along in λ_{spec} the intensity of our images changes. This is because the bulb of our white light source is a real bulb, and so to first approximation the emission spectrum is similar to a blackbody emitter. This is quite problematic for us, as most bulbs available have a very weak signal in the infrared portion of the spectrum, starting at about 775 nm, right where our data is. Fortunately, there are infrared enhanced bulbs available, which manage to create a somewhat flat emission spectrum with strong IR components. We use the OSL2BIR bulb from ThorLabs. The unfortunate downside is that it after a few hundred hours of use, these bulbs start to melt and burn the fiber optic cable we use to position and collimate the white light. However, the cable is only about \$150, which is a fraction of the cost of a tank of liquid helium, so we can accept this as an operating cost considering its longevity.

Still, the signal intensity changing as we step along remains a problem. For our step-and-glue we can not leave the camera integration settings the same. They must dynamically tune as we step. Fortunately, we've programmed an optional feature into our software named lock-signal-to-noise. Programmed into our software is a table of camera settings, based on the available equipment. An example of this is shown in Table 3.

Table 3: A table representative of the camera and filter wheel settings I used to lock the signal to noise ratio during the step-and-glue scans. The true table has many more steps.

OD Filter	Gain	Integration Time (ms)
3	1	100
2	1	100
1	1	100
0	1	100
0	3000	100
0	5000	100
0	5000	2000
0	5000	10000

Here the gain is a variable controlling the CCD camera's electronic gain. The OD filter setting is an electronic motorized wheel with 5 reflective neutral density filters on it, the densest being an OD5 filter and the weakest being an OD1, the 0 setting is an empty hole in the wheel (no filtering). When the lock-signal-to-noise feature is enabled, as our step-and-glue algorithm steps along λ_{spec}, it dynamically monitors the intensity of the signal, and moves up or down in this table. If the signal is too weak it moves down to increase the signal, and if the signal is too bright, it moves up. Critically, it records these settings in the data file that contains the N images.

Recording the settings of each image is essential, because after taking the N images, the data file is used as a template by the copy-prior-scan algorithm. This algorithm returns to the exact settings which went into the N images file, takes a new image, and saves it into the original file. This way, when we calculate our reflectivity using equation 2.10, we do not need to normalize the images by integration time, camera gain, filter wheel setting, etc.

It is important to collect the N images first because they will always be the brightest images. If we collect the S images first, when we collect the N images afterwards, we may overexpose the camera. In the 5-2 reflectivity routine discussed in Section 2.5.5, the S images are collected first, but an additional set of reference images N_r are collected first in order to establish the optimal settings which the copy-prior-scan algorithm will use. As mentioned, the 5-2 method is what we recommend as the best method.

So, to summarize, the step-and-glue algorithm is used with the lock-signal-to-noise feature enabled to collect our N images, then the copy-prior-scan algorithm is used to collect the B_N, B_S, and S images. Equation (2.10) is used to calculate the reflectivity from these images. To produce the single unified image shown in Figure 24, all the images (which are just matrices) are concatenated into one giant matrix. The vector of the wavelength axis values for each image are also concatenated into one giant vector. These are both then sorted by the concatenated wavelength vector. This sorting is necessary because the images will have significant overlap at the edges.

Finally, a pixel binning procedure is used along the wavelength axis for publishable data. This is necessary because an image with 512 by 21,000 pixels can not be adequately displayed or printed. Furthermore, we will discuss in Chapter 4 fitting our transfer matrix method (TMM) simulator to this data, and fitting the full data set would be far too slow. This binning can be tailored to the desired resolution. The data in Figure 24has been binned to a resolution of 0.1 nm, which leaves us with a more manageable 2245 pixels in the wavelength dimension.

2.5.2 Standard Alignment Procedure

Here we describe the standard alignment procedure for imaging anything. This processes is the cornerstone of our reflectivity measurements, and in some cases is used multiple times (once for the sample, and once for the mirror).

We begin with Figure 25, which is a more thorough representation of the optical setup in our lab. Some elements are still omitted for clarity. Specifically, Lens 2 and 3 are two separate lenses on flip mounts at different locations. Lens 2 reimages the Fourier plane onto the spectrometer (angle resolved imaging), and Lens 3 reimages the sample plane onto the spectrometer slit (real space imaging).

I present here a step-by-step guide for the basic alignment procedure. This procedure is done at the start of each day's experiment. It corrects for small amounts of drift that occur over time. All the mirrors are on standard kinetic mounts, which allow the mirror to be rotated both horizontally and vertically for alignment, without moving the mirror significantly along the optical axis. The position of the mirrors is never adjusted during the alignment procedure, only the tip and tilt. All lenses are on dual micrometer stages, which allow them to move in both directions perpendicular to the optical axis. The tip and tilt of the lenses are never adjusted during the standard alignment procedure. All apertures may be opened or closed down, typically closed down when used for alignment, then opened once alignment is complete.

Before beginning the alignment procedure, Mirror 2 is flipped out of the laser path, usually the white light source is powered down. Lenses 1, 2, and 3 are all flipped out of the imaging path. The microscope objective is removed from the laser path. The CW laser is turned on. Once the system is prepared in this state, the following steps are followed:

1. Mirror 5 is adjusted to center the laser on aperture 1. Mirror 4 is adjusted to center the laser on aperture 2. This process (walking the beam) is repeated until the laser is centered on both apertures. Typically very small (if any) adjustments are needed.
2. Mirror 1A is adjusted to center the laser passing through the beamsplitter on aperture 3. Mirror 1B is adjusted to center the laser on aperture 4. This process is repeated until the laser is centered on both apertures. Typically very small (if any) adjustments are needed.
3. The beamsplitter is not aligned on a daily basis. As a general rule of thumb, it is never moved. If it is moved, it maybe be realigned by tuning its position and angle so that the laser coming off it passes through an alignment aperture placed approximately where the cryostat normally mounts, and terminates on a mark on the wall approximately 4 m away.
4. Whatever we are imaging is mounted in approximately the sample plane. This can be a mirror for our N images or the sample inside or outside the cryostat. From here forward we will refer to it as the sample, even if it is the mirror or a sample outside of a cryostat.

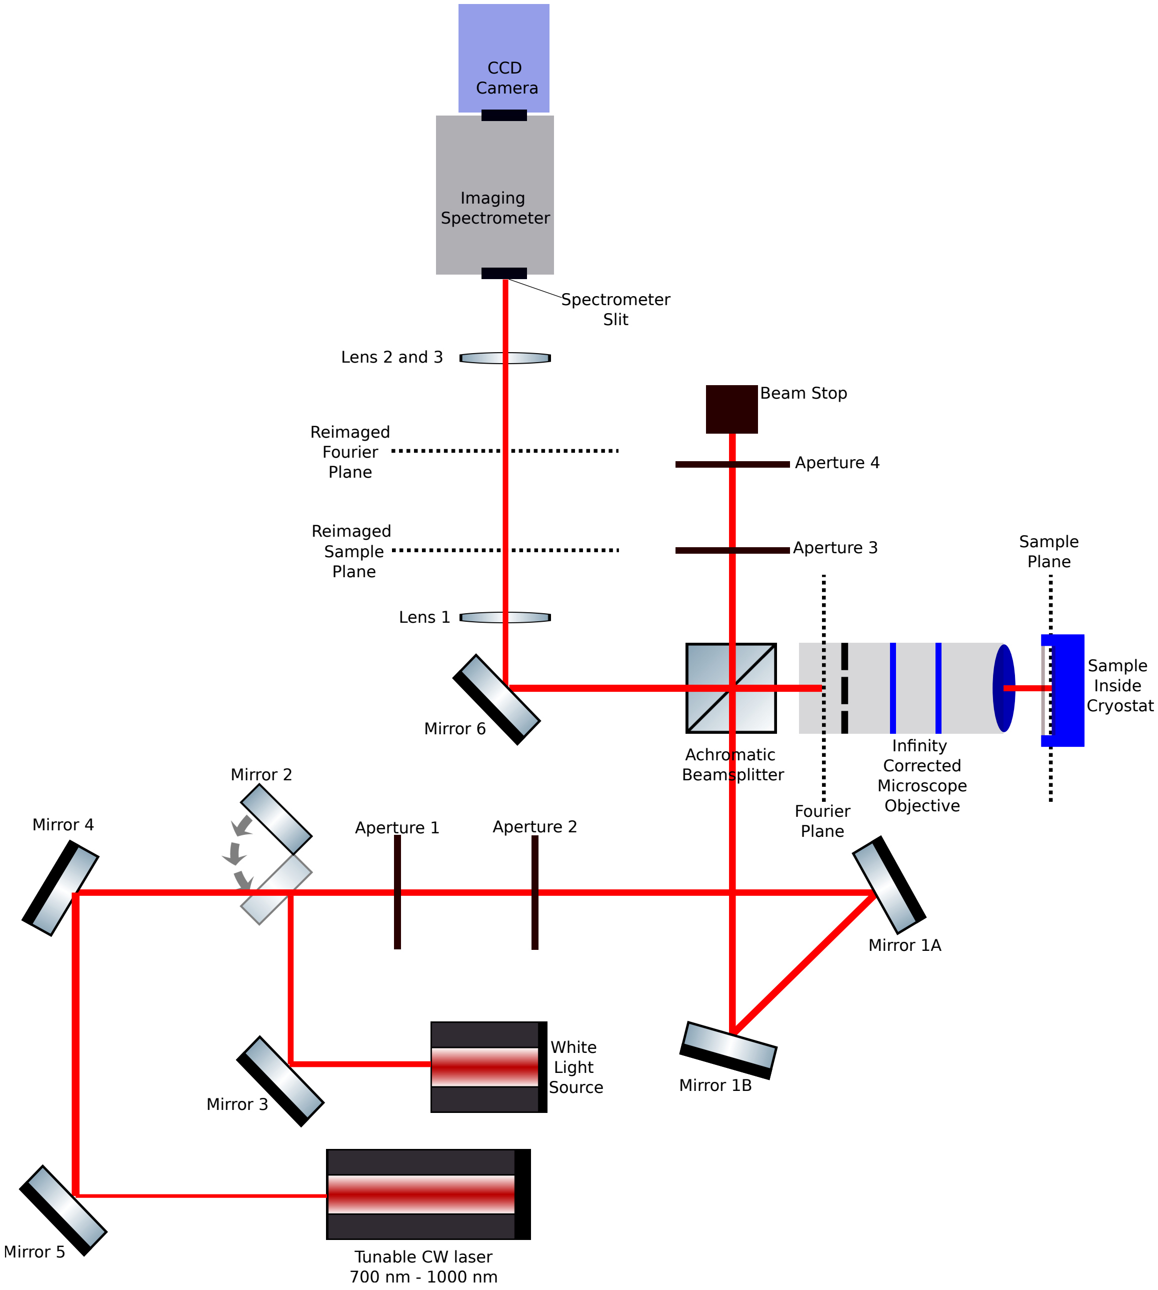

Figure 25: Illustration of the optics necessary for reflectivity measurements. Credit to Alexander Franzen who created the vector graphics library for optical components used in our figures [34].

5. The sample's tilt is adjusted so that its back reflection is going back through aperture 1 and 2. This step can be difficult. Our cryostat is not mounted on any type of rotational mounts, because they are not so practical. This alignment involves loosening the mounting screws and manually adjusted the angle of the cryostat. While the screws are loose, small pieces of paper may be wedged between the cryostat and the mounting apparatus in order to slightly adjust how the cryostat sits once the mounting screws are tightened. This step ensures that the sample is normal to the optical axis.
6. The software for controlling the CCD camera is turned on, and the camera is turned to auto imaging. The auto-tune feature is enabled, which will increase filtering if the signal becomes too bright. The spectrometer slit is closed, and then a digital crosshair is placed on the screen which is displaying the camera's images. The crosshair is in the middle of the slit, both vertically and horizontally. The digital crosshair will act as an alignment aperture.
7. The spectrometer slit is fully opened.
8. Mirror 6 is tuned so that the laser is centered on the digital crosshair on the screen.
9. Lens 3 (the real space imaging lens) is flipped into the optical path and adjusted so that the laser is centered on the digital crosshair on the screen.
10. Lens 1 is flipped into the optical path and adjusted so that the laser is centered on the digital crosshair on the screen.
11. The microscope objective (MO) is inserted into the optical path. The MO is on a gimbal mount. A bubble level is used to ensure the MO is level, if it is not, the gimbal mount may be adjusted. Precaution must be taken, because the MO is quite heavy, the springs of the gimbal mount sometimes do not return to the proper resting position. If the MO is not level, one can try manually assisting the springs and see if that corrects the issue before adjusting the knobs on the gimbal mount. The position of the MO should not need to be adjusted, as its double optical posts are both collared with anti-rotation adapters as well.
12. The cryostat is mounted on a lab jack for vertical motion, and two micrometers for horizontal motion. One of the micrometers is along the optical axis, named the focal micrometer. The focal micrometer is adjusted until the laser image on the screen is

brought to a minimal point size. Typically this will be one pixel on the screen. Care must be taken to avoid overloading the camera with too much signal or burning a pixel.

13. If necessary, Lens 3 is slightly adjusted to center the the laser spot on the digital crosshair.
14. The setup at this point is in real-space imaging. For this, the white light source is turned on, and flip mirror 2 into the optical path. Mirror 2 should be adjusted to ensure that white light is reaching the sample. In general, the white light beam is broad and diverging, so not much alignment is required.
15. For real-space imaging, one must slightly adjust the focal micrometer of the cryostat until the image on the screen is crisp. One may need to move the sample around to find a structure to image. A little spec of dirt is a good structure to image here. The purpose of this step is to approximately move the cryostat into correct focal position to make later adjustments easier. (This step can be skipped in the system is well aligned.)
16. There is a real-space aperture located in the reimaged sample plane after Lens 1, which is mounted on a dual micrometer stage to allow movement in the two directions perpendicular to the optical path. This aperture is now closed down, and its position adjusted until it is centered on the digital crosshair. This aperture should be closed for all images taken for angle-resolved reflectivity.
17. Lens 3 is flipped out of the optical path, and Lens 2 is flipped into the optical path. The image on the screen should be a large round circle. This circle is then centered on the digital crosshair.
18. In general, the real-space imaging setup and the k-space imaging setup can have slightly different alignment. Adjust the focal micrometer to make the large circle on the screen as large as possible. When our sample is in the real space plane, the numerical aperture is maximized. However, if the image begins to break apart into an inner circle and an outer ring, you've gone too far. It should be the biggest solid circle that can be formed.
19. The slit is now closed down, and the spectrometer adjusted to the wavelength of interest (e.g. 775 nm). The system is now ready to take images.

2.5.3 The 4-2 Reflectivity Routine

The 4-2 reflectivity routine is the first routine we found to work well. This process involves doing the full standard alignment twice, once for the mirror, and then again for the sample in the cryostat. The result is four sets of images; hence the name 4-2. The 4-2 routine is generally easy to execute well, but perhaps a bit slow. The 4-2 routine will not produce the most pristine data, but it will never produce bad data.

This routine uses two important algorithms from our software, the step-and-glue algorithm with the lock-signal-to-noise feature engaged, and the copy-prior-scan algorithm. These are described in Section 2.5.1.

This routine has the advantage that one does not have to align the sample or the mirror "blind," as in the 4-1 or 5-2 algorithms. However, the repeatability of the alignment comes into play. Extreme care must be taken when aligning the sample and mirror to be normal to the laser beam (when the reflected beam is aligned onto Apertures 1 and 2).

The process is as outlined below:

1. Perform the full standard alignment procedure with a good mirror in the sample plane.
2. Take the N images using the step-and-glue algorithm with the lock-signal-to-noise feature enabled.
3. Remove the mirror and let the white light terminate at infinity. Use the copy-prior-scan algorithm to take the B_N images using the N images as the template
4. Perform the full standard alignment procedure again, including removing all the optics, this time with the cryostat mounted in the sample plane. The digital crosshair does not need to be re-aligned.
5. Highly absorbing optical tape should have been placed inside the cryostat window ahead of time. Move the cryostat so the optical tape is being imaged. If tape was not placed in the cryostat window ahead of time, some can be placed in between the microscope objective and the cryostat. (In this case, the paper from the back side of the tape is not removed. Under no circumstances should adhesive from the optical tape ever touch the cryostat window.)
6. Use the copy-prior-scan algorithm with the N images as the template to take the B_S

7. Move the cryostat so the region of interest on the sample is in view. Use the copy-prior-scan algorithm to take the S images.
8. Use the Reflectivity-Calculator algorithm inside the Data-Analyzer GUI to process the data.

2.5.4 The 4-1 Reflectivity Routine

The 4-1 reflectivity routine is the hardest routine to successfully execute, but it has the potential for very pristine images as it avoids the problem of repeatability in the alignment procedure. However, the 5-2 reflectivity routine can achieve images that are just as good, and is significantly easier to execute. This processes involves doing the full standard alignment once with the mirror. Afterwards, the mirror is removed and the cryostat is moved into the sample plane. None of the optics are moved while aligning the cryostat. This makes it very difficult to determine if the cryostat is properly aligned. This is what we mean by aligning the system "blind."

After collecting data, it may be processed live in the lab in just a few minutes. The resulting reflectivity image will reveal if the system is misaligned or not. With bad luck, you may have to iterate multiple times to get the system into alignment.

The process is as outlined below.

1. Perform the full standard alignment procedure with a good mirror in the sample plane.
2. Take the N images using the step-and-glue algorithm with the lock-signal-to-noise feature enabled.
3. Remove the mirror and let the white light terminate at infinity. Use the copy-prior-scan algorithm to take the B_N images using the N images as the template
4. Move the cryostat into the sample plane. The focal micrometer may be adjusted to maximize the numerical aperture of the system, making the image on the screen as large as possible. The angle of the cryostat may be adjusted by flipping mirror 2 out of the optical path and looking for the back reflected laser on Aperture 1 or 2. This can be very difficult to find. Once aligned, flip Mirror 2 back into the optical path.

5. Highly absorbing optical tape should have been placed inside the cryostat window ahead of time. Move the cryostat so the optical tape is being imaged. If tape was not placed in the cryostat window ahead of time, place some in between the microscope objective and the cryostat (if you do this, do not remove the paper from the backside of the tape. Under no circumstances should adhesive from the optical tape ever touch the cryostat window).
6. Use the copy-prior-scan algorithm with the N images as the template to take the B_S
7. Move the cryostat so the region of interest on the sample is in view. Use the copy-prior-scan algorithm to take the S images.
8. Use the Reflectivity-Calculator algorithm inside our Data-Analyzer GUI to process the data.

2.5.5 The 5-2 Reflectivity Routine

The 5-2 reflectivity routine is a great compromise between the 4-1 and 4-2 routines. It has the potential to produce images that are just as pristine as the 4-1 routine, while also being much easier to execute, although still more difficult to execute than the 4-2 routine.

This routine involves first taking a set of reference images N_r. These reference images are only used as the template for the copy-prior-scan algorithm, and are not used in the final reflectivity produced. Therefore, if the system is not changing much from day to day, one may skip the first alignment in the procedure and just use a set of images from a prior day (in which case, this would be named the 5-1 routine). Afterwards, the full standard alignment procedure is performed on the sample and the S and B_S images are taken. Afterwards, the cryostat is removed and a mirror is moved into the sample plane. This mirror must be aligned blind similar to the 4-1 routine. However, because the mirror has light weight and is on a kinematic mount and micrometer stage (and also doesn't have the extra reflection of the cryostat window), it is far easier to align this mirror blind than it is to align the cryostat blind.

The process is as outlined below.

1. Perform the full standard alignment procedure with a good mirror in the sample plane.

2. Take the reference images N_r using the step-and-glue algorithm with the lock-signal-to-noise feature enabled.
3. Skip to this step if using images from a prior experiment for your reference images. Remove the mirror and necessary optics, and perform the full standard alignment procedure with the cryostat in the sample plane.
4. Highly absorbing optical tape should have been placed inside the cryostat window ahead of time. Move the cryostat so the optical tape is being imaged. If tape was not placed in the cryostat window ahead of time, place some in between the microscope objective and the cryostat (in this case, do not remove the paper from the backside of the tape. Under no circumstances should adhesive from the optical tape ever touch the cryostat window).
5. Use the copy-prior-scan algorithm with the N_r images as the template to take the B_S
6. Move the cryostat so the region of interest on the sample is in view. Use the copy-prior-scan algorithm with N_r as the template to take the S images.
7. Remove the cryostat, let the white light terminate at infinity and use the copy-prior-scan algorithm with N_r as the template to take the B_N images.
8. Move the mirror into the sample plane. The focal micrometer may be adjusted to maximize the numerical aperture of the system, making the image on the screen as large as possible. The angle of the mirror may be adjusted by flipping mirror 2 out of the optical path and looking for the back reflected laser on aperture 1 or 2. This can be very difficult to find. Once aligned, flip mirror 2 back into the optical path.
9. Use the copy-prior-scan algorithm with N_r as the template to take the N images.

2.5.6 Leveling The Stopband

As discussed above, the alignment of the optical system in the collection of our data is critically important, and not always trivial to accomplish. Often when measuring reflectivity off a sample in a cryostat, we see very peculiar behavior, such as stopbands that are not flat, but rather, have strange curves in them. I have sought to discover whether there may be a way to process the data to correct for misalignment present in the system during the collection of either the N. Specifically we are interested in correcting N images collected in the 5-2 reflectivity routine described in Section 2.5.5. In order to investigate this, we took some data in which we purposefully misaligned the system.

We began with our standard alignment procedure to take the N images off of a mirror. Then, we slightly adjusted the angle of the mirror by a few degrees and took those images again using the copy-prior-scan algorithm. We also took background images, and then calculated our relative reflectivity. In a properly aligned system, the relative reflectivity of a mirror to itself should be unity at all wavelengths and angles. Indeed, we found a uniform reflectivity near unity along with some statistical noise. Therefore we concluded the issue is probably not the angle of the sample or mirror.

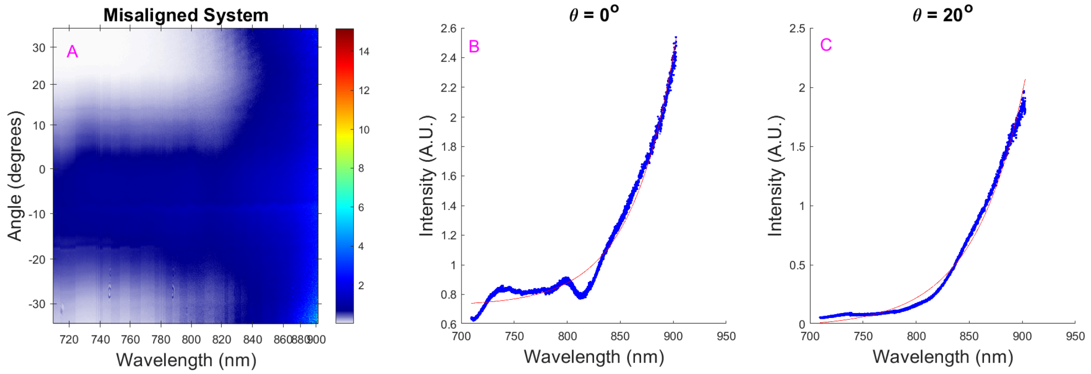

Figure 26: In image A we have the relative reflectivity of a mirror measured against itself, in which we have purposefully changed the focus of our optics system in between measurements. Image B and C show the profiles at $\theta = 0^o$ and $\theta = 15^o$

Next, we performed the standard alignment procedure again, and again took the N images off of a mirror. This time, rather than tilting the mirror, we moved it slightly out of focus using your micrometer stage, and we found a much more systematic error in our results. The relative reflectivity is shown in Figure 26A. Profiles at 0 and 15 degrees are shown in images B and C. We find these profiles can be reasonably fit by the form

$$y = A + Be^{\frac{P}{\tau}}, \tag{2.11}$$

where A, B, and τ are fitting parameters, and P is the horizontal pixel number. We notice that the fits are not the same at all the angles. Still, this seems to indicate that the bulk of the error stemming from focal misalignment may be divided out of the data if only we had the appropriate exponential curve. We can attempt to extract the curve by recognizing where the stopband is in a bad data set. One such bad reflectivity data set is shown in Figure 27A. Although the data is poor, we can clearly see where the stopband is approximately, we conservatively isolate this region of data and fit Equation (2.11) to it. We reason that clearly this region should be flat, so we can divide by our best fit. The fit is extrapolated to the entire spectrum, and divided out of it. The result is shown in Figure 27B.

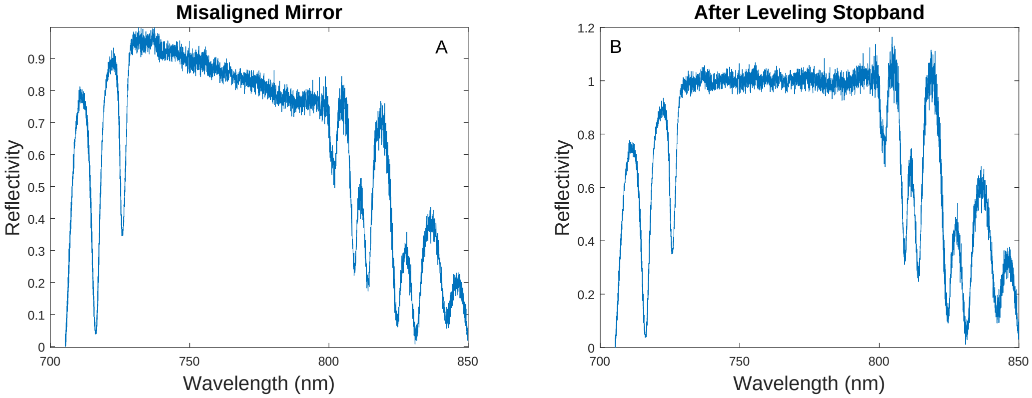

Figure 27: In image A we have a measured reflectivity curve in which some misalignment has occurred. In image B we attempt to fix this error by fitting the data in the stopband to equation 2.11, and dividing the result out of the entire curve.

Overall, it certainly does look a lot more like we expect. We know that when we perform reflectivity measurements on these samples, it is possible to get the pristine curves shown in Figure 24B. Therefore we think this method of fixing a measured reflectivity curve is possibly justified.

2.6 Active Stabilization

For our work in transition metal dichalcogenide monolayer (TMD) systems, we often found ourselves working with very small samples, i.e., flakes on the order of 10 to 100 microns [83, 80, 82]. This was especially prevalent in our work leading up to and including TMD quaternions [79, 81]. In these experiments, we often wished to perform some type of scan while monitoring PL from the sample; for example, electrical voltage scans and temperature scans. However, because the flakes were so small, there was a serious risk that some random movement would cause the flake to change its location on our screen and ruin our data. Standard precautions were taken, such as floating the optics table and cushioning the cryostat with movement dampening materials. But the cryostat is still ultimately hooked up to a vacuum pump in order to pull liquid helium through it, which is a source of vibrations.

Therefore I developed a system which allowed us to monitor and move the sample back into position in real time very quickly. I first had to modify our experimental setup to a dual imaging setup, with one image being used by the spectrometer and CCD camera to collect our PL data, and the other image being used by a second USB camera to monitor our sample's position. A simple illustration of this setup is shown in Figure 28.

Generally with our TMD work, we have a surplus of laser power. Therefore the achromatic Beamsplitter 2 may be chosen to be a 50-50 beamsplitter, and half our laser power simply terminates at the beam stop. The image exiting the backside of the microscope objective is collimated in real space. A portion of this beam is diverted back to Beamsplitter 2, and a portion of that beam is sent to the USB camera. As the beam is collimated in real space, it will form an image 1 focal length after passing through Lens 3, on the screen of the USB camera.

With this setup, we can monitor the sample's position while collecting data on the sample. I programmed the CCD camera and the USB camera to be controlled by two separate instances of MATLAB on the same computer, and so they may both be used in parallel. The MATLAB instances can also communicate through network protocols with each other, and so every time a CCD image is captured, a corresponding USB camera image is captured and stored with it.

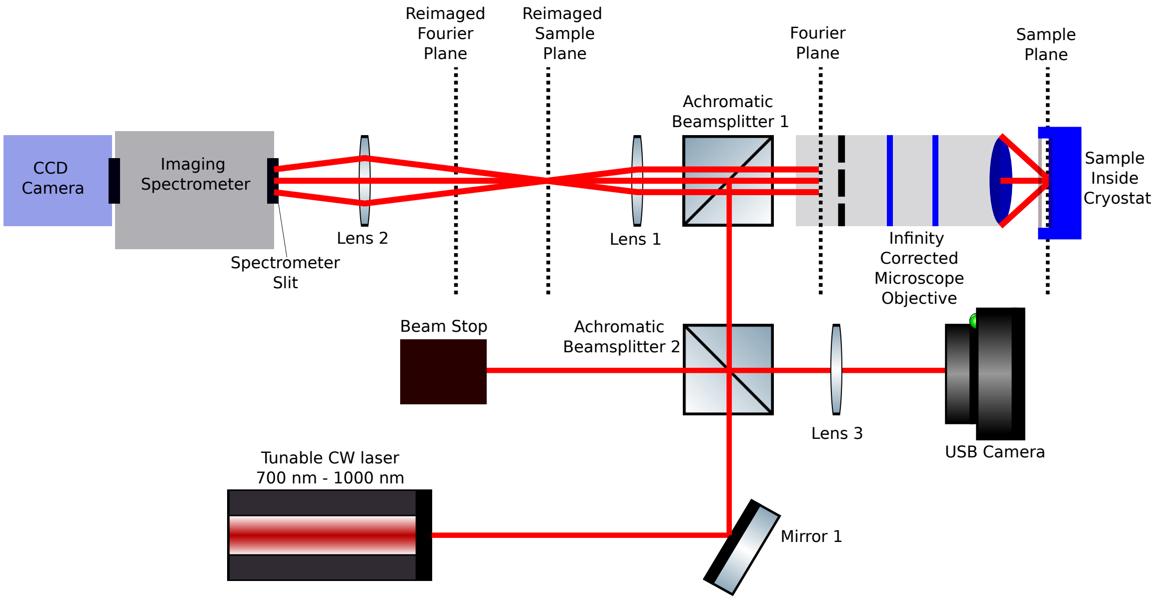

Figure 28: A simple illustration of our optics setup for dual imaging. Credit to Alexander Franzen who created the vector graphics library for optical components used in our figures [34].

The experiments involving the quaternions often involved lengthy periods of data acquisition because the signal to noise ratio was quite low compared to our GaAs samples. Additionally, because a shift of a few microns was problematic for us, a human being touching the optical table in order to fix the alignment, if the sample moved due to random vibration, often caused more problems than it solved. For this reason, we automated the repositioning of the sample, a process we've termed "active stabilization."

The cryostat is mounted on two electronic stepper motors, which handle the movement. Not pictured in the simple illustration of our optics setup is that the cryostat is mounted flat, and a lab jack is utilized for vertical motion to bring the sample into focus. This is significant because our electronic stepper motors are not strong enough to move the cryostat vertically, but horizontally is easy for them.

The major challenge I faced was how to program a computer to recognize when the sample had moved out of position due to a random fluctuation, and how to move it back into position. Fortunately, our samples have a well (enough) defined structure, especially the ones with gold conductive pads placed near them for electrical measurements. So then, the problem is to identify how much structures in an image have moved, and this is a well-studied problem. This may be efficiently accomplished using discrete Fourier transforms. For a more thorough analysis see Ref. 39. I utilized their published image registration library [38] for our work. The end result of their work is to tell us how many pixels the sample has moved. Often times, there are markers with numbers printed near our sample which may also be used to help with this process. If there are not, the process still works fine.

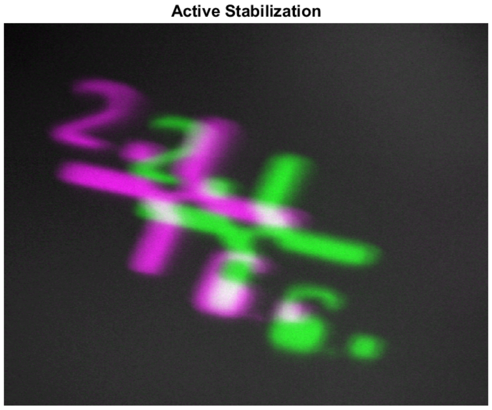

Figure 29: A gold pad which has been shifted by a random vibration in our system. The active stabilization system will recognize this, and move the sample back into alignment using the stepper motors.

Figure 29 shows an example of a shifted gold marker. The purple image is the original location of the gold marker, for which our system is configured, and we wish to take PL data of the sample (off screen here). The current position is shown by the green image. Our image registration library tells us that the green image is 86 pixels down and 113 pixels to the right of the purple image.

A calibration has to be made to convert pixels to micrometers, as our stepper motors operate in micrometers not pixels. Additionally, our stepper motors are not necessarily aligned perfectly with the horizontal and vertical axes of the USB camera. So, we are essentially doing a 2D vector basis transform. But this is easy enough. Utilizing the stepper motors, I moved the sample a known large amount in both the vertical and horizontal directions. Using the image registration library, I was then able to create the calibration between micrometers and pixels. Once we calculate how much to move our sample, we move it back into position.

The code uses a tolerance so that the sample is not moved constantly. The image registration requires about 2 seconds of computation time on our setup, so the instance of MATLAB running the USB camera (and stepper motors) is able to check for sample movement every 2 seconds. The stepper motors can then be tuned in about 3 more seconds. Therefore if sample movement occurs, it will be fixed within roughly 5 seconds. The PL data collected by our CCD camera takes place over the range of minutes, so these random movements will mostly get averaged out, leaving us with good PL data. In our experience, a random sample vibration occurs every 30 seconds to one minute. However, if the sample moves out of alignment, there is no PL data to collect, as the signal comes from a very small region of the sample. Therefore, active stabilization is important to facilitate these measurements.

2.7 Calibrating the Stress Needle

One of the major challenges in our pursuit of TMD polaritons was determining a method of adjusting the detuning, to bring an exciton into resonance with a cavity. Unlike our GaAs samples, which have centimeters of usable sample, our TMD samples have flakes the size of 10 to 100 microns. As a result, the cavity wedge we use in GaAs samples is not applicable to TMD polariton samples.

One avenue we explored was using strain to tune the the exciton energy through bandgap renormalization [80]. To do this, we built a sample mount to hold our TMD sample in place, while stressing it from behind by pressing a needle into its substrate. It is very easy to devise a stress assembly which will shatter a sample. This is because the sample itself moves very little. So, rather than attaching the needle to a micrometer, and pushing it forward, we instead kept it in position with pressure from a spring connected to a micrometer. By tuning the micrometer we were able to compress the spring, and increase the force applied by the needle, and the resulting strain.

We needed an accurate calibration of the force applied by the needle as a function of the micrometer setting. So we mounted the stress assembly to a board with a SHIMPO 20 Newton force gauge. By taking a series of measurement we were able to measure our spring constant. The exerted force calculated from this calibration was then used as input data for a finite element computation in the commercial ANSYS software, giving us the induced strain as a function of micrometer position.

3.0 Transfer Matrix Method Simulator

In Chapters 4 and 5, I will make extensive use of an electromagnetic simulation I coded in Matlab. This code is controlled with a series of graphical user interfaces (GUIs), and Appendix B serves as a software manual for those GUIs. Those GUIs are available online for download [12]; additionally I have published a series of videos on YouTube demonstrating their use [11]. In this Chapter, I focus on the underlying physics that goes into that simulation. In particular, Maxwell's equations are formatted into a transfer-matrix methodology. I also present an accurate derivation of Snell's law for materials with complex index of refraction, which is frequently misreported in various literature and online discussions.

I also will discuss efficient implementation of this simulation into Matlab. The difference between optimized and optimized code in this case is very drastic, and optimizing the code was an essential piece of making the simulation usable for the fits I perform in later chapters.

3.1 Complex Snell's Law

In this section I will derive the equations which describe light reflection and refraction at the boundary between mediums. To begin, I start with the complex wave equation for electromagnetic waves in an isotropic absorptive medium [75],

$$\nabla^2 \vec{E} = \frac{1+\chi_r+i\chi_i}{c^2}\frac{\partial^2 \vec{E}}{\partial t^2}. \tag{3.1}$$

Here, I have explicitly broken the complex susceptibility into real and imaginary parts. The complex index of refraction is given as $n = \sqrt{1+\chi_r+i\chi_i}$ where the square root is the typical complex extension of the square root. The complex nature of these quantities is what encodes into the equations that absorption is occurring. Now we consider this equation as it applies to a plane wave traveling through an absorptive medium. Specifically, it is easiest to consider this near an interface with a lossless material (no absorption).

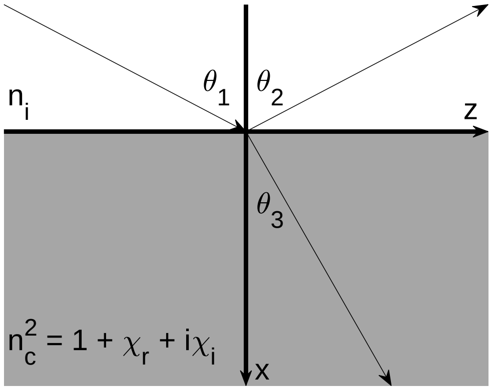

Figure 30: A simple illustration of a wave incident on a boundary between air and some absorptive medium.

A simple illustration of this scenario is shown in Figure 30, keeping the notation in line with Reference 19. Here, the top material is lossless and described with just a purely real and positive optical index n_i. The absorptive material underneath is described with a complex index n_c which may be related to the complex susceptibility χ_c as

$$n_c^2 = 1 + \chi_c = 1 + \chi_r + i\chi_i. \tag{3.2}$$

Slabs of material such as the one shown can be approximated as infinite slabs, and waves as infinite plane waves, which may be written as

$$\vec{E} = \vec{E_o} e^{-i(\vec{k}\cdot\vec{r} - i\omega t)}. \tag{3.3}$$

In this equation, $\vec{k}$ is complex valued, and the imaginary piece of it is ultimately what leads to a decaying amplitude due to the absorption; it can be explicitly broken into real and imaginary parts as $\vec{k} = \vec{k_r} + i\vec{k_i}$. The vector $\vec{E_o}$ determines the polarization. Typically one works with $\vec{E_o}$ either being entirely in the y direction (into and out of the page), which is a transverse electric wave (TE), or with the magnetic field being entirely in the y direction which is a transverse magnetic wave (TM). In a TM wave, the electric field is then in the x-z plane. So then in TM waves we tend to work with the B field instead, which are also then written as plane waves.

One can see by symmetry, there will only be decay while moving in the x- direction. This means that $\vec{k_i}$ may only have a component in the x direction. The value of $\vec{k_r}$ on the other hand can have both x and z components. Therefore, it must be the case that $\vec{k_r}$ and $\vec{k_I}$ are not necessarily parallel. It is of course possible for the incident wave to be normal to the boundary, in which case $\vec{k_r}$ will also be in the purely x direction.

Plugging the plane wave solution into the complex wave equation yields

$$\vec{k} \cdot \vec{k} = k_r^2 - k_i^2 + 2i\vec{k_r} \cdot \vec{k_i} = \frac{1 + \chi_r + i\chi_i}{c^2}\omega^2. \tag{3.4}$$

Next, the equation may be separated into real and imaginary parts, yielding

$$k_r^2 - k_i^2 = (1 + \chi_r)\frac{\omega^2}{c^2} \tag{3.5}$$

and

$$2\vec{k_r} \cdot \vec{k_i} = 2k_r k_i \cos\theta_3 = \chi_i \frac{\omega^2}{c^2}. \tag{3.6}$$

Here, the angle between $\vec{k_r}$ and $\vec{k_i}$ is θ_3 from the illustration. This is a system of equations which may be solved for k_r and k_i. When the dust settles we are left with

$$k_r = \frac{\omega}{\sqrt{2}c}\left(\sqrt{(1+\chi_r)^2 + \frac{\chi_i^2}{\cos^2\theta_3}} + (1+\chi_r)\right)^{1/2}, \tag{3.7}$$

and

$$k_i = \frac{\omega}{\sqrt{2}c}\left(\sqrt{(1+\chi_r)^2 + \frac{\chi_i^2}{\cos^2\theta_3}} - (1+\chi_r)\right)^{1/2}. \tag{3.8}$$

I will now use this result to derive the proper complex Snell's law. Going back to Figure 30, I now assign the incident wave the amplitude $\vec{E_1}$, the reflected wave an amplitude $\vec{E_2}$ and the transmitted wave an amplitude $\vec{E_3}$. The $\vec{k}$ for Waves 1 and 2 both have no imaginary component because they exist in lossless material; all three waves have a real component in the z direction. The notation here can get a bit messy. The components of the k-vectors will have three subscripts, in order referencing; ray number, real or imaginary component, then spatial dimension.

Even with absorptive materials, we still have the boundary conditions that the parallel components (y components) of the E field and H field are continuous at the boundary. In our case, it is a very good approximation that the magnetic permeability is equal to one, meaning the parallel components of the B field are also continuous. Taking the boundary to be where $x = 0$, for TE waves this means we have

$$\vec{E}_1 e^{-ik_{1rz}z} + \vec{E}_2 e^{-ik_{2rz}z} = \vec{E}_3 e^{-ik_{3rz}z}. \tag{3.9}$$

Where all the E-vectors are in the purely y direction. For TM waves, we instead work with the parallel component of the B field, leading to

$$\vec{B}_1 e^{-ik_{1rz}z} + \vec{B}_2 e^{-ik_{2rz}z} = \vec{B}_3 e^{-ik_{3rz}z}. \tag{3.10}$$

Here, all the B-vectors are in the purely y direction. However, in order for either of these equation to be true at any location z along the interface, we have the standard phase matching constraint, written as

$$k_{1rz} = k_{2rz} = k_{3rz}. \tag{3.11}$$

After this point, all the arguments for TE and TM are exactly the same. Meaning our final result for a complex Snell's law will not have a polarization dependence. For Rays 1 and 2, because they are in a medium with purely real optical index, this leads immediately to the standard law of reflection $\theta_1 = \theta_2$. However, the refracted ray is where things begin to change from the standard real Snell's law derivation. The quantity k_{3rz} may still be calculated as $k_{3r}\sin\theta_3$, where k_{3r} is the vector magnitude of the real component of ray three's k-vector. Because Ray 3 is inside an absorptive medium, we must use Equation (3.7) to calculate the magnitude, giving us

$$k_{3rz} = \sin\theta_3 \frac{\omega}{\sqrt{2}c}\left(\sqrt{(1+\chi_r)^2 + \frac{\chi_i^2}{\cos^2\theta_3}} + (1+\chi_r)\right)^{1/2}. \tag{3.12}$$

Ray 1 is inside a lossless material, and so we may easily calculate its z component as

$$k_{1rz} = k_{1r}\sin\theta_1 = \frac{\omega n_i}{c}\sin\theta_1. \tag{3.13}$$

Note that the standard law still applies that the frequency ω in both slabs is the same. One or both slabs of material having absorption does not change this. From the phase matching condition, these two equations are now set equal to each other. After some minor cancellations and rearranging this leads to

$$\frac{\sqrt{2}n_i \sin\theta_1}{\sin\theta_3} = \left(\sqrt{(1+\chi_r)^2 + \frac{\chi_i^2}{\cos^2\theta_3}} + (1+\chi_r)\right)^{1/2}. \tag{3.14}$$

This is an equation to be solved for θ_3 in terms of θ_1 and the material parameters. I won't write out all the algebra here, but I will describe the steps. First, it is squared, and then $1+\chi_r$ is moved to the left hand side, and it is squared again. Then, substitute $\cos^2\theta_3 = 1 - \sin^2\theta_3$. This is now a quadratic in $1/\sin^2\theta_3$. Utilizing the quadratic formula, we arrive at (after a lot of cancellations and simplifications)

$$\frac{2}{\sin^2\theta_3} = 1 + \frac{1+\chi_r}{n_i^2 \sin^2\theta_1} \pm \sqrt{\left(\frac{1+\chi_r}{n_i^2 \sin^2\theta_1}\right)^2 + \left(\frac{\chi_i}{n_i^2 \sin\theta_1^2}\right)^2}. \tag{3.15}$$

The plus or minus solution from the square root may be resolved. If $\chi_I = 0$, the material is now lossless and so the standard Snell's law must be recovered. This would also imply $1+\chi_r = n_3^2$. One can see the standard Snell's law falls right out if we select the positive root. The quantity inside the square root is strictly positive as this notation avoids complex angles. If it wasn't, it would be necessary to utilize the Sign function because $\sqrt{x^2} = |x|$ not just x. Thus, we are left with our final generalized Snell's law

$$\frac{2}{\sin^2\theta_3} = 1 + \frac{1+\chi_r}{n_i^2 \sin^2\theta_1} + \sqrt{\left(\frac{1+\chi_r}{n_i^2 \sin^2\theta_1}\right)^2 + \left(\frac{\chi_i}{n_i^2 \sin\theta_1^2}\right)^2}. \tag{3.16}$$

If desired, the inverse sine function may be utilized to solve for θ_3 explicitly. However, the normal Snell's law is usually left in terms of sines. Note that this equation was derived assuming the top slab of material was lossless, with a purely real optical index. This result has been reported before [26]; however, the knowledge does not seem to be widespread.

The phase matching condition $k_{1rz} = k_{3rz}$ may be used transitively for layers beneath the layer solved for here. The full simulation is for a layered sample in a cryostat. The data taken is rays of light in the air above that cryostat, and so simulating samples with

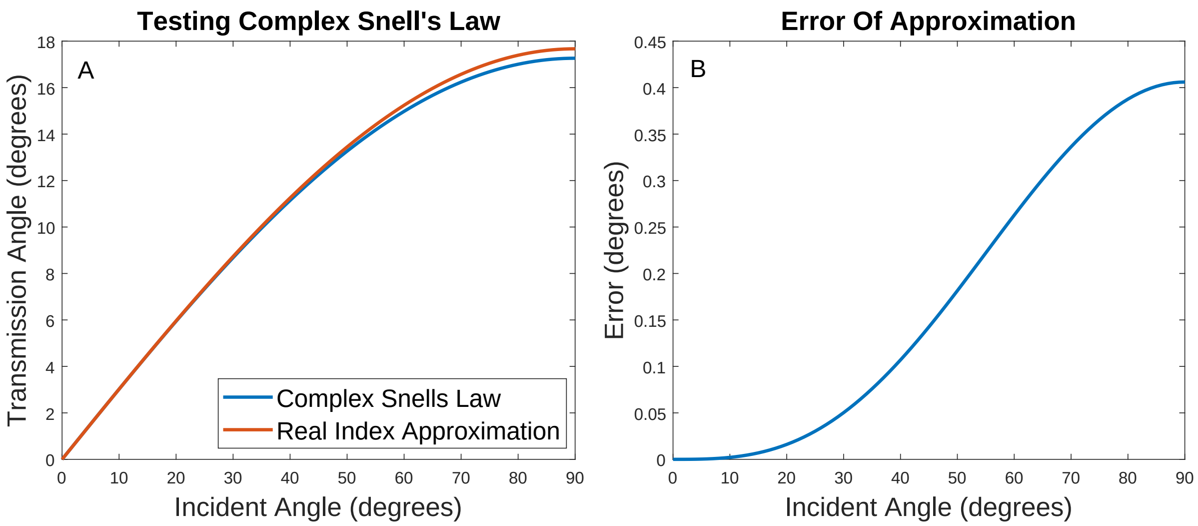

Figure 31: Testing the complex Snell's law and comparing it to the real index approximation

an air layer over top of them is not a problem. If one were working with an oil immersion microscope, some additional considerations might be required here.

Lets now look at an example of this equation in action. Consider the top slab of material to be vacuum, $n_i = 1$, with the bottom slab of material being GaAs at a temperature of 5 K, and the vacuum wavelength of the light equal to 775 nm. In my simulations there are times where I've considered using angle dependent index functions. For example, in Section 5.5, I discuss alternatives to using the light-hole exciton in the simulations. One of these options was allowing the excitons inside the GaAs quantum wells to have a finite mass, which would make the index of refraction of the quantum wells angle-dependent. In this case, we can just use Ref. 37 to calculate the non-angle-dependent complex index of bulk GaAs, giving us $n_c = 3.2944 + 0.0288i$. This may be converted into an electric susceptibility resulting in $\chi_r = 9.8521$ and $\chi_i = 0.1899$. Using Equation (3.16), the angle of the transmitted ray may be calculated over a range of initial angles. These results are shown in Figure 31.

Additionally, I have plotted what I've termed the "real index approximation." In this approximation, only the real part of the index is utilized along with the standard Snell's law to calculate the transmission angle. One can see from Figure 31A that the difference between the complex Snell's law and this approximation is quite minimal, especially at lower angles of incidence. The error between the true and approximate calculation is shown in Figure 31B.

Another approximation one finds commonly when searching online for how to handle this situation is to use the standard Snell's law with the complex index of refraction. This ultimately makes use of the inverse sine function extended to complex arguments, resulting in a complex angle. The transmission angle is then simply the real part of this complex angle. The results of this approximation are almost exactly identical to the results of the real index approximation.

These approximations are so close to the true calculation, it is no wonder that they have spread through the community as the proper way to handle refraction with complex indexes. However, the hard work of calculating the complex Snell's law is already done, so in my simulation I do not use either of the two approximations. Furthermore, my hope is for my TMM simulator to be broadly useful to other students with other projects, that is why it has been conveniently coded into a user friendly GUI available online.

3.2 The TMM Equations - TE

In this section, I will derive the equations for the transfer matrix method for transverse electric (TE) waves. These are electromagnetic waves where the electric field is always parallel to the interface, even when the ray is at oblique incidence. In Figure 32 the electric field is oscillating into and out of the page. This derivation is covered well in Reference 19; however, they skip the transverse magnetic (TM) case in favor of the duality principle. The duality principle meaning that they say to swap E with B and ϵ with μ. However, there are some additional precautions to take when applying the duality principle in this case. So, I will first derive the TE equations with a focus tailored towards my purposes, then I will expand upon the reference material and derive the TM equations.

A simple illustration of the problem at hand, defining some of the essential variables, is given in Figure 32. A critical assumption is that the materials are all isotropic. Of course, the materials are actually made up of a lattice of atoms, and so we would expect that in the most general case that rays of light reflect slightly differently based on orientation relative to that crystal lattice. However, the end result of my TMM simulation under the isotropic assumption seems reasonably accurate.

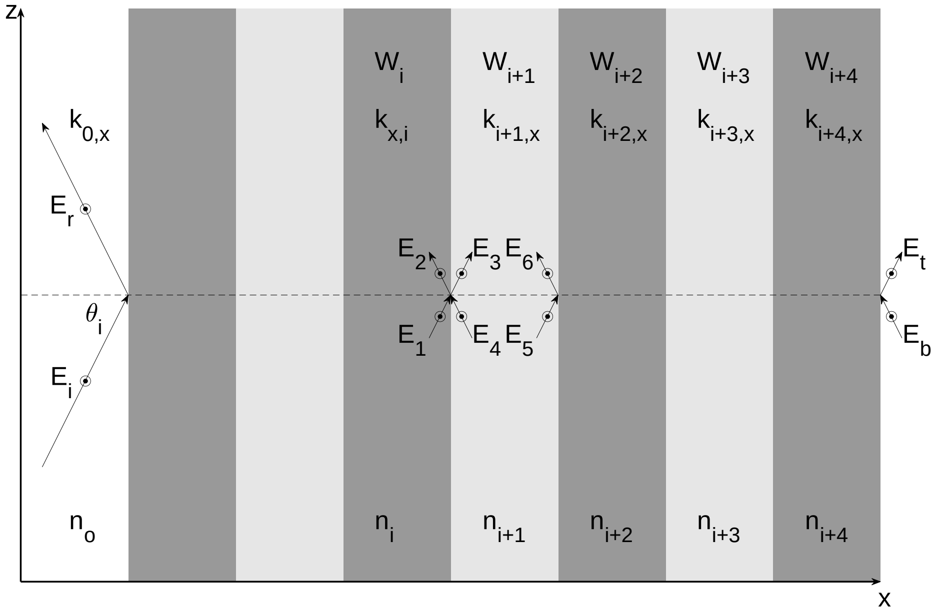

Figure 32: A simple illustration of the transfer-matrix method problem. The polarization is represented by the circles, which are arrow heads out of the page.

There is a TE wave with amplitude E_i of known frequency ω incident on the topmost layer of the layered medium, at some known angle θ_i. The initial medium has known index n_o, typically air. The widths W of all the layers are known, as are the indexes of refraction of all the layers. There is a reflected wave with amplitude E_r and a transmitted wave with amplitude E_t. The wave incident on the backside of the sample has an amplitude E_b which will ultimately be set to zero at the end, but is kept for consistency. The task at hand is to solve for the complex reflection coefficient $r = E_r/E_i$ and transmission coefficient $t = E_t/E_i$.

First, notice that once again there is a phase matching condition at every boundary, and so all the waves will have the same Z component for their k-vectors. This may be calculated for the incident wave to obtain $k_{z,o}$, and so it is calculated for all layers.

Using the complex Snell's law from equation (3.16), the angle of every ray of light may be calculated throughout the structure. Then the magnitude of the real and imaginary parts of the k-vectors k_r and k_i may be calculated using Equations (3.7) and (3.8) for every layer of the material, where the susceptibility has been calculated by squaring the complex index. Recall, as argued above by symmetry, that $\vec{k}_I$ is purely in the x-direction. The x component of $\vec{k}_r$ in every layer may be calculated using the already calculated angle as $k_{r,x} = k_r \cos\theta$. The sign is opposite for the rays moving to the left.

So then, the entire k-vector may be calculated throughout the entire structure for all the rays of light. Of particular interest for the TMM equations will be the x component of the i^{th} layer k_{ix}, which is complex valued in general.

Next, I relate the amplitudes of the waves from one layer to the next. This will then be applied throughout the entire structure to relate E_i and E_r to E_t and E_b. To do so, we first consider the boundary between the i^{th} layer and the $i+1^{th}$ layer as indicated in Figure 32. The boundary conditions in Maxwell's equations are that the parallel components of the E and H fields are continuous. Inside our semiconductors it is a very good approximation to set the magnetic permeability equal to one. In this case, the parallel component of the B field is continuous. For the electric field, taking $x = 0$ to be the boundary, we have

$$E_1 e^{ik_z} + E_2 e^{ik_z} = E_3 e^{ik_z} + E_4 e^{ik_z}. \tag{3.17}$$

This reduces down simply to

$$E_1 + E_2 = E_3 + E_4. \tag{3.18}$$

Next, the

$$\vec{k} \times \vec{E} = \omega \vec{B}. \tag{3.19}$$

One can see that because the TE wave has an electric field in the y direction, and $\vec{k}$ lies in the plane, B will have z and x components. The z component of B will be continuous, and may be calculated simply as $B_z = k_x E_y / \omega$. The requirement of continuity then leads to

$$k_{i,x}(E_1 - E_2) = k_{i+1,x}(E_3 - E_4). \tag{3.20}$$

Here, I have divided out ω because it is the same in every layer, and recognized that the forward and backward rays in the same layer have equal and opposite x components in their k vectors. Rearranging this leave us with

$$(E_1 - E_2) = \frac{k_{i+1,x}}{k_{i,x}}(E_3 - E_4). \tag{3.21}$$

In keeping with Reference 19, I introduce the definition

$$P_{i,TE} = \frac{k_{i+1,x}}{k_{i,x}}. \tag{3.22}$$

It is critically important to be aware that in Section 3.3, when I derive the equations for TM polarization, the definition of P will change. The definition here for P only applies to TE polarization. For this reason, I have placed a subscript TE on the P variable in the above definition. However, I will drop that subscript moving forwards as all the equations for TM polarization are essentially the same after this point.

Placing P into the prior equation now results in

$$E_1 - E_2 = P_i E_3 - P_i E_4. \tag{3.23}$$

Utilizing equation (3.18) to solve for E_1 and E_2 we have

$$E_1 = \frac{1 + P_i}{2} E_3 + \frac{1 - P_i}{2} E_4 \tag{3.24}$$

and

$$E_2 = \frac{1 - P_i}{2} E_3 + \frac{1 + P_i}{2} E_4, \tag{3.25}$$

which may be recast in matrix form as

$$\begin{pmatrix} E_1 \\ E_2 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1 + P_i & 1 - P_i \\ 1 - P_i & 1 + P_i \end{pmatrix} \begin{pmatrix} E_3 \\ E_4 \end{pmatrix}. \tag{3.26}$$

This is the first piece needed to solve this problem. We have written propagation through a boundary in matrix form. Next, we need to handle propagation through the slabs. We recognize that E_5 and E_6 may be simply related to E_3 and E_4, as they are the same waves just propagated a distance W_{i+1} through layer i. Cast in matrix form, this is written as

$$\begin{pmatrix} E_3 \\ E_4 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} e^{-ik_{i+1,x}W_{i+1}} & 0 \\ 0 & e^{ik_{i+1,x}W_{i+1}} \end{pmatrix} \begin{pmatrix} E_5 \\ E_6 \end{pmatrix} . \tag{3.27}$$

Now, by combining these two matrix equations, E_1 and E_2 are directly related to E_5 and E_6, written as

$$\begin{pmatrix} E_1 \\ E_2 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1+P_i & 1-P_i \\ 1-P_i & 1+P_i \end{pmatrix} \begin{pmatrix} e^{-ik_{i+1,x}W_{i+1}} & 0 \\ 0 & e^{ik_{i+1,x}W_{i+1}} \end{pmatrix} \begin{pmatrix} E_5 \\ E_6 \end{pmatrix} . \tag{3.28}$$

This can be simplified down to

$$\begin{pmatrix} E_1 \\ E_2 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} (1+P_i)e^{-ik_{i+1,x}} & (1-P_i)e^{ik_{i+1,x}} \\ (1-P_i)e^{-ik_{i+1,x}} & (1+P_i)e^{ik_{i+1,x}} \end{pmatrix} \begin{pmatrix} E_5 \\ E_6 \end{pmatrix} . \tag{3.29}$$

Now, I define the propagation matrix for layer i as B_i as

$$B_i = \begin{pmatrix} (1+P_i)e^{-ik_{i+1,x}} & (1-P_i)e^{ik_{i+1,x}} \\ (1-P_i)e^{-ik_{i+1,x}} & (1+P_i)e^{ik_{i+1,x}} \end{pmatrix} . \tag{3.30}$$

I adopt the more generalized notation that the amplitude of the forward-propagating wave in layer i at the right most edge is written as $E_{f,i}$ and the backward-propagating wave is written as $E_{b,i}$. This means that $E_1 = E_{f,i}$ and $E_2 = E_{b,i}$. So, our matrix equation is now more generally written as

$$\begin{pmatrix} E_{f,i} \\ E_{b,i} \end{pmatrix} = B_i \begin{pmatrix} E_{f,i+1} \\ E_{b,i+1} \end{pmatrix} . \tag{3.31}$$

Now, through repeated application of this matrix multiplication, the final desired result relating E_i and E_r to E_t and E_b is within our grasp. We have the matrix equation

$$\begin{pmatrix} E_i \\ E_r \end{pmatrix} = B_o B_1 B_2 ... B_N \begin{pmatrix} E_{f,N} \\ E_{b,N} \end{pmatrix} . \tag{3.32}$$

This propagates us all the way to the final layer of the material, at the right most boundary of the entire sample. However, the amplitudes are still within the sample. Now, I could just use the matrix equation I derived above to get to the right side of the final boundary. However, a simple trick I utilized in my code was to add an extra layer of air to the bottom of the sample. The thickness of this layer isn't important, but it can be zero. This adds one more B matrix into the equation which I label B_{air}, leading to

$$\begin{pmatrix} E_i \\ E_r \end{pmatrix} = B_o B_1 B_2 ... B_N B_{air} \begin{pmatrix} E_t \\ E_b \end{pmatrix}. \tag{3.33}$$

Now, all the matrices in the middle are multiplied together, resulting in a final transfer matrix M for the sample, written out this is

$$\begin{pmatrix} E_i \\ E_r \end{pmatrix} = \begin{pmatrix} M_{1,1} & M_{1,2} \\ M_{2,1} & M_{2,2} \end{pmatrix} \begin{pmatrix} E_t \\ E_b \end{pmatrix}. \tag{3.34}$$

Now we are finally able to utilize $E_b = 0$, which implies $E_i = M_{1,1}E_t$ and $E_r = M_{2,1}E_t$. Immediately we have the complex transmission coefficient $t = 1/M_{1,1}$. These two results may be divided by one another, to end at

$$\frac{E_r}{E_i} = r = \frac{M_{2,1}}{M_{1,1}}, \tag{3.35}$$

where r is the complex reflection coefficient for the whole system. The power reflectivity (the measurement discussed in Section 2.5) is now calculated as $R = |r|^2$.

Using these equations, the problem of solving Maxwell's equations has been conveniently turned into matrix algebra. In a typical case, one wants to perform this procedure at many different angles and wavelengths of light. Efficiently implementing an algorithm to do that is important and a bit tricky. I discuss this in Section 3.4.

3.3 The TMM Equations - TM

I will now perform the same derivation from Section 3.2, except for transverse magnetic (TM) waves. In a TM wave, the magnetic field is always parallel to the boundary at the interfaces. In Figure 32 the letter E is replaced with the letter B in each location. As I explained above, I am assuming the relative magnetic permeability μ is equal to one. This is a good approximation in our semiconductors. Looking at the right boundary of layer i leads immediately to

$$B_1e^{ik_z} + B_2e^{ik_z} = B_3e^{ik_z} + B_4e^{ik_z}. \tag{3.36}$$

If I were deriving Snell's law, B_3 would be zero. The phase- matching condition falls directly out once again. This is why Snell's law is not polarization-dependent for isotropic materials. Dividing out the phased-matched exponential terms leaves us with

$$B_1 + B_2 = B_3 + B_4. \tag{3.37}$$

Next, the E-field must be calculated using

$$\vec{E} = \frac{c^2}{\omega n^2}\vec{k} \times \vec{B}. \tag{3.38}$$

Here, n is complex-valued. One can see that E will thus have components in the x and z direction. The boundary condition applies only to the parallel component, so we need only calculate the z component of E,

$$E_z = \frac{c^2}{\omega n^2}k_xB_y. \tag{3.39}$$

Calculating this for all four waves and imposing the continuity condition leads to

$$\frac{k_{i,x}}{n_i^2}(B_1 - B_2) = \frac{k_{i+1,x}}{n_{i+1}^2}(B_3 - B_4). \tag{3.40}$$

Here, I have divided out ωc^2 because it is the same in every layer, and recognized that the forward and backward rays in the same layer have equal and opposite x components in their k-vectors. Rearrange this leave us with

$$B_1 - B_2 = \frac{n_i^2k_{i+1,x}}{n_{i+1}^2k_{i,x}}(B_3 - B_4). \tag{3.41}$$

Just as in Section 3.2 when working with TE waves, I now introduce the definition

$$P_{i,TM} = \frac{n_i^2 k_{i+1,x}}{n_{i+1}^2 k_{i,x}}. \tag{3.42}$$

Comparing this to the definition of P for TE waves in Equation (3.22), a very important difference is present. In the TE waves, the optical index n did not show up in the definition directly. But here, it does. In many electromagnetic books, the reader is told to swap E and H and to swap μ and ϵ in these sorts of derivations in order to avoid deriving the TM case. That is, they are told to utilize the duality principle. However, because I assumed $\mu = 1$ before arriving at Equation (3.22), the duality principle would not yield the correct result. Caution must be exercised to ensure earlier assumptions do not contradict this duality principle trick.

From here forward, all the equation manipulations are the same as the TE case, except the value of P is different, so I will not write them out as thoroughly. Once again, I drop the TM subscript from the P. The method of relating B_3 and B_4 to B_5 and B_6 is exactly the same, just introducing an exponential term, and so I once again define the propagation matrix for layer i as

$$B_i = \begin{pmatrix} (1+P_i)e^{-ik_{i+1,x}} & (1-P_i)e^{ik_{i+1,x}} \\ (1-P_i)e^{-ik_{i+1,x}} & (1+P_i)e^{ik_{i+1,x}} \end{pmatrix}. \tag{3.43}$$

All the same steps lead to

$$\begin{pmatrix} B_i \\ B_r \end{pmatrix} = B_o B_1 B_2 ... B_N B_{air} \begin{pmatrix} B_t \\ B_b \end{pmatrix}. \tag{3.44}$$

Here the multiple uses of the letter B is unfortunate, but I am trying to keep my definitions in line with Reference 19. The B's inside the vectors are the field magnitudes, and the B's outside the vectors with an index subscript are the propagation matrices for each individual layer given by equation (3.43). Just as with TE waves, these matrices may be multiplied together, leaving us with

$$\begin{pmatrix} B_i \\ B_r \end{pmatrix} = \begin{pmatrix} M_{1,1} & M_{1,2} \\ M_{2,1} & M_{2,2} \end{pmatrix} \begin{pmatrix} B_t \\ B_b \end{pmatrix}. \tag{3.45}$$

As before, I assume no incident ray from the right side of the sample, meaning $B_b = 0$, and so through the same algebra I arrive at

$$\frac{B_r}{B_i} = r = \frac{M_{2,1}}{M_{1,1}}, \tag{3.46}$$

where the power reflectivity may be calculated as $R = |r|^2$. So then, one can see that the calculation for TE and TM waves is nearly identical, with the sole exception that for TE waves, P is given by equation (3.22) and for TM waves P is given by Equation (3.42).

3.4 Efficient Implementation Of TMM In Matlab

The equations I derived in the previous sections were used to program a TMM simulator which I use in Chapters 4 and 5 to fit models to data. Additionally, the simulator is utilized to predict and design samples, including TMD samples discussed in Chapter 6. My simulator is available for download [12] along with videos showing how to use it [11]. This code is conveniently wrapped up inside of a GUI for the user. Specific details about how to utilize that GUI are given in Appendix B which acts as a software manual. This section however, will focus on giving a broader overview of the code, and discussing some of the breakthroughs I made in improving the efficiency.

As I have mentioned, computation time for these simulations can be problematic. However, the computation time when efficiently implemented in MATLAB is a fraction of the time when inefficiently implemented. The equations derived in the previous sections are not particularly cumbersome. However, I am not only looking to simulate the reflectivity of our samples at a single wavelength at a single angle. I am interested in simulating hundreds of wavelengths at dozens of angles all simultaneously. Furthermore, I am interested in doing that thousands of times in order to fit simulation results to experimental results.

MATLAB is a programming language that heavily favors formatting tasks into larger matrix operations rather than things such as nested for loops. This has to do with some of the mechanisms going on inside of the software. For this reason, a lot of my MATLAB code is formatted in terms of very large matrices that are a bit difficult to understand. I've done my best to include comments in the code to help make it more readable.

Of course, I also had the option of writing the simulator in a different programming language. However, a goal of mine was to leave this software in a state where other students in the lab would have the ability to read it and modify it as needed for their own purposes. MATLAB is generally easier to read for beginning programmers, fairly common for undergraduates to learn, and perhaps most relevant to me, it has very nice built in tools for the construction of graphical user interfaces (GUIs). In general, training a new student on how to use a GUI is much easier than teaching them to modify portions of code. So, MATLAB was the language I decided to use.

The majority of the variables are three-dimensional arrays, such as $n_{i,j,k}$. This is the complex optical index. The first index refers to angle, the second to wavelength, and the third is the layer within the structure. However, the code equivalent of the B matrices in Equations (3.30) and (3.43) are four-dimensional, in which the fourth index is the position in the matrix, expressed as

$$B = \begin{pmatrix} B_1 & B_3 \\ B_2 & B_4 \end{pmatrix} . \tag{3.47}$$

Below is a simplified breakdown of what the code does. The user specifies the wavelengths and angles for the simulation. Additionally, the structure of the sample is specified by loading in a structure file (conveniently made with another one of my GUIs, see Appendix B). Note, the first layer and last layer of the sample have thin layers of vacuum added on to them to make calculations simpler. For simplicity here, I will only include the steps for the TE waves.

1. Use user specified angles and wavelengths to calculate k_{rz} the z component of k_r in the first vacuum layer.
2. Evaluate the optical index function at all wavelengths and angles in each layer and construct the 3d array N.
3. Use N along with Equation (3.16) to calculate the angle of the rays for each wavelength and angle inside each layer.

4. Use Equation (3.7) to calculate k_r, the magnitude of the real part of the k vector for all wavelengths and angles inside each layer.
5. Use Equation (3.8) to calculate k_i, the magnitude of the imaginary part of the k vector for all wavelengths and angles inside each layer.
6. Calculate the x component of k_r as $k_{r,x} = \sqrt{k_r^2 - k_{rz}^2}$. Here, we utilize the phase matching constraint. The value of k_{rz} is preserved across boundaries.
7. Use equation (3.22) to calculate P for the TE case for all wavelengths and angles inside each layer
8. Use Equation (3.30) to construct the four dimensional B array. The first three indices are angle, wavelength, and layer. The fourth is the position as defined in Equation 3.47.
9. Preallocate the four dimensional array M. Begin with $M = B$
10. Step through one layer at a time, and explicitly perform the two by two matrix multiplication between the current layer of M and the next layer of B for all wavelengths and angles at once. Example code is given below (note the use of Matlab's dot multiplication). This vectorization scheme improved the performance of the simulation massively.

    ```
    for j = 1:numer_of_layers
    M(:,:,j,1)=...
    B(:,:,j,1).*M(:,:,j+1,1)+B(:,:,j,3).*M(:,:,j+1,2);
    M(:,:,j,2)=...
    B(:,:,j,2).*M(:,:,j+1,1)+B(:,:,j,4).*M(:,:,j+1,2);
    M(:,:,j,3)=...
    B(:,:,j,1).*M(:,:,j+1,3)+B(:,:,j,3).*M(:,:,j+1,4);
    M(:,:,j,4)=...
    B(:,:,j,2).*M(:,:,j+1,3)+B(:,:,j,4).*M(:,:,j+1,4);
    end
    ```

    Credit to my personal friend Benjamin Liu who helped think up this vectorization scheme.
11. The final slice of the four dimensional array M is the transfer matrix M in equation (3.34).

Not included in the example code is the calculation of the E and B fields inside of each layer. This is an optional part of the code that only runs if the user requests it. The four dimensional array may be used to quickly calculate the fields inside each layer of the sample. An example of this is shown in Figure 33.

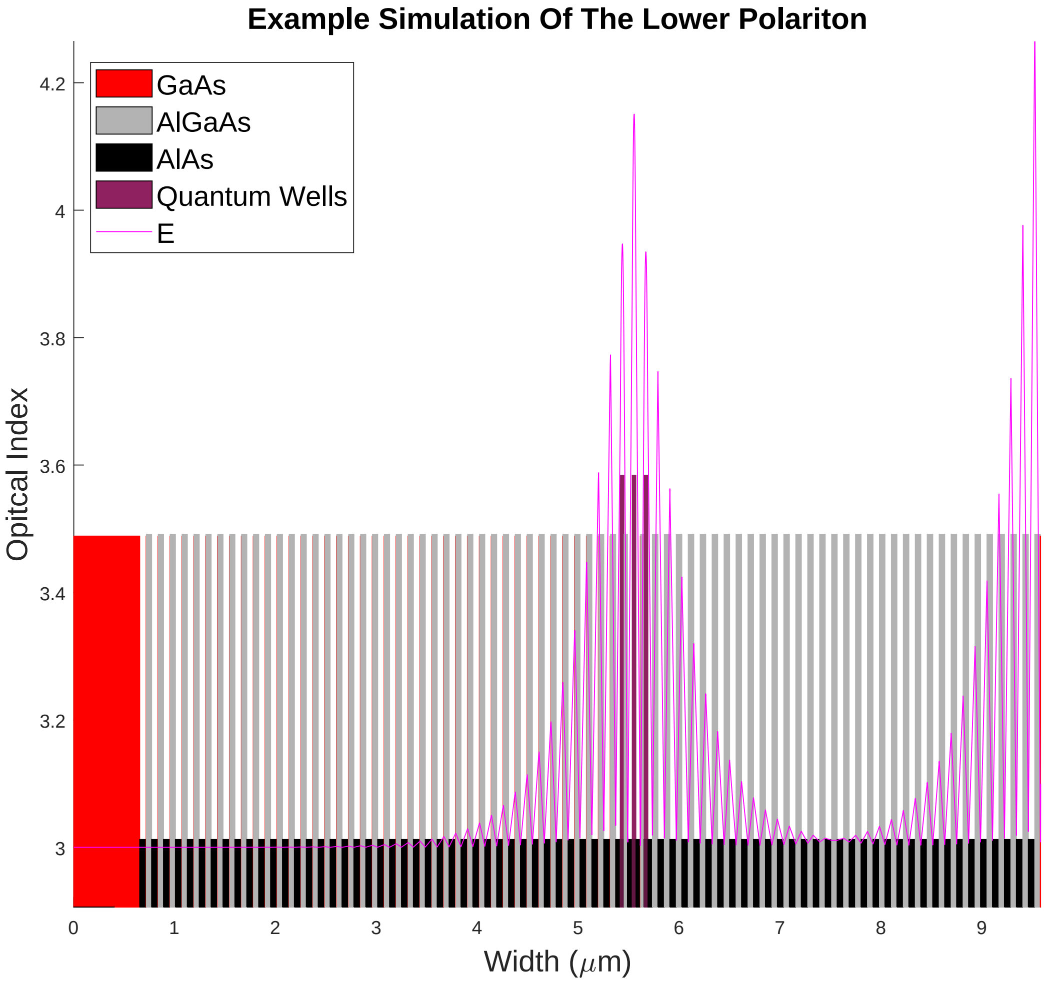

Figure 33: The electric field inside of a long-lifetime sample, simulated at the lower polaritons wavelength.

It is interesting to note how much the electric field leaks out of the microcavity. As discussed in Section 1.6, one can think of the design as either a $3\lambda/2$ design or a $\lambda/2$

design in which additional quantum wells have been placed outside of the microcavity, in the mirrors. The electric field leaks so far out of the cavity, it might be possible to create even stronger coupling by adding more quantum wells. Some groups have taken this to the limit, and incorporated many more quantum wells [5]; however they were working with Bragg-polaritons instead of microcavity polaritons.

4.0 Long-lifetime Sample Characterization

I will now present the main goal of this work, the characterization of our long-lifetime microcavity samples. Here I present a careful calibration of the exciton fraction of polaritons in high-Q ($\sim 300,000$), long-lifetime (~ 300 ps), GaAs/AlGaAs microcavities. Although we have made estimated of this parameter, we have never had an exact calibration of this crucial parameters, which is critically important for many calculations including polariton-polariton interaction strength [78, 72] as well as getting accurate measurements of the absolute polariton density from photoluminescence (PL) [58, 77, 17, 3]. It is much harder to establish this number in high-Q structures compared to low-Q structures, because the upper polariton is nearly invisible in high-Q cavities. This is because in long-lifetime samples the linewidths of the polariton branches are so narrow as to be unresolvable in reflectivity spectra. Furthermore, the upper polariton's PL is so weak as to be unobservable in many regions of interest. A simple way to understand why the upper polariton PL is so weak is to think in terms of the branching ratio for decay processes of the upper polaritons. In short-lifetime samples, emission of PL photons from the upper polariton states occurs at about the same rate as emission of phonons to jump down into a lower polariton state, but in long-lifetime samples, jumping down into lower polariton states occurs much faster than photon emission from the upper polaritons.

This makes accurate determination of the exciton fraction very difficult. I will utilize a combination of photoluminescence, photoluminescence excitation, and reflectivity measurements to highly constrain an electromagnetic simulation of the sample utilizing the transfer matrix method. In Chapter 5, I will apply these methods to a low-Q sample to confirm they are valid. My main goal is a fitted curve of exciton fraction as a function of the lower polariton energy which I may share with our collaborators. This analysis was performed on many samples and a table of the results is provided in Table 4.

4.1 Sample Structure

The structures used in our experiments are created by putting two highly reflective distributed Bragg reflectors (DBRs) in close proximity with each other in order to form a standing wave of light with 3 antinodes. At the antinodes of the standing wave there are a series of quantum wells. This design results in large electric fields interacting with the excitons inside the quantum wells. The entire sample is grown through molecular-beam epitaxy (MBE) by our collaborators, the Loren Pfeiffer group at Princeton University and the Zbig Wasilewski group at the University of Waterloo. A simplified but to scale representation of the sample is shown in Figure 34.

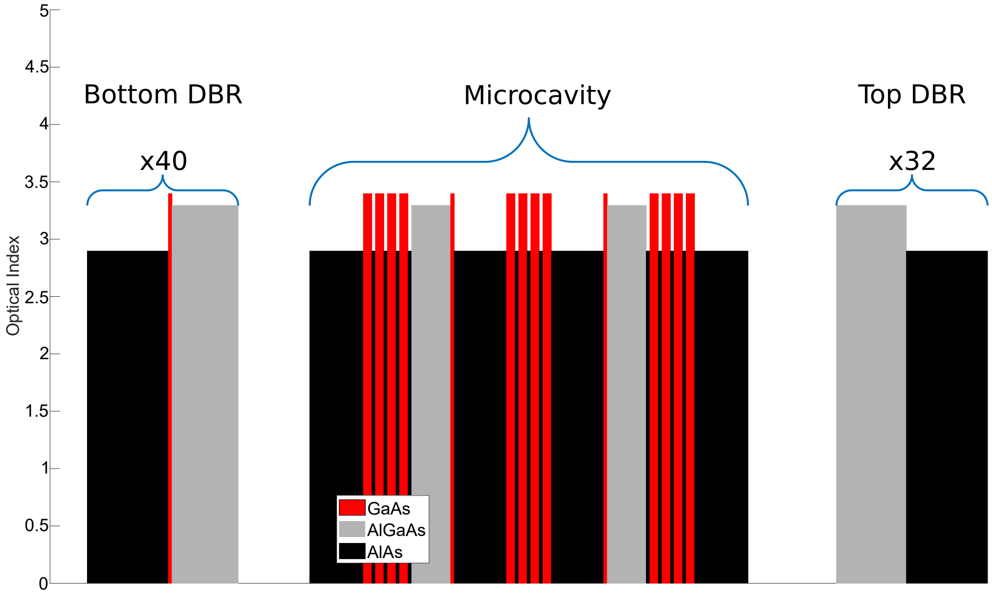

Figure 34: A simplified to scale illustration of our long-lifetime sample. The widths are the physical thicknesses of the design, and the heights are the optical index at relevant conditions.

In that illustration, the widths of the bars are to scale with our real design, and the heights are the optical index evaluated at 775 nm at cryogenic temperatures. As I have mentioned, the DBRs are created with alternating layers of high index and low index material, where the thicknesses of the materials are chosen such that each slab has an optical path length equal to one quarter of the design wavelength. The AlGaAs is a mixed alloy properly written as $Al_xGa_{1-x}As$. In my case, x is 0.2 approximately.It can also be interesting to show such a plot replacing the vertical axis with the band gap values from 2, as shown in figure.

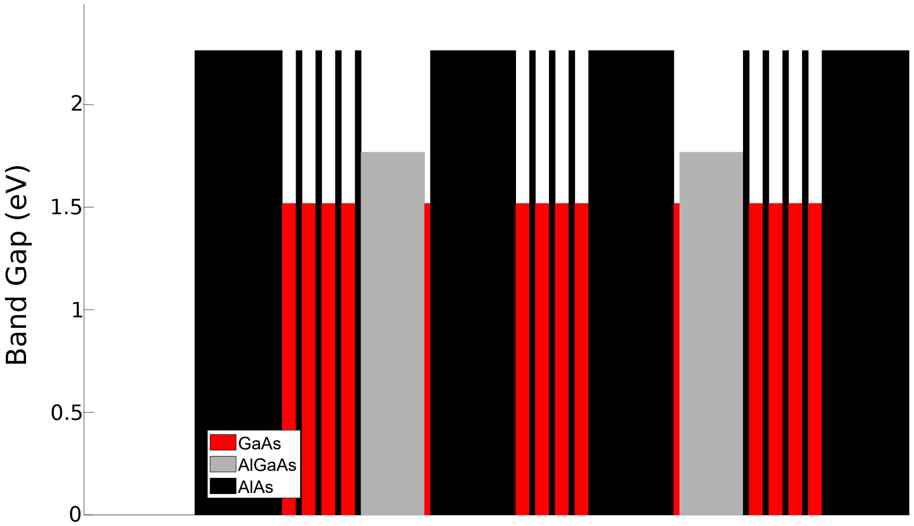

Figure 35: A simplified to scale illustration of our long-lifetime sample's microcavity. The heights are the band gap of our materials. We see the quantum confinement which modifies our GaAs excitons.

We see that the AlAs layers surrounding the thicker GaAs layers (the quantum wells) will act like the potential barriers in the well known 1D particle in a box problem common in many introductory quantum texts [50]. However, the wetting layers which are asymmetrically surrounded by AlGaAs and AlAs are also quantum wells. However, because they are so much thinner, their excitons will generally be at much higher energies, and we will not need to worry about them.

Not pictured in these simple illustrations is the substrate. This sample was grown on a GaAs wafer. Additionally, there is always a thin layer of GaAs placed on the very top, typically 3 nm thick. The capping layer, which helps prevent oxidation, as AlAs is susceptible to oxidation.

4.2 Photoluminescence Measurements

The first measurement is a very straightforward PL measurement, done by pumping the sample with a laser at an energy higher than the excitons. This creates many excited states (hot carriers) which scatter down through phonon scattering and become polaritons. The polaritons eventually annihilate and release a photon out of the cavity, which ultimately leaves the sample entirely and is collected by the optics. One careful consideration is what wavelength to choose for the laser. Because of the top DBR in the structure, some wavelengths will be able to enter the cavity and create hot carriers much more efficiently. Figure 36 shows a typical reflectivity curve on the photonic side of resonance, which guides the decision.

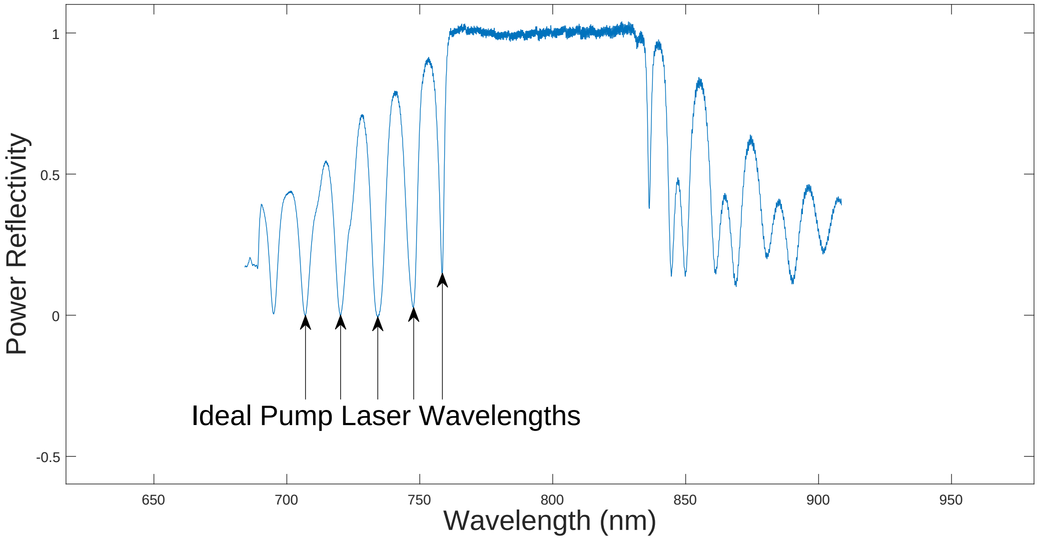

Figure 36: How to select a laser wavelength for non-resonant pumping to produce polaritons.

All the Bragg minimums at energies above the quantum well excitons are good candidates for non-resonant pumping. However, our Msquared tunable laser can not easily be tuned below 700 nm, so I have omitted the Brag minimum below 700 nm as an ideal pump wavelength. Pumping at higher energies (shorter wavelengths) is possible with other lasers in our lab, such as a 632 nm HeNE laser or a 532 nm green diode laser. All three lead to similar results.

The standard book value for the band gap of $Al_{0.2}Ga_{0.8}AS$ at cryogenic temperatures is right around 700 nm. This approximately lines up with one of my experiences in dealing with these samples. Generally the Bragg oscillations are not visible in the reflectivity structure at wavelengths shorter than 690 nm. This is likely because our DBRs are absorbing all the light rather than reflecting it. We can see this beginning at the short wavelength limit of the reflectivity data shown in Figure 36.

With the polaritons created, the photoluminescence (PL) may now be measured. An optical longpass filter is placed in the imaging path of our collection system to filter out the laser light. Typical PL is shown in Figure 37. In image A we see the full 2D image captured by the camera, and in image B the profile at normal incidence is shown. I've tried fitting the profile with both a Lorentzian and a Gaussian distribution and find that neither are completely correct, but both are mostly correct. I've tried converting the data into intensity versus energy instead, and still find that neither the Gaussian nor the Lorentzian distribution completely capture the behavior. Additionally, Voigt distributions which are a convolution of a Lorentzian and Gaussian distribution, did not quite work either. In fact, despite much investigation, I have never found a function with a small number for fitting parameters which completely fits the distributions. Likely this is because this is not a simple phenomenon. There is inhomogeneous line broadening due to disorder, which is readily visible in the exciton, which ultimately will contribute to the line shape of the polariton.

Still, I need to perform a fit at every angle in the image to extract our dispersion curve $E_{lp}(\theta)$ to proceed. Ultimately I settled on using the Lorentzian fit. The extracted curve is shown in Figure 38A as Position 1. Our samples are grown on a 2 inch GaAs wafer, and in this case the MBE machine produced a sample which is thickest at the center, and thinner at the edges. This effectively means the cavity mode is lower energy at the center and higher

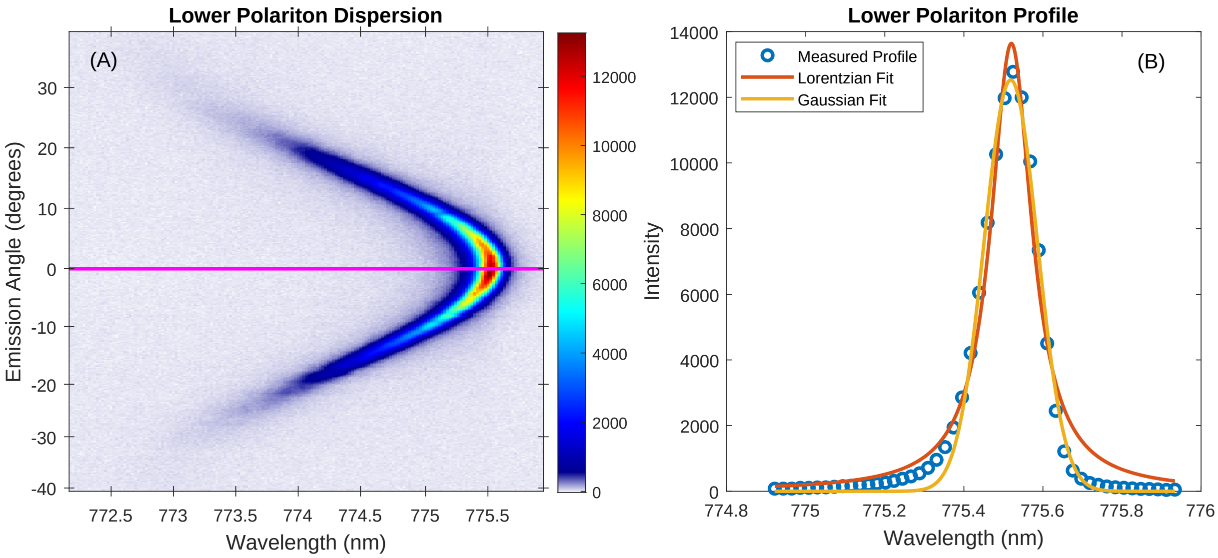

Figure 37: In image A we have typical lower polariton photoluminescence. In image B we have the profile along $\theta = 0$. We see Gaussian and Lorentzian fits are not quite right, but close.

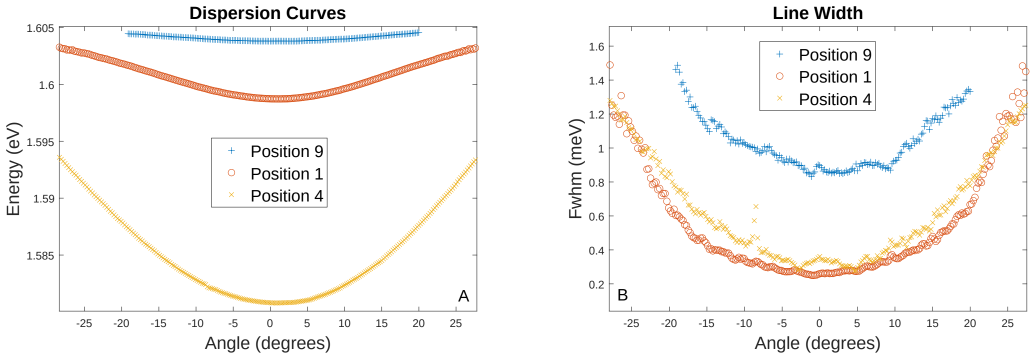

Figure 38: The dispersion curves and linewidths of the polaritons measured from three different locations on the sample.

energy at the edges. The excitons also vary slightly in the same manner due to an increased quantum confinement at the edges pushing their energy higher. However, in practice the energy shifting of the excitons as I move across the sample is quite minimal compared to the shifting of the cavity mode, within our region of interest. As such, movement across the sample modifies the detuning, and so a range of exciton fractions are observable.

In Figure 38 the point of observation has moved to a more photonic and excitonic location, where the PL of the polaritons is measured again. Additionally, I show the extracted line width's full width at half maximum (FWHM) in image B. Although our optics can collect data at approximately ±45 degrees, often PL from the polaritons is difficult to extract at higher angles. This is due to two separate effects. At very photonic detunings, the lower polariton is essentially a cavity mode, and one can see a very strong TE-TM splitting. This can of course be corrected by including polarizers in the setup, but it slows down the data collection, and ultimately is not necessary as we will see throughout this chapter. At very excitonic detunings, the lower polariton is essentially an exciton with a flat dispersion $E(\theta)$, the lower energy (longer wavelength) Bragg modes begin to crash into this relatively flat dispersion curve. This occurs because as I move across the sample, the entire reflectivity curve shifts with the thickness as well, including the Bragg modes. As a result, often the curve covering the largest emission angles occurs near a slightly photonic location, approximately 35% excitonic, although this is not a hard rule.

One interesting thing to note is that the linewidths measured at Position 1 and 4 are nearly identical, even though their energies are very notably different. I believe this says more about the resolution limit of our spectrometer than it does about the polaritons. Generally, the linewidth data from long-lifetime samples is hard to trust.

For my work, I wanted to collect the most pure measurements of the lower polariton possible. There were two main effects to avoid, which can be caused by our laser transmitting too much power. First, if the laser is absorbed and heats the sample significantly there will be a thermal red shift. As the sample is heated, the band gaps of the materials shift towards lower energies (longer wavelengths). Additionally, the reflectivity curve shifts towards lower energies (longer wavelengths) when the sample is warmed up. Therefore both the exciton and the cavity photon will shift towards lower energies when the sample is warmed up. Actually, these both shift for the same reason, ultimately, because the optical index is directly related to the band gap of the materials.

The other effect to avoid is a density-dependent blue shift. Excitons can efficiently scatter with one another; as a result some of this behavior is inherited in the polariton-polariton scattering strength. This scattering will show up as a shift of the energy of the polaritons, based on population density squared. In general, one can be safely outside of both of these regimes. I confirmed this by filtering the power of our pump laser down, and observing that the resulting dispersion curves experience almost no change. If the full power of the laser were used, it would certainly be possible to see one or both of these effects.

For a demonstration, see the data shown in Figure 39. The two images are from different samples and different days, but they illustrate the effect. In the data of Image A, I tried pumping the sample with a medium power laser, and then I introduced an optical filter to cut the laser power by one order of magnitude. In Image B I increased the temperature of the cryostat. The effect in Image A of a higher-powered laser is similar to the higher temperature curve in Image B; both result in red shifting. Our measurements are performed at low enough powers to minimize this shift.

Some of the results of a more thorough temperature study I conducted are shown in

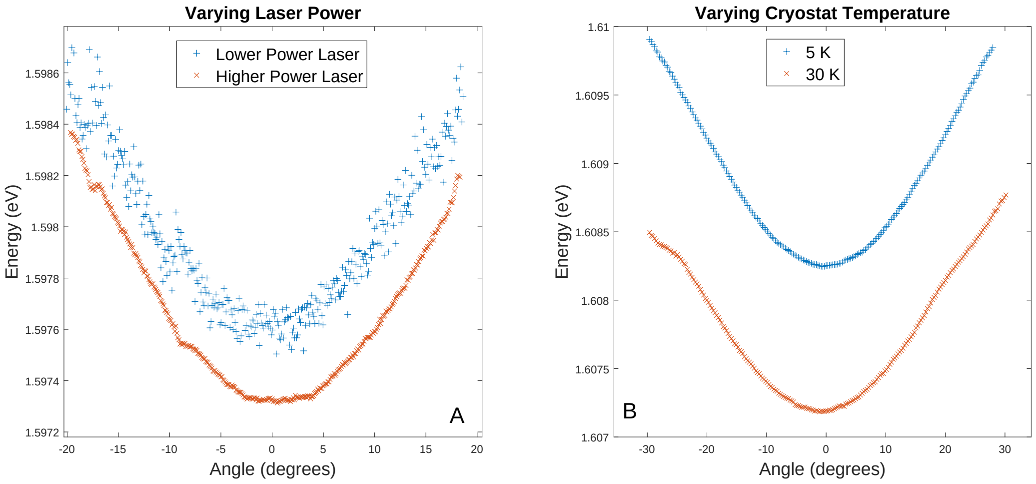

Figure 39: Red shift caused by heating the sample.

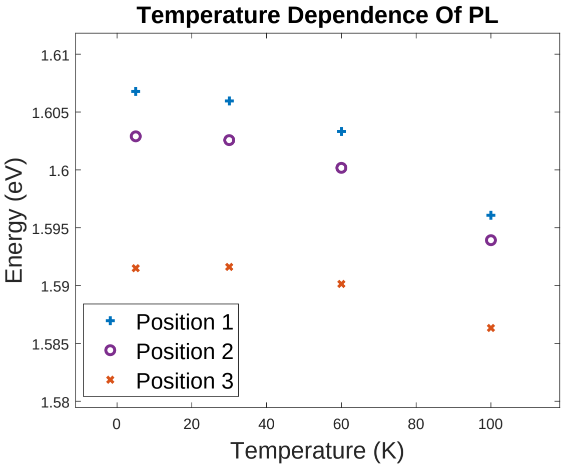

Figure 40: The effect of varying temperature on the lower polariton photoluminescence.

Figure 40. I collected the PL measurements at multiple temperatures and multiple locations. Only three locations are plotted for visual clarity. The energy shown is specifically at normal incidence, but full angle-dependent curves were collected. As expected, the lower polariton drops in energy as the temperature increases.

4.3 Photoluminescence Excitation Measurements

As discussed above, it is not possible to measure the upper polariton in our long-lifetime samples directly through PL measurements or reflectivity measurements. Therefore, I developed a method of photoluminescence excitation measurements (PLE). PLE measurements are not a new concept [87], but they were a new method for our lab. So far, PLE measurements are the only method to directly measure the upper polariton in our samples.

The basic idea of a PLE measurement is to sweep the wavelength of a laser across some system resonance while observing the photoluminescence of a lower-energy resonance. For my work, that means sweeping the laser across the upper polariton while measuring the lower polariton's PL. The laser is at higher energy than the lower polaritons, so downscattering will result in a lower polariton population, and corresponding lower polariton PL. When the laser is resonantly tuned to the upper polariton, a local maximum in energy absorption should occur, resulting in an observable increase in the lower polariton PL Intensity. Ultimately, we are creating upper polaritons which scatter down through phonon emission until they populate lower polariton states.

Particular care must be taken concerning background subtraction, but I have devised a scheme which I've found works very well and reliably. PLE measurements may be performed conveniently using algorithms in our software outlined in Appendix A. This also includes the data processing.

The optics setup in Figure 16 is converted to a reflectivity geometry by flipping Mirror 1 into the laser path. The sample is mounted inside our cryostat, and positioned at the sample plane of our microscope objective. A diagram showing the optics setup is shown in Figure 41.

The first step is to obtain the most pristine image of the lower polariton dispersion curve possible. This is best achieved through non-resonant pumping, discussed in Section 4.2. In short-lifetime samples this non-resonant pumping will result in both lower polariton and upper polariton PL being visible. Unfortunately, as discussed above, in long-lifetime samples, the main decay channel for the upper polaritons is scattering down into lower polariton states, rather than annihilation. As a result, we see virtually no upper polariton

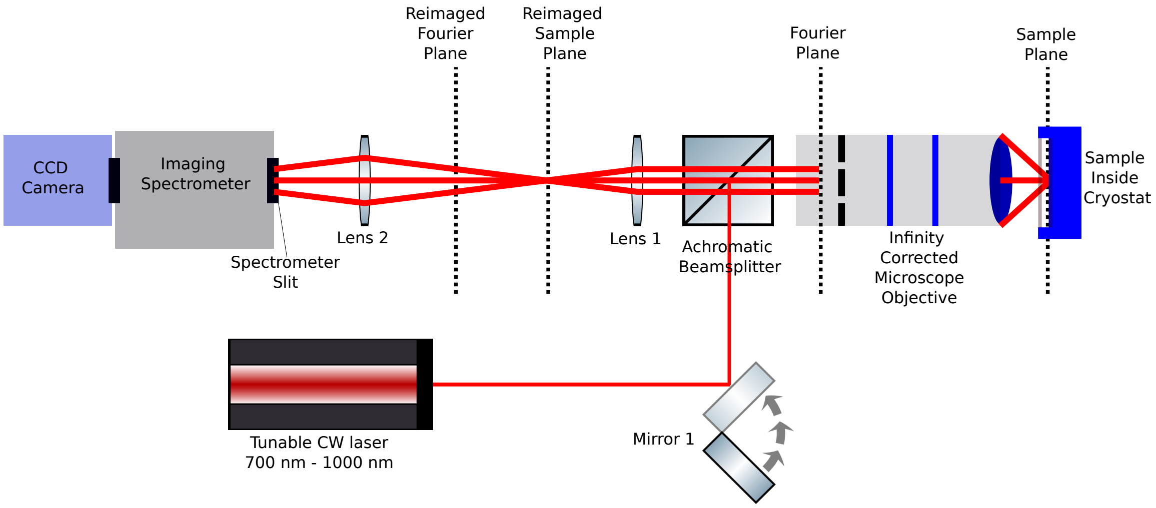

Figure 41: A simple illustration of the optics setup for PLE measurements. Not pictured are polarizers and wavelength filters in the imaging path, near two reimaged planes.

PL in long-lifetime samples.

Typical lower polariton PL is shown in Figure 37. This image will be an essential part of our background processing. The biggest challenge I faced in the PLE measurements was that because the upper polariton's wavelength is only about 10 nm away from the lower polariton, the noise from the laser light scattering can easily overwhelm the signal of the measurement. To avoid this, I first used basic optical methods to filter out scattered laser light. The laser is vertically polarized, and the lower polaritons produced through the phonon emission are unpolarized. Therefore, I put a horizontal polarizer into the imaging setup to remove a large amount of the scattered laser light while only removing half of our PL signal. Furthermore, I had a tunable longpass filter from Semrock, which I put into the imaging setup. This filter is tuned by simply rotating the angle of the filter slightly. I tuned it so that most of the lower polariton PL may pass, while most of the laser is filtered out. Then, without changing any optics, I swept the wavelength of the laser while taking a series of images. One such image is shown in Figure 42A.

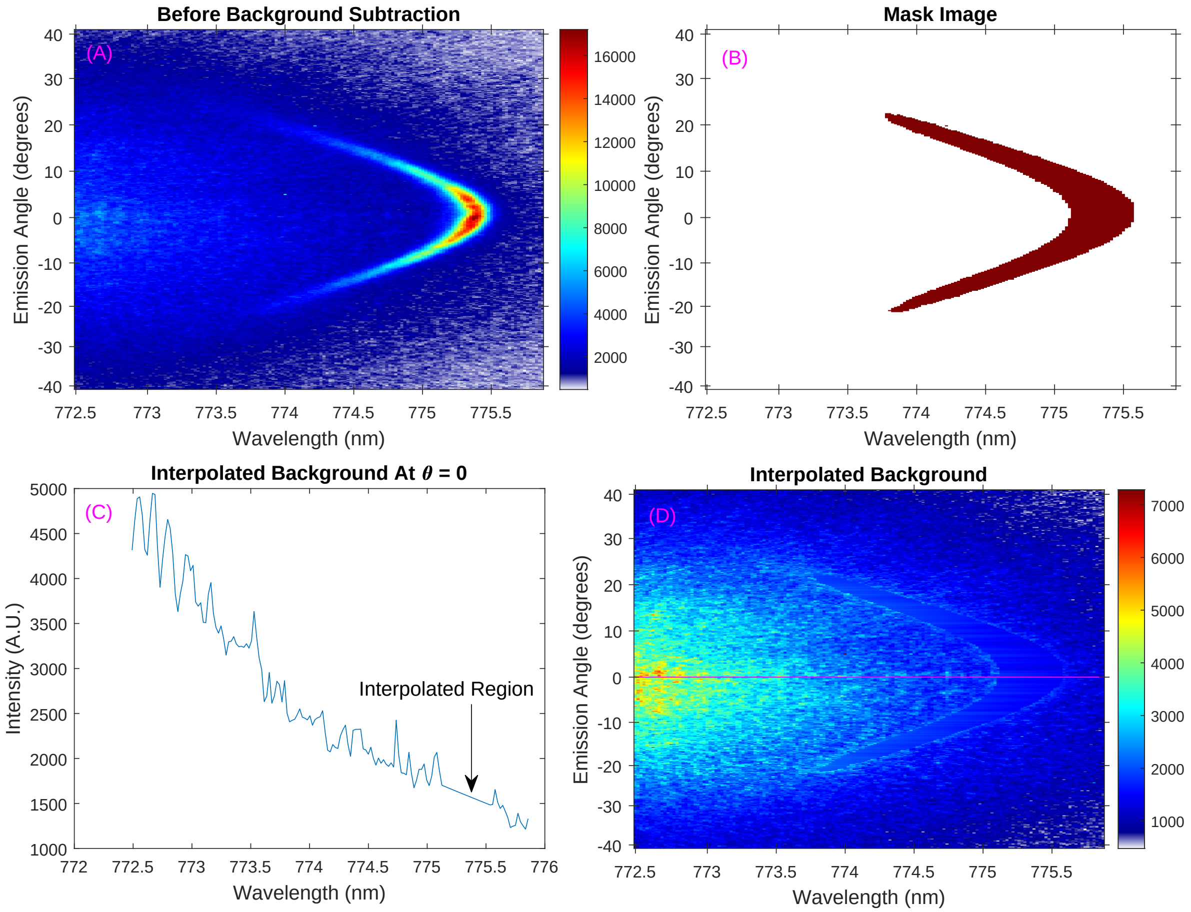

Figure 42: These images show the background subtraction scheme for PLE measurements.

Comparing Image 42A to 37, we see that even with our optical filtering, a large amount of laser light backgrounf is still present in the image. Worse yet, this laser light travels horizontally across the image as I sweep the laser. The majority of this comes from laser light reflected off the sample, which means it also changes with position on the sample. Therefore, there is no simple way to just take a background image, and subtract off the noise. Instead, I devised a method to interpolate all the pixels which do not contain PL to give the value of the background in the pixels which do contain PL.

I started by using the clean image of the PL created through non-resonant pumping shown in Figure 37A to determine which pixels contain PL signal. That information will then carry over to Figure 42A, telling me which pixels contain both PL signal and background versus which pixels contain only background. I will call the intensity of the clean image in Figure 37A $P_c(\lambda, \theta)$ and the intensity of the noisy image in Figure 42A $P_n(\lambda, \theta)$.

I defined a threshold intensity T, such that pixels in Figure 37A with intensities above this threshold are considered to have polariton PL in them, while pixels with intensities below this threshold are assumed to contain only background and no PL. The value of the threshold is determined by experience using my data processing interface [12], which allows the user to control this choice by producing the mask image shown in 42B.

The threshold T has a region of usable values. If the threshold is too high, we have no pixels with PL in them. If the threshold is too low, we have no pixels without PL, which means we have no pixels to interpolate to determine the background inside the pixels with PL. Generally, this means it is better to use a T value which is lower rather than higher. Therefore, I erred towards lower threshold values, which is why the red pixels in the mask image seem larger than the PL in the clean image. Mathematically this can be expressed as

$$M(\lambda, \theta) = \begin{cases} 1 & P_c(\lambda, \theta) \geq T \\ 0 & P_c(\lambda, \theta) < T. \end{cases} \tag{4.1}$$

Pixels above the threshold are shown as the red pixels and pixels below the threshold as white pixels. Along each row of the image in Figure 42A, I fit the white pixels with some curve, then use that curve to interpolate the value of the background within the red pixels. A third order polynomial is sufficient to capture the intensity behavior of the white pixels. I denote this best fit as $\text{Polyfit}(M = 0, \lambda, \theta)$, where $M = 0$ indicates that the fit is created using only the white pixels. A separate fit is performed for each row of the data. The constructed background is this written as

$$B(\lambda, \theta) = \begin{cases} P_n(\lambda, \theta) & M(\lambda, \theta) = 0 \\ \text{Polyfit}(M = 0, \lambda, \theta) & M(\lambda, \theta) = 1. \end{cases} \tag{4.2}$$

The constructed background corresponding to the image in Figure 42A is shown in figure 42D, and the intensity profile along the line drawn at $\theta = 0$ is shown in Figure 42C. The interpolated region is labeled and quite visible, as the random fluctuations seen in the real noise are not present in the interpolated region. However, the integrals that follow in the next steps reduce the impact of fluctuations. My interpolation scheme has produced a background image which is sufficient for my purposes. I can now subtract off the background image to produce an image of just the lower polariton PL. This background subtracted image's intensity is $P(\lambda, \theta)$, giving

$$P(\lambda, \theta) = P_n(\lambda, \theta) - B(\lambda, \theta). \tag{4.3}$$

This procedure is repeated at each wavelength in the laser sweep. Examples of these background subtracted images are shown in Figures 43A, 43B, and 43C.

Now that the bulk of the laser noise has been removed from the PL images, the images can be integrated to obtain the total intensity of lower polaritons as a function of the pump laser's wavelength. The images are actually two-dimensional slices of the full three-dimensional k-space. Because the dispersion relationship is radially symmetric in the in-plane momentum directions, I performed this integral using circular symmetry. Numerically, this integrated lower polariton intensity is

$$I = \sum_{\lambda} \sum_{\theta} P(\lambda, \theta)\ 2\pi\theta\ \Delta\theta\ \Delta\lambda, \tag{4.4}$$

where $\Delta\theta$ and $\Delta\lambda$ are determined by the resolution of our setup.

In order to convert from counts on a camera to absolute polariton density, it is necessary to correct by the lifetime of the polaritons. This lifetime is both angle and wavelength dependent. In principle, if the temperature of the lower polaritons changed during the measurement, this could lead to a change of the calibration of the lower-polariton density. However, the process of sweeping the pump laser wavelength does not significantly change the distribution of the polaritons. This means that the distribution of $P(\lambda, \theta)$ does not change from image to image. Since the goal is not to determine absolute polariton populations, but rather only to see relative differences in this integrated intensity, this correction may be ignored for the measurement.

The integrated intensity as a function of the laser wavelength is shown in Figure 43D. A clear peak is present in the center of the data range, which is well fit by a Lorentzian distribution. Care was taken to increase the density of our data points in the most important

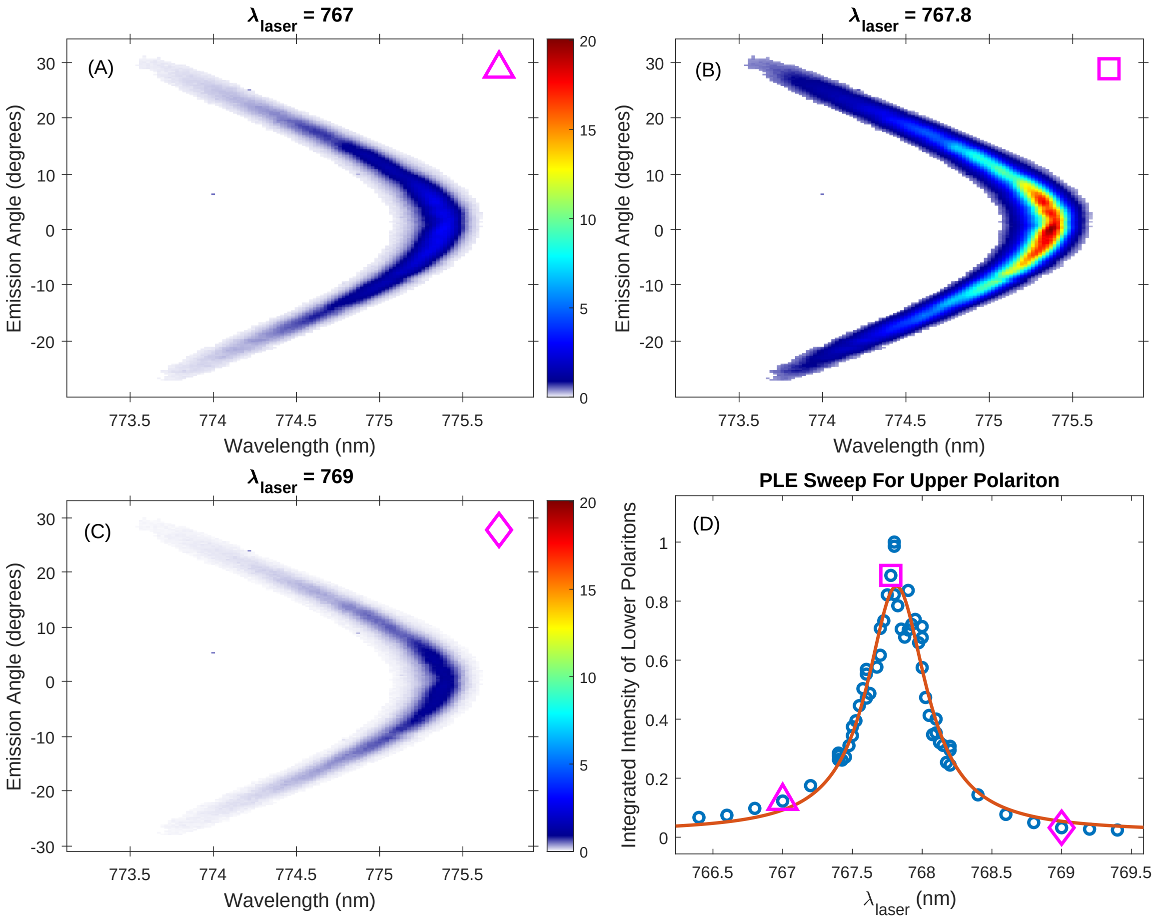

Figure 43: Images of the lower polariton as the pump laser wavelength is swept across the upper polariton's resonance. The background subtraction is handled as demonstrated in Figure 42. Images A, B, and C are then integrated to produce the data points shown in image D. We see a clear peak form corresponding to the upper polariton's resonance.

spectral region. Three data points are selected and enclosed by a triangle, square, and a diamond. The background subtracted images corresponding to these three data points are shown in Figures 43A, 43B, and 43C. One can easily see that Figure 43B is noticeably brighter than the other two images. From the Lorentzian fit I extract the line center, and conclude this is the wavelength of the upper polariton $\lambda_{up,ple}$. Furthermore, because the pump laser was a cone centered on normal incidence ranging over ± 5 degrees, I conclude specifically that I have measured the upper polariton at $k_{\parallel} = 0$.

In Chapter 5 I will report these same measurements on short-lifetime samples. In such samples, the upper polariton is visible in reflectivity measurements, and so one can compare the PLE measurement to a direct reflectivity measurement of the upper polariton. There is a small discrepancy; that discrepancy may be factored into an overall error bar. The maximum discrepancy observed in the short-lifetime samples was 0.4 meV, in which the PLE result was at higher energy than the reflectivity result.

I also investigated the effects of temperature on PLE. I found that PLE measurements still worked very well at temperatures of 100 K. I did not investigate at higher temperatures as polaritons typically don't survive at higher temperatures, although some recent results in our lab [4] indicate that polaritons can persist in the weak coupling limit at room temperature. It would be interesting for someone to try replicating my methods on polaritons at room temperature. Some of this temperature-dependent data is shown in Figure 44. The upper polariton shifts towards lower energy as the temperature increases, as expected.

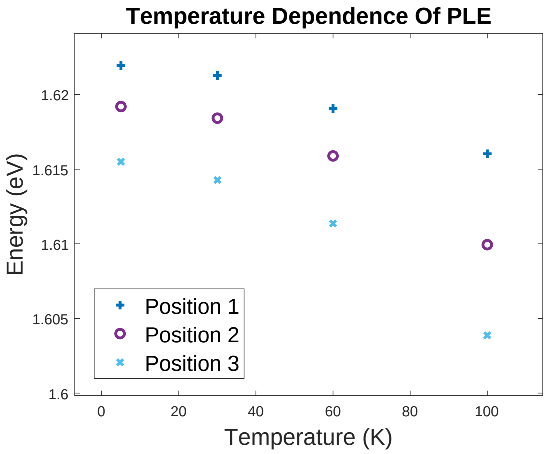

Figure 44: The effect of varying temperature of the sample on the result of a PLE measurement.

One alternative to the PLE measurement which I attempted involved the same sweeping of the laser, but instead of measuring the lower polariton, I measured the power of the laser coming off the sample. Theoretically, if a significant portion of the laser were absorbed by the sample when resonant with the upper polariton, a noticeable dip in the power of the reflected beam would be absorbed. This is quite appealing because it would mean measuring a very bright signal, rather than a weak one surrounded by laser noise. Unfortunately, I found this was not a successful endeavor. The measured reflected laser power did not show a noticeable dip, most likely because total power absorbed is too small.

4.3.1 PLE Measurements Of The Lower Polariton

In this subsection, I present my attempts to measure the lower polariton via PLE. These results are not essential to the overall characterization of long-lifetime samples. As discussed earlier, the lower polariton is not measurable in reflectivity spectra in our high-Q samples. It would still be good to know what such a measured value would be, to know how much shift there is between PL and reflectivity measurements of the lower polariton.

This would allow direct observation of the Stokes shift, which is a shift caused by disorder, typically causing the PL measurement to be observed at lower energy than the reflectivity measurement [33]. As discussed in Chapter 5, I actually observed the opposite in the short-lifetime samples: I observed PL measurements at higher energy than the reflectivity measurements, which I ultimately attribute to some sort of thermal shifting.

The PLE measurement is very similar to a reflectivity measurement, and so I wondered if it was possible to measure the lower polariton via some modifications to my PLE measurements. Sweeping the laser over the lower polariton is easy enough; the problem is that there are not any nearby states for the lower polaritons to decay into. Indeed, one of the appealing aspects of the lower polaritons is that their primary decay channel is emitting photons, which I collected for my various measurements.

One thought was to look for PL emitted from one of the Bragg modes. Looking at the reflectivity curve in Figure 36, there are clearly Bragg modes at energies below the polaritons. Sometimes when I non-resonantly pumped a sample, some PL is observed at these Bragg modes. However, when I tried sweeping the laser over the lower polariton, I did not find the PL from these Bragg modes to noticeably change.

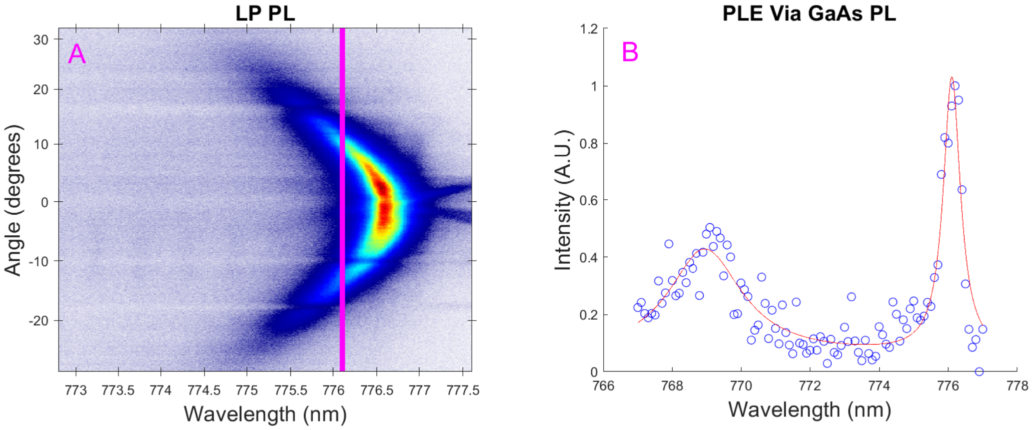

Figure 45: PLE measurement of the upper and lower polariton by monitoring GaAs PL Intensity.

Another thought was to try looking for the bulk GaAs PL, which predominantly originates from the substrate. Of course, the substrate is buried deep below two very large DBRs. GaAs PL is regularly observed when I pump the sample non-resonantly. Typically the substrate PL is around 815 nm (1.5213 eV). For the most part this did not work. However, on Sample P9-10-20.1 from the Princeton group, it did. This sample is an outlier, however. It arrived in the mail shattered. There were still large pieces of the sample, which I was able to mount in the cryostat for experimentation after cleaning them with methanol. Some of the data from this experiment is shown in Figure 45.

The thick pink line in Image A is the result of the PLE fitting shown in Image B. The GaAs PL did increase in brightness when the laser was resonantly pumping either of the two polaritons. This may be an accurate measurement of the upper and lower polariton, in which case it should give a value close to the result we would see in a reflectivity measurement. In this case, rather than a thermal shift, the PL is at lower energy than the PLE. This would be more in line with a standard Stokes shift caused by disorder. However, the fact that this measurement only seems to work in one specific sample is suspicious. The fact that the results are believable, however, suggests that there is a technique to discover here.

Another idea I had to measure the lower polariton via PLE was to try resonantly pumping the sample to create a condensate. In principle, when the laser was resonantly tuned to the lower polariton, a condensate would form, and at non-resonant settings it would not. Of course, the scattered light from the laser can easily overwhelm the camera, and the condensate would not be visible. One solution was to pump the sample with a ring, then use real-space filtering to cut out the laser, and collect only data from the center of the ring. One such ring pump profile is shown in Figure 46A.

The optics leading into the axicon reshaped the beam so that the final laser shown in 46A is almost all at normal incidence. This is of course necessary for a PLE experiment, as the angle of the pump laser gives the polariton's dispersion curve. Ultimately this experiment did not work. However, while I had this setup built; I tried once again measuring the lower polariton via PLE by observing the GaAs PL. I found that this worked pretty well, but only at very excitonic detunings. Figure 46B shows the angle-resolved PL of the lower polariton, with the PLE result for normal incidence plotted on top as the thick pink line. There is a small but notable Stokes shift. However, as evidenced by the very flat polariton dispersion, these polaritons are essentially excitons. Ultimately this method is not versatile enough to help in long-lifetime sample characterization.

4.4 Failings Of The Rabi Model

Given a successful measurement of the upper polariton via PLE, is this alone enough to calculate an exciton fraction for the polaritons using the Rabi model outlined in Section 1.5? As I will show here, it is not, at least not accurately. To begin, I will rearrange the polariton equations to give something more easily fit to the data. The energy of the lower polaritons is

$$E_{lp}(k_\parallel) = \frac{1}{2}\left(E_{ex}(k_\parallel) + E_c(k_\parallel) - \sqrt{4g^2 + \left(E_{ex}(k_\parallel) - E_c(k_\parallel)\right)^2}\right). \tag{4.5}$$

As mentioned earlier, the mass of the excitons is about five orders of magnitude above the mass of the cavity photons. One can therefore approximate the exciton dispersion curve as not having any $k_\parallel$ dependence. The excitons only couple to the photons that have the

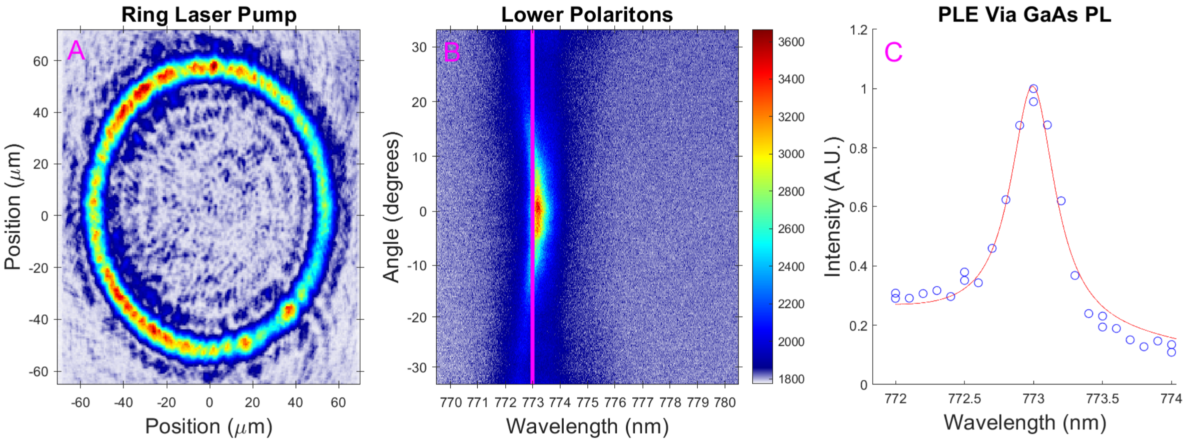

Figure 46: A: Pump laser spatial profile for the condensate measurement. B: The lower polariton angle-resolved PL; the thick pink line is the $k_{\|} = 0$ result of the PLE measurement of the lower polaritons using the GaAs substrate luminescence. C: The PLE measurement of the lower polariton, along with a Lorentzian fit of the data.

same $k_\parallel$ value. However, because the excitons do not have $k_\parallel$ dependence, I can now simply replace the $k_\parallel$ dependence with θ dependence. This is advisable because the raw data is in angle, not momentum. This leaves the dispersion of the upper and lower polaritons as

$$E_{up,lp}(\theta) = \frac{1}{2}\bigg(E_{ex} + E_c(\theta) \pm \sqrt{4g^2 + \big(E_{ex} - E_c(\theta)\big)^2}\bigg), \tag{4.6}$$

where the cavity photon dispersion is

$$E_c(\theta) = \frac{E_{co}}{\sqrt{1 - \left(\frac{\sin\theta}{n_c}\right)^2}}. \tag{4.7}$$

Here, $E_c(0)$ is written as E_{co}.

Equation (4.6) can be viewed as a fitting equation for our measured lower-polariton data, with four fitting parameters, E_{ex},E_{co},n_c, and g. In order to show the shortcomings of the Rabi model, it is most convenient to recast these equations so that the exciton fraction at normal incidence $|X_o|^2$ and our PLE result $E_{up,o}$ are fitting parameters. This can be done by writing E_{co} and g in terms of E_{ex} , $E_{up,o}$, and $|X_o|^2$. Deriving the necessary relationships is most conveniently done by returning to the original diagonalization problem which lead to the polaritons as the new eigenstates. The equations for the eigenvectors of the $k_\parallel = 0$ state in this case are

$$\begin{pmatrix} E_{ex} & g \\ g & E_{co} \end{pmatrix} \begin{bmatrix} X_o \\ C_o \end{bmatrix} = E_{lp,o} \begin{bmatrix} X_o \\ C_o \end{bmatrix} \tag{4.8}$$

for the lower polariton state and

$$\begin{pmatrix} E_{ex} & g \\ g & E_{co} \end{pmatrix} \begin{bmatrix} C_o \\ -X_o \end{bmatrix} = E_{up,o} \begin{bmatrix} C_o \\ -X_o \end{bmatrix} \tag{4.9}$$

for the upper polariton state.

The energies are solved for by setting the determinant of the system equal to zero; so each matrix only has one independent equation. One can first solve for g in terms of the new parameters. The top row of the upper polariton's matrix equation is

$$E_{ex}C_o - gX_o = E_{up,o}C_o. \tag{4.10}$$

This can be manipulated into

$$-gX_o = (E_{up,o} - E_{ex})C_o. \tag{4.11}$$

To utilize $|C_o|^2 = 1 - |X_o|^2$, one can multiply each side by its own complex conjugate. Here, I assume g is purely real. That leaves us with

$$g^2|X_o|^2 = |E_{up,o} - E_{ex}|^2|C_o|^2 = |E_{up,o} - E_{ex}|^2(1 - |X_o|^2). \tag{4.12}$$

Finally, g may be isolated as

$$g = |E_{up,o} - E_{ex}|\sqrt{\frac{1 - |X_o|^2}{|X_o|^2}}. \tag{4.13}$$

This is the first needed transformation equation to move to the new set of parameters.

Next, the equation for E_{co} can be derived. This is most easily done by returning to the part of the diagonalization where the determinant is set equal to zero. Using Equation (4.9) and setting the determinant equal to zero yields

$$(E_{ex} - E_{up,o})(E_{co} - E_{up,o}) = g^2. \tag{4.14}$$

This equation may be quickly solved for E_{co}, leading to

$$E_{co} = E_{up,o} + \frac{g^2}{E_{ex} - E_{up,o}}. \tag{4.15}$$

This is the second transformation equation needed to shift to the new fitting parameters.

Now equations (4.7), (4.13), and (4.15), can be substituted into equation (4.6) to yield the final functional form to fit to the data, which can be expressed as the function $E_{lp}(E_{ex}, |X_o|^2, E_{up,o}, n_c, \theta)$. However, the result is a very messy equation which is not particularly illuminating. Since in my MATLAB code I just use a multi step calculation instead, it is more helpful to show the MATLAB code which calculates the dispersion in multiple steps. We can define

$$E_{LP}(\theta) = E_{lp}(E_{ex}, |X_o|^2, E_{up,o}, n_c, \theta). \tag{4.16}$$

where the function is computed by the following code:

```
function E_lp_output = E_LP(E_ex,ex_frac,E_up_o,n_cav,theta)
%E_ex is exciton energy (infinite mass approximation)
%ex_frac = |X_o|^2
%E_up_o is our PLE result, the upper polartion at \theta = 0
%n_cav is the optical index of the cavity
%theta is the indepdent variable (x-axis), in degrees
%
%E_lp_out is the lower polartion dispersion curve

%calculate g
g = abs(E_up_o - E_ex) * sqrt( (1-ex_frac)/ex_frac );

%get cavity zero energy
E_c_o = E_up_o - (g^2) / (E_up_o - E_ex);

%get cavity at all angles
E_c = E_c_o ./ sqrt( 1 - ( sind(theta)/n_cav ).^2 );

%make lower polaritons
E_lp_output = .5*(E_ex + E_c - sqrt(4*g*g + (E_ex - E_c).^2));

end
```

The true fitting function also includes a possible shift to the angular axes. This is because in day-to-day alignments the $\theta = 0$ line isn't always perfectly centered on the screen. The angular calibration of the camera still applies; it is just shifted a few pixels up or down.

In Section 4.4.1 I will fit this dispersion curve to the experimental data. The advantage of the variable transform I performed is that now the exciton fraction and the PLE result are fitting parameters. This means I can do fits with and without the PLE parameter freely varying. I will use this to show that PLE is a necessary constraint. Additionally, the exciton fraction being a free fitting parameter means that I can pin down the exciton fraction value in the fits, and fit the other free parameters to get the best constrained fit. This will allow me to demonstrate how tightly constrained the exciton fraction is, which is of critical importance.

4.4.1 Fitting The Dispersion Curve With And Without The PLE Constraint

The dispersion curve was extracted from Figure 37A, and previously plotted in Figure 38A. I can fit equation (4.16) to it without utilizing the PLE result we measured. All the fitting parameters are free to vary unconstrained. Unsurprisingly, a very nice fit of the data can be obtaoined, shown in Figure 47A.

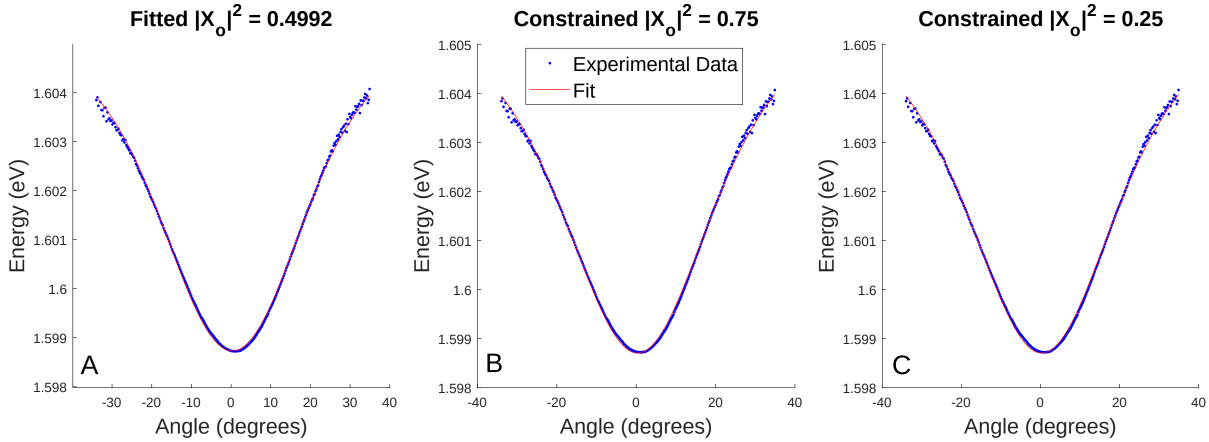

Figure 47: Three different fits of the same data with very exciton fractions.

However, if the exciton fraction is to 0.75 or to 0.25, and the other parameters are allowed to vary freely, the fits shown in Figures 47B and 47C are obtained. All three fits are extremely good. This is consistently the case when fitting lower polariton dispersion data without a PLE constraint. To more thoroughly investigate this, I systematically swept the exciton fraction while allowing the other parameters to vary unconstrained. In order to quantify the quality of the fits, I define a generic error for fitting curves to one dimensional data sets, given as

$$\text{Error} = \sum_x (\text{Fit}(x) - \text{Data}(x))^2. \tag{4.17}$$

The results are shown in Figure 48; the error is normalized by the minimal error from the sweep. The two fits from the most extreme locations in the sweep are shown, corresponding to exciton fractions of 0.05 and 0.95. The extremely-low-exciton-fraction fits still look very good. The extremely-high-value fits are notably worse fits even to the human eye. By eye, the fits ranging from 0.05 to .85 are all very good fits.

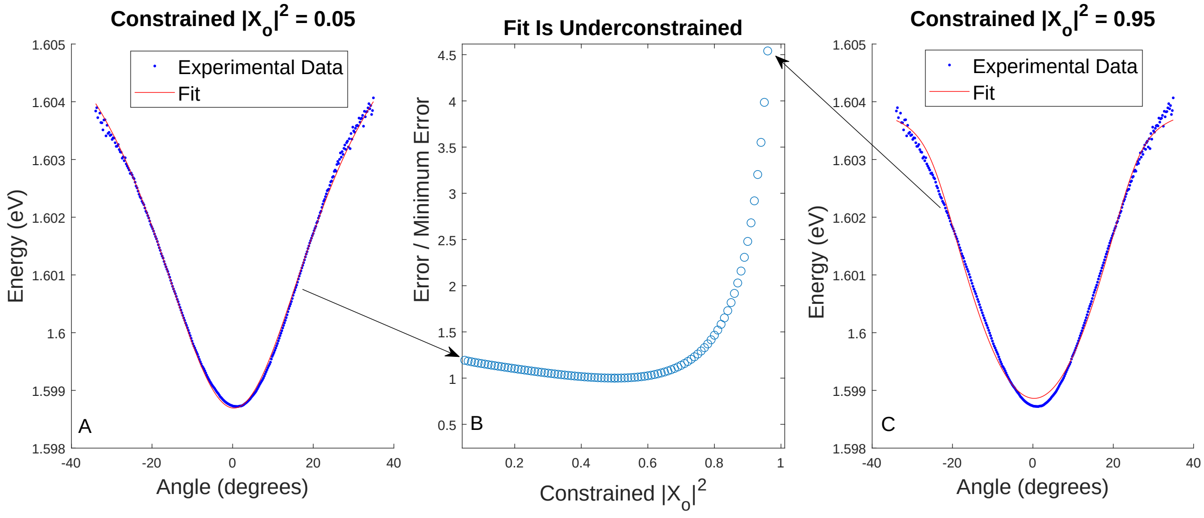

Figure 48: The results of systematically sweeping the exciton and obtaining the best fit with that value. Here we do not have any PLE data constraining our fit.

These results are compelling evidence that more measurements are required to pin down the exciton fraction accurately. Therefore I utilized the PLE result so that the $E_{up,o}$ parameter is not a free fitting parameter. The results of this sweep are shown in Figure 49. Image D shows the error as the exciton fraction is varied, keeping the $E_{up,o}$ parameter fixed to the value measured in the PLE experiment. Image C is the result of allowing the exciton fraction to be a fitting parameter, rather than selecting its value and fitting all the other parameters around it. That is to say, Image C has the globally best fit. Comparing the scale of the vertical axis in Figure 49D to the scale in Figure 48B, it is readily apparent that the fits are much more constrained now. Three of the data points have pink geometric shapes placed around them to indicate that the full fits are shown in Images A B and C. Although

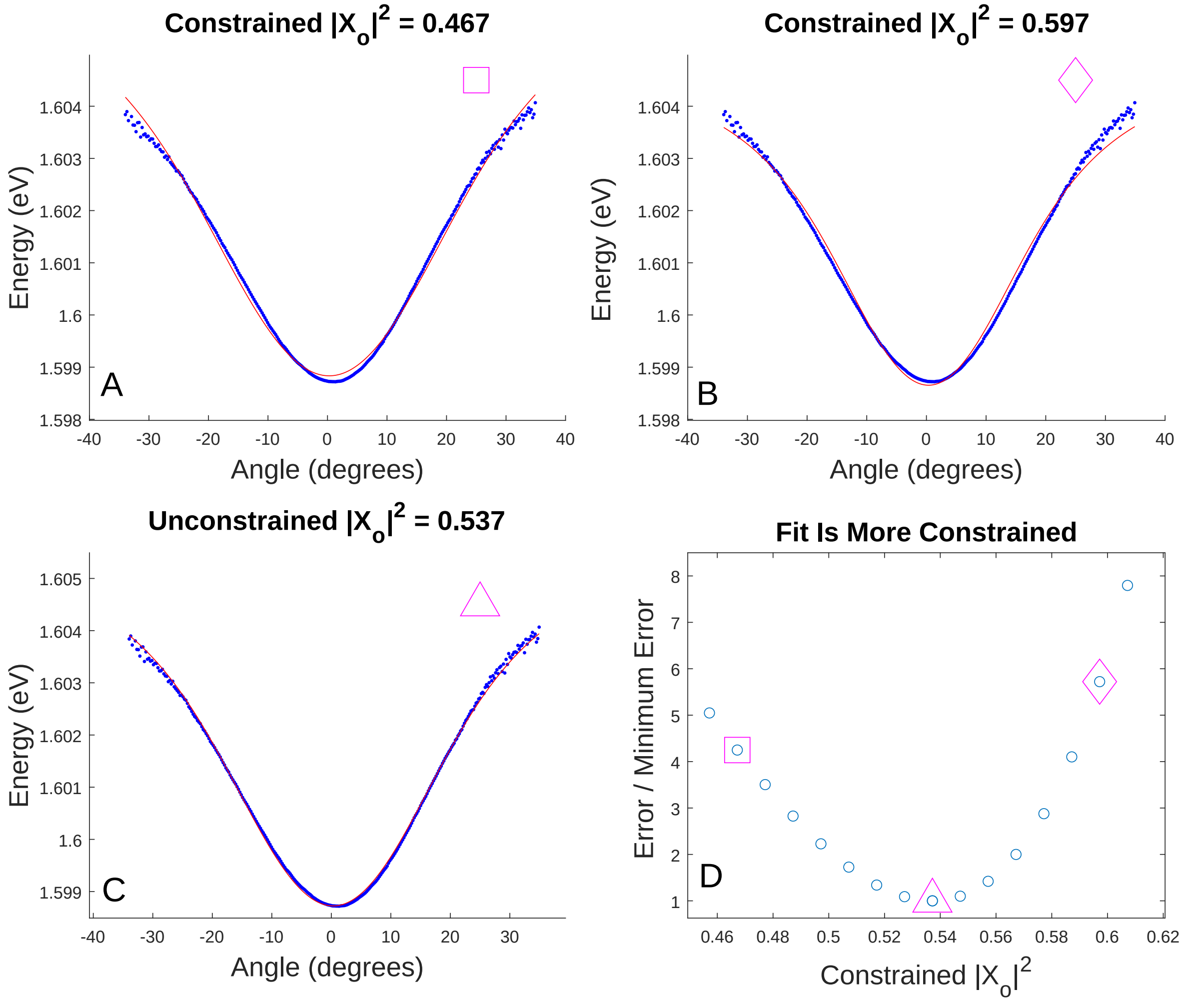

Figure 49: The results of systematically sweeping the exciton and obtaining the best fit with that value. Here we utilize our PLE result as a constraint.

the difference is faint, Images A and B represent what the limit that the exciton fraction could be changed while just beginning to become visibly unacceptable fits. Therefore the exciton fraction is constrained between 0.46 and 0.6 in this case.

Although this seems like a good fit, this section is titled “Failings of the Rabi Model” because of the problems discussed below. Just because a model is well constrained does not mean it is accurate.

The first issue to address is that of linewidth. The equations presented in Section 1.5 come from purely real-valued energies. This is equivalent to assuming quantum states with infinite lifetimes or zero linewidth. Real exciton and cavity states have finite lifetimes. In the absence of inhomogeneous broadening due to disorder, this may be included in the Rabi model by simply adding an imaginary piece to the energies. The equations do not change so much that they can not be used.

However, the number of parameters has gone up. There is now a linewidth associated with the exciton and the cavity. That are two more fitting parameters. Unsurprisingly, if I fit only the dispersion curve with the now complex Rabi-model equations, the fit is very loosely constrained, just as it was when I didn't use the PLE data point as a constraint. Naturally, the measured linewidths are to be considered. However, as discussed in regard to Figure 38B, it seems likely that the line broadening caused by the spectrometer is playing an important role in our measurements. I could try to account for that instrumental resolution, but it becomes yet another parameter in the model. The PLE measurement can possibly serve as a measurement of the linewidth of the upper polariton, and I've tried incorporating it into a complex Rabi model. But ultimately the fit was still poorly constrained, leading to exciton fractions between 0.35 and 0.58. Fundamentally this is a problem of information. In moving from purely real energies to complex energies I introduced more additional fitting parameters, and simply do not have enough reliable measurements to tightly constrain these fits.

Additionally, although complex energies are sufficient for handling homogeneous line broadening,there is also broadening due to inhomogeneous disorder. Our sample consists of 12 quantum wells all at slightly different locations inside of the microcavity. Differences between these quantum wells would need to be incorporated into the Hamiltonian as additional exciton states. This would be the simplest version of incorporating inhomogeneous disorder. There is almost certainly some number of localized exciton states caused by small defects in the growth of the sample as well.

Even worse, even if I create a model that takes these issues into account, a question of usefulness arises. I am not only interested in analyzing samples that already exist, but also modifying existing sample designs to produce new samples. For example, at one point I wished to iterate on our designs to produce samples that are close to 50% excitonic at the flat region in the middle of the sample so as to create polariton structures with minimal energy gradient [88]. Simply using the equations from a two-state model may be sufficient for determining an exciton fraction; however, it tells nothing about how to modify a design to achieve a desired sample. For these reasons, I pursued an electromagnetic simulation of the sample utilizing the transfer-matrix method (TMM), discussed in the next section. This model will ultimately give exciton fractions from its fits, and also provide insight into future sample design.

4.5 Reflectivity And Transfer Matrix Method Simulations

A thorough derivation of the equations for a transfer-matrix method (TMM) simulation is given in Chapter 3 based partially on Ref 19. In Section 2.5 I discussed how to perform reflectivity measurements and showed some example reflectivity data in Figure 24. More reflectivity data is shown below in Figure 50. In this section, I focus on applying the TMM simulation and tailoring it to the problem of characterization. My goal was to produce a working TMM simulation of our samples which can be used to determine exciton fraction of the polaritons. I also wished for the simulation to be accurate even with slight modifications of the sample design, so as to predict the polariton characteristics of proposed future samples.

The central aspect of a TMM simulation is that the sample is treated as a series of infinite slabs of material stacked upon one another. The Fresnel equations are utilized along with some linear algebra, and ultimately with knowledge of the complex refractive index of the materials in the sample, the reflection and transmission coefficients of an electromagnetic plane wave incident on the sample can be calculated. The TMM equations can handle various angles, wavelengths, and polarizations. Thus, a working TMM simulation allows broad reflectivity measurements to be useful constraints on the model of our sample.

The process here is broken down into two steps. First, I discuss how to fit the reflectivity simulation to the broad reflectivity measurements performed in Section 2.5. Then, I discuss fine tuning the quantum well and cavity parameters so that the simulated polariton branches match the PL measurements of the lower polariton and the PLE measurement of the upper polariton.

4.5.1 Fitting TMM Model To The Broad Reflectivity

To create a TMM simulation that matches the data shown in Figure 24, I needed accurate data for the index of refraction of the main materials in the DBRs; GaAs, AlAs, and the alloy AlGaAs (20% aluminum) at cryogenic temperatures. I tried various published data sets and models for the index functions [62, 6, 32, 67, 35, 63, 37, 66]. Ultimately, I decided to use the index functions given by Reference 37 because they allowed me to tune the temperature and alloy percentage. This is appealing because it is often quite useful to first measure the properties of a structure at room temperature before going to cryogenic temperature. The index functions given in Reference 37 covers all three of the DBR materials over a wide range of temperatures and wavelengths. This particular library, however, does have the downside that the imaginary component of the index functions seem to be consistently larger than what I believe is the case for our samples, based on the fits of my TMM simulation to the data. Most likely, the discrepancy comes from the difference between high purity MBE grown materials for our structures compared to materials grown through other methods. Higher purity materials will have less impurities, which should reduce absorption, especially in the transparent region below the band gap where absorption is already low, which is the most important spectral region for my model. I found that to have consistency with our data, the imaginary part of the index of refraction of the DBR materials needs to be reduced by about a factor of 100 compared to the values reported in Ref. 37, in order to produce the characteristic flat stopbands shown in Figure 24. This results in the imaginary component of the modified index functions of Reference 37 being in reasonable agreement with Reference 63, which reported index measurements on MBE-grown materials specifically.

No matter which index reference I use, I will need to tune parameters of the simulation to match the data. I find that simply uniformly scaling the thicknesses of all the layers of the sample in the simulation relative to the designed, nominal values of the widths is not sufficient to bring the simulation into good agreement with the data. I therefore tuned various parameters of the simulation to fit the broadband reflectivity data. This TMM simulation was coded into a convenient graphical user interface (GUI) based on the MATLAB platform, which is available online for download [12], along with video tutorials of its use [11]. This GUI also includes an optimizer to tune variables in order to align simulation with measurement. Below, I outline what the variables are which I tuned in order to produce the good agreement shown in Figures 50B, 50C, and 50D.

The first set of variables I introduced are three scaling factors for the thickness of all GaAs, AlAs, and AlGaAs layers. These variables account for the possibility of systematic overgrowth/undergrowth during the growth process. I tuned these three variables in order to fit the simulation's reflectivity to the measured reflectivity utilizing Matlab's built in least-squares curve fitting function.

However, in a problem such as this, many local minimums in the χ^2 function exist. Therefore, after running the least-squares function to find a local minimum, I then randomized the variables within 0.2% of their optimized values to generate the starting conditions for the next run, and ran the least squares fitting function again, in an annealing-type method. The percentage 0.2% used here can be termed the "aggressiveness;" I found 0.2% worked best through trial and error. An aggressiveness that is too large will waste a lot of optimization time as the starting conditions for each iteration are often very bad. Similarly, a low aggressiveness will converge towards a global minimum much more slowly. The algorithm repeats the randomization process a few hundred times, each time randomizing the best minimum found so far.

Varying only the thicknesses of the layers did not give satisfactory fits; the result was a simulation in which the photonic stopband was narrower than I measure. The width of a photonic stopband is directly proportional to the difference between the index of the two DBR materials [46]. Therefore, I also allowed the index functions of AlAs and AlGaAs to vary in the fitting procedure. I found that good fits could be obtained by simply introducing two scaling factors α_{AlAs} and α_{AlGaAs}, and modified the index functions as

$$\begin{aligned} n'_{AlAs}(\lambda) &= \alpha_{AlAs} n_{AlAs}(\lambda) \\ n'_{AlGaAs}(\lambda) &= \alpha_{AlGaAs} n_{AlGaAs}(\lambda), \end{aligned} \tag{4.18}$$

where n is the real part of the index function of the two materials. Similarly, I introduced scaling factors β for the imaginary piece of the index functions,

$$\begin{aligned} \kappa'_{AlAs}(\lambda) &= \beta_{AlAs}\kappa_{AlAs}(\lambda) \\ \kappa'_{AlGaAs}(\lambda) &= \beta_{AlGaAs}\kappa_{AlGaAs}(\lambda), \end{aligned} \tag{4.19}$$

where κ is the imaginary piece of the index function. Typically I found that the thicknesses of the layers needed to be scaled between 0.95 and 1.05 of their designed values, and the α variables need to be scaled between 0.95 and 1.05, in order to produce the agreement shown in Figure 50. However, as mentioned above, the β variables were typically between 0.01 and 0.001, likely caused by the high purity of our samples, resulting in low absorption.

My GUI is capable of fitting multiple reflectivity curves simultaneously. For example, three line profiles at 0, 15, and 30 degrees were taken from the data shown in Figure 50A, and are shown in Figures 50B, 50C, and 50D. These three curves are being fit with equal weighting. The GUI is quite versatile and capable of including in its fits profiles at various angles, temperatures, polarizations, and even other types of data such as ellipsometric and transmission. However, I found that including three profiles at cryogenic temperatures was sufficient for the fitting, and resulted in a reasonably fast optimization time. Its almost painful to have worked so hard to make such beautiful reflectivity data, and then only use three rows of the entire image, but the cun time for the optimization is a very serious consideration here. Letting the simulation run over night is enough time to produce the fits shown in Figure 50 using a standard modern computer.

Standard MATLAB measures were taken to optimize the code. Specifically it is critically important to minimize the frequency of calls to the basic linear algebra subprograms (BLAS), as the overhead cost of transferring data into the BLAS is significant. Instead, the problem should be formatted as a single large matrix, and executed as a single call if possible. The famous MATLAB advice to vectorize the code is in full effect here. Doing this properly reduces computation time to a fraction of an inefficiently coded execution. This has been discussed in more depth in Chapter 3.

I found that tuning the thicknesses of the layers and the index functions results in a simulation which largely agrees with our measured reflectivity over a fairly broad wavelength

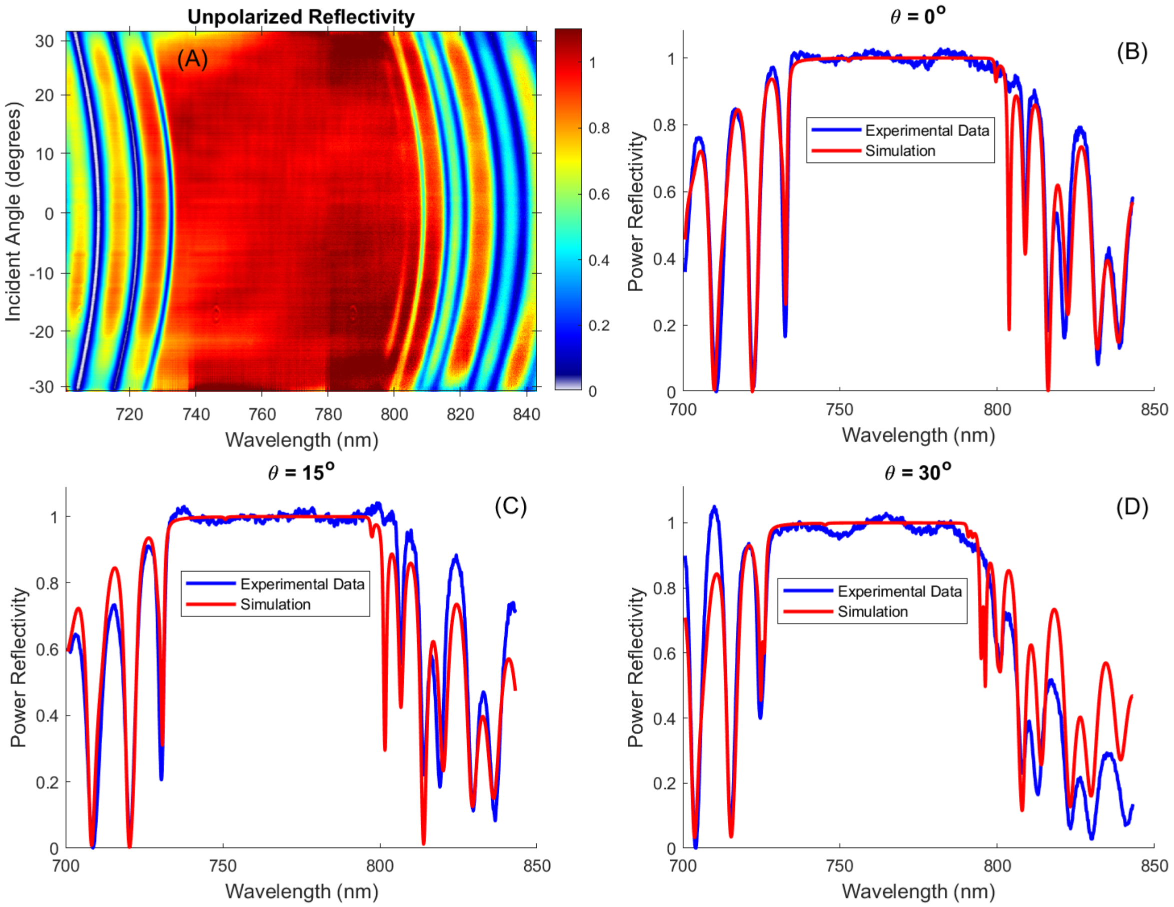

Figure 50: Image A: Reflectivity measurement of our long-lifetime samples. This image is calculated from four measurements as described in Equation (2.10). Images B, C, and D show the agreement between the TMM simulation and the data in image A after the fitting process at three different angles. There is reasonably good agreement. The discrepancy near 805 nm is likely caused by the bulk excitons in the substrate of the sample.

range. The fit is especially good at normal incidence, except at approximately 805 nm. This region is heavily influenced by the band gap of our GaAs substrate at cryogenic temperatures. Indeed, using a non-resonant laser, it is possible to see some PL from the substrate around 805 nm. When the sample is pumped with a white light source, some of the white light is likely absorbed in the substrate and scatters down to be re-emitted around 805 nm. This would serve to fill in the dip that should be present in the reflectivity. The simulation is not capable of accounting for absorption at one wavelength causing emission at another wavelength, and so I believe this explains the discrepancy.

It could in principle be possible to utilize an optical index with a negative imaginary component, which acts as a source of gain rather than loss, to effectively mimic light being emitted by those excitons near 805 nm. However, because physicists and engineers often use a flipped sign conventions for the imaginary piece of the index of refraction, sometimes published index values are negative rather than positive. To make things more user friendly, I've coded into my simulation a check which always forces the imaginary piece of the index to be positive. If a future student wishes to try this, they would need to go into the code and disable that, but it isn't hard to find in the code.

At higher angles, the numerical aperture of the system begins to limit the reliability of the data. The optics are capable of handling up to ±45 degrees of light cone, and room-temperature measurements of the reflectivity look quite good over this entire range. However, when imaging through our cryostat window it is not possible to collect good reflectivity data beyond 30 degrees.

I performed these measurements at various locations on the samples. For a rotated MBE growth, all layers should scale by the same percentage as point of observation moves outward radially. When I took the simulation that produces the curves in Figure 50, and scaled only the thickness of the entire sample uniformly, I found similarly good agreement with experimental reflectivity data measured at other locations on the wafers. This means that I only needed to perform the fitting procedure outlined above at one location on the sample, and then I had a working broad reflectivity simulation for all locations of interest on the sample.

4.5.2 Fitting The TMM Model To PL And PLE

With a TMM simulation that matches the broad-spectrum reflectivity measurements, I can now turn to the finer details governing the polaritons. Generally, these two parts of the fitting are somewhat independent. The variables tuned in the simulation in this section have almost no effect on the broader reflectivity results shown in Figure 50. Similarly, the tuning performed in Section 4.5.1 has a modest impact on the polaritons I will simulate in this section.

In Figure 50, no traces of the polaritons are seen in the experimental data nor the simulation. This is because the linewidth of the cavity is too narrow for the spectrometer to resolve, and the mesh of the simulation is too coarse to resolve them. However, if I significantly zoom in on the simulation results in Figure 50B, and increase the density of the simulation mesh, I obtain the red curve in Figure 51A. Only the bare photon mode is seen because so far, I have ignored the exciton resonance in the quantum wells inside of the microcavity. I must now determine how to include the excitons in the simulation.

Some works have accounted for the excitons in the index of refraction of the GaAs quantum wells [55]. I decided, however, to use a simple charged-oscillator model for the excitons [73]. The equation for the complex electronic susceptibility of the excitons is

$$\chi_{ex}(E) = A\frac{(E_{ex}^2 - E^2) + i\Gamma E}{(E_{ex}^2 - E^2)^2 + \Gamma^2 E^2}, \tag{4.20}$$

where E_{ex} is the natural energy of the exciton, Γ describes local damping but is also equal to the full width at half max (FWHM) of the exciton resonance due to homogeneous broadening, and A is an amplitude related the density of oscillators in our quantum well. These three parameters can be tuned in order to bring our simulation into agreement with both the lower polariton PL measurement and the upper polariton PLE measurement.

However, as I will discuss in Chapter 5, I also performed this fitting on the data obtained from low-Q samples where I directly measured both polariton branches through reflectivity. I consistently found that the curvature of the upper polariton of the simulation was larger than what I experimentally measured in reflectivity. That is, equivalently, the simulation produced an upper polariton with too small a mass compared to the experimental data. To fix this, I tried convolving Equation (4.20) with a Gaussian distribution in order to account for inhomogeneous disorder (i.e. a Brendel-Bormann oscillator). I also tried creating a finite series of exciton states to account for the 12 slightly different quantum wells in the sample.

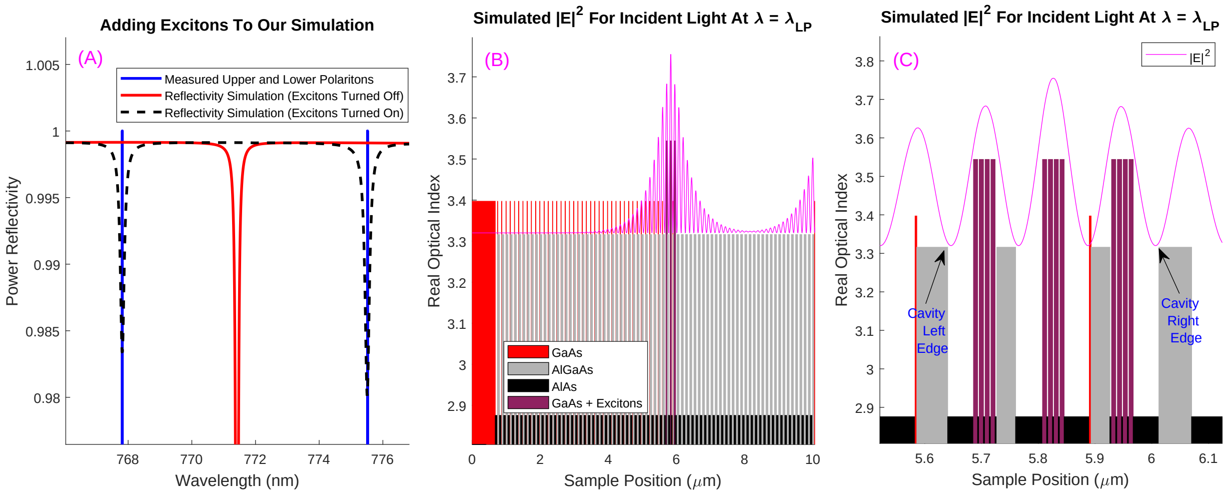

Figure 51: Image A: The simulation with the excitons turned off, and then with the excitons turned on. The cavity mode splits into our upper and lower polariton modes. The vertical blue lines are the experimentally measured values. Image B: The electric field plotted on top of a scale drawing of the sample. The electric field is confined within the microcavity. Image C: a zoom-in of image B, to highlight the antinodes of the electric field that lie on top of the quantum wells, which helps give our samples their strong coupling.

I also tried making E_{ex} dependent on angle in order to account for the non-infinite mass of the excitons. However, all these methods required unbelievable numbers in order to work.

Ultimately, I found the most convincing approach to be to recognize that a second exciton state exists at an energy approximately 15-20 meV above the first exciton, namely the light-hole exciton. This second resonance serves to bend the upper polariton downward into agreement with the measurements. This is a reasonable model, because as I discussed in Section 1.2, it is well known that there are both heavy-hole and light-hole excitons in the quantum wells, with heavy-hole excitons being the lower energy particle. In this picture with both heavy-hole and light-hole excitons present, the so-called upper polaritons should perhaps be called middle polaritons. However, for consistency, I will continue to call them upper polaritons as I will have no need to refer to the third polariton state which lies above the light-hole excitons.

I added the two excitons into the quantum wells by simply adding their susceptibilities to the susceptibility of the bulk GaAs already in the simulation:

$$\chi_{QW}(E) = \chi_{GaAs} + \chi_{ex,hh} + \chi_{ex,lh} \tag{4.21}$$

If all of the parameters of the two exciton resonances are allowed to vary freely, the system is underconstrained; therefore I imposed two conditions and still obtained good fits. I imposed that $A_{hh} = A_{lh}$ and that $\Gamma_{hh} = \Gamma_{lh}$, which means there are now four parameters governing the excitons in the quantum wells, namely the two exciton energies $E_{ex,hh}$ and $E_{ex,lh}$, and the amplitude and linewidth they both share, A and Γ. The two exciton energies are roughly known from the physics of GaAs bare quantum wells [10]; for the other two parameters, the simplest assumption is that the oscillator strengths and damping of the two exciton resonances are comparable, thus justifying setting them equal.

When I add just the heavy-hole exciton, I see that the single dip produced in the solid red curve of Figure 51A splits into two dips as shown in the dashed black curve. When I also include the light-hole exciton, a third dip is present at shorter wavelength outside the field of view. These dips are the polariton branches. The simulation may be run over a range of angles to create the polariton dispersion curves, as shown in Figure 52.

In order to produce good agreement between the simulation and the measured data over a range of angles, I found it necessary to introduce a fifth fitting term at this point. The layers which make up the cavity are illustrated in Figure 51C. I introduce a scaling term which uniformly scales the thickness of all non-quantum well layers inside the cavity, relative to the rest of the sample. Inside the cavity there are very many thin layers. Therefore, issues such as systematic overgrowth/undergrowth or extra disorder at the interface may be present, and this scaling term allows us to partially account for such effects.

With these five tunable parameters, I can now tune the simulation to match the measurements of the polaritons. Note that the measurements of polaritons in our high-Q samples are not reflectivity measurements, while the simulation is a reflectivity simulation. This means there are two shifts to consider. These are the difference of the peaks of the lower polariton reflectivity measurements and the PL measurements, and the difference of the peaks of the upper polariton PLE measurements and reflectivity measurements. In Chapter 5 I show that, by good fortune, these two shifts tend to produce errors that offset one another, and so they are not too concerning.

The PL and PLE data are relatively pristine peaks, but the TMM simulation gives dips on top of a background stopband which may have some shape to it. In order to compare the two, at each incident angle of the simulation (each column in Figures 52 A and B) I fit the polariton dips with a Lorentzian line shape subtracted from a background function. Typically the background is a second-order polynomial. I performed this at every angle until I had extracted the dispersion curve of the lower polariton $\lambda(\theta)_{lp,sim}$. I similarly performed the same routine on the upper polariton to obtain $\lambda_{up,sim}$, recalling that for high-Q samples I only have to perform this at $\theta = 0$ corresponding to the PLE measurement.

In order to fit the simulation to the data, I need to define a metric to minimize. First there is the error associated with the lower polariton:

$$e_{lp} = \frac{1}{n_\theta} \sum_\theta |\lambda(\theta)_{lp,sim} - \lambda(\theta)_{lp,pl}|, \tag{4.22}$$

where n_θ is the number of terms within the summation. Similarly there is the error for the upper polariton:

$$e_{up} = |\lambda_{up,sim} - \lambda_{up,PLE}|. \tag{4.23}$$

A best fit will involve the minimization of both of these errors; however, I found that sometimes the error of the lower polariton can dominate the upper polariton. So I introduced a weighting factor α_{up}, which enters into the total error as

$$e_{total} = e_{lp} + \alpha_{up} * e_{up}. \tag{4.24}$$

Through trial and error I have found that $\alpha_{up} = 5$ works well. Users who download my code and use it for their own characterization may wish to experiment with this number, which can be easily set in the GUI. The fitted simulation for the high-Q sample is shown in Figure 52B. The inverted power reflectivity $1 - R$ is plotted in the background image with a tight color scale to make the polariton dips in the stop band visible, and the LP PL data and the PLE data point are plotted on top as bars. The size of the bars corresponds to the FWHM of the data, so the reader may compare the linewidths of the simulation to the linewidths of the data. There is very good agreement overall.

I experimented with alternative forms for Equation (4.22) with the goal of making the linewidth of our simulation match the linewidth of the experimental data. However, I found this detrimental. The simulation can not capture several essential factors affecting the linewidths. For example, high-momentum polaritons may scatter down into lower momentum states. Additionally, upper polaritons may scatter down into lower polariton states. These effects will result in line broadening [72]. Hence, I believe the simulation will always undershoot real linewidths. The Γ fitting parameter has almost no effect on the shape of the dispersion curves, and almost exclusively determines the linewidths of the polaritons. So, I removed it from the pool of fitting parameters and fix it to 0.5 meV, which is in good agreement with measurements I performed on bare GaAs quantum well samples shown in figure 6. I would expect the linewidth of the lower polariton near $\theta = 0$ to be the least tainted by these scattering effects, and indeed I generally find the linewidth of the simulation to be in modestly good agreement with the PL measurements in that region.

With the TMM simulation now fitted to the data sets, I may use Equation (4.25) to extract the exciton fraction via a numerical derivative. I did this by shifting the exciton energies E_{ex} up and down by a small amount ΔE_{ex}, then I extracted the corresponding curves of the lower polariton from the simulation, E_{lp_+} and E_{lp_-} respectively. Finally, the

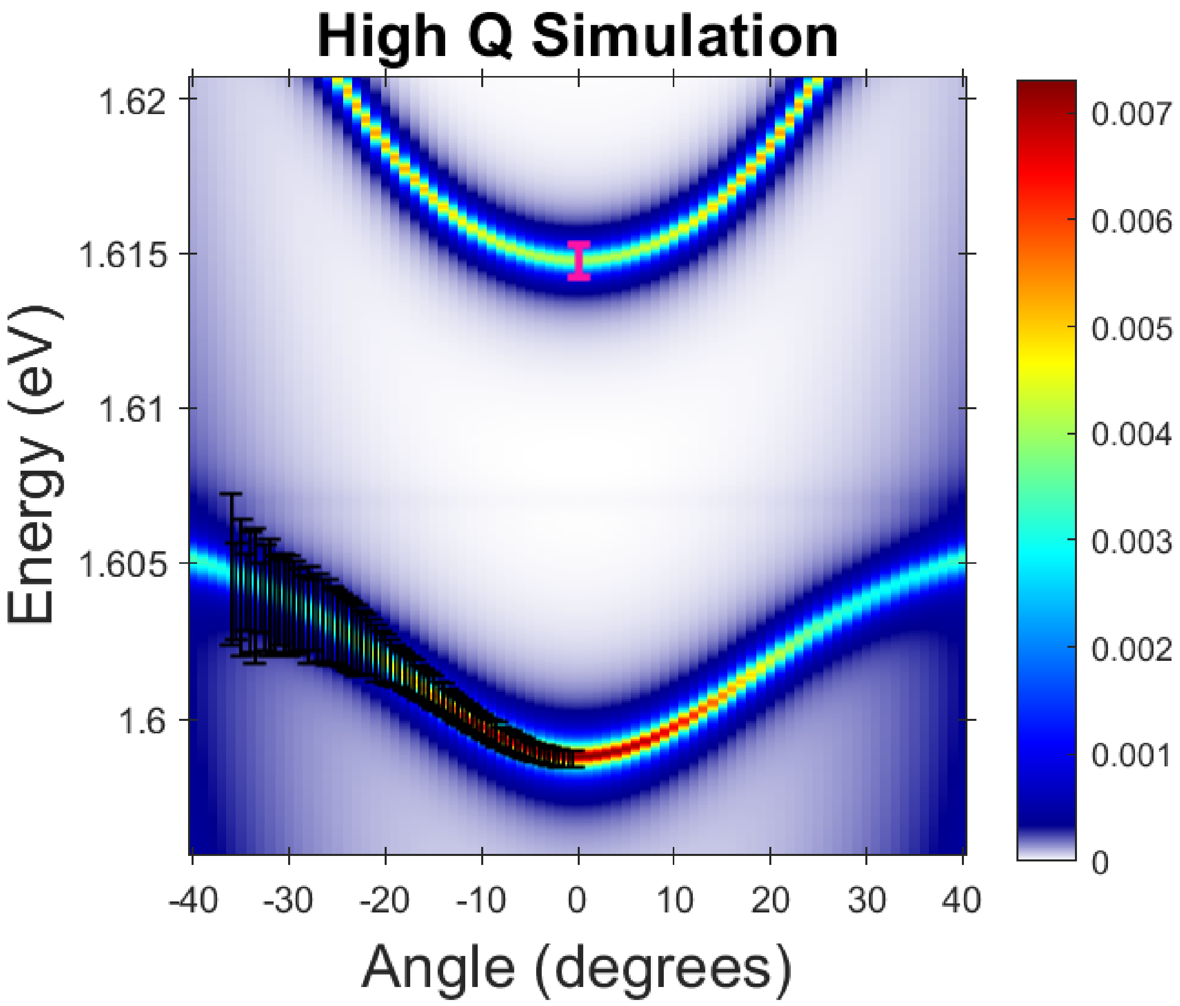

Figure 52: In the background, the unpolarized inverted reflectivity $1 - |r|^2$ from the fitted TMM simulation of the high-Q samples with a tight color scale to make the polaritons more visible. Plotted on top of the upper polariton is the PLE measurement. Plotted on top of the lower polariton is half the lower polariton data as measured through photoluminescence via non-resonant pumping. The width of the bars represents the FWHM of the measurements.

exciton fraction may be calculated:

$$|X_{k_\parallel}|^2 = \frac{d}{dE_x} E_{lp}(k_\parallel) \approx \frac{E_{lp_+}(k_\parallel) - E_{lp_-}(k_\parallel)}{2\Delta E_{ex}}. \tag{4.25}$$

As ΔE_{ex} becomes smaller, the simulation mesh needs to become smaller to accurately extract the derivative of the lower polariton curves. I found $\Delta E_{ex} = 0.1$ meV to work well enough. Additionally, because there are two excitons in the simulation, I performed this numerical derivative on both excitons separately to arrive at the heavy-hole and light-hole exciton fractions $|X_{k_\parallel,hh}|^2$ and $|X_{k_\parallel,lh}|^2$. I found that the light-hole exciton fraction is generally quite small, which is sensible as the light-hole excitons are further away from the lower polaritons. For the heavy-hole exciton, I saw the expected behavior; namely, the lower polaritons are more excitonic at larger angles.

The quantity of interest for most calculations is the total excitonic fraction at a given location on the sample, and at a given in-plane momentum, which can be written as

$$|X_{k_\parallel,total}|^2 = |X_{k_\parallel,hh}|^2 + |X_{k_\parallel,lh}|^2. \tag{4.26}$$

Specifically, the value at normal incidence $|X_{0,total}|^2$ is the final result to report. Thus, the problem is solved. I've shown how to extract the exciton fraction from the three measurements; PL, PLE, and reflectivity. However, the process is a bit cumbersome. In Section 4.6 I will show how I created fitted curves that can easily be shared with our collaborators.

4.6 Final Result: The Characterization Curve

I've now shown how to extract the exciton fraction from TMM simulation after fitting it to the PL, PLE, and reflectivity data. I performed those measurements and the associated TMM fitting at many locations on many samples. The results for nine locations on sample 4-6-15.1 are shown in Figure 53A. The numerical derivative method gives an exciton fraction at all angles simulated. I inverted the dispersion relationship $E(\theta)$ to express $|X|^2(\theta)$ as $|X|^2(E)$ at each of those locations, and the curves are shown in Image B.

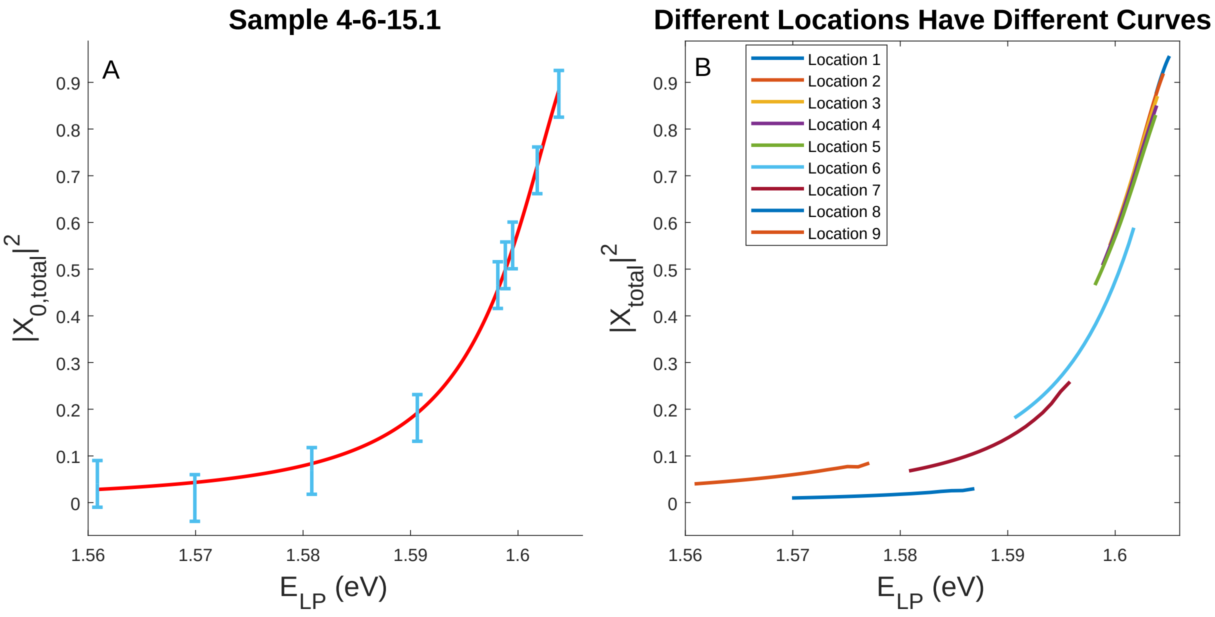

Figure 53: A: For nine locations on sample 4-6-15.1, the results from the PL, PLE, and reflectivity measurements described in this thesis, and then fit with the TMM simulation to extract the total exciton fraction at $k_{\|} = 0$. These nine data points are fit with Equation (4.29). The values of the two fitting parameters are given in Table 4. B: The exciton fraction curves created through the numerical derivative method at finite angles using the simulations of the nine locations.

The characterization problem is now fully solved. By performing a reflectivity measurement, lower polariton PL measurement, and upper polariton PLE measurement I have shown how to extract the total exciton fraction at a single location on the sample, and we have shown this method works over a wide range of detunings. However, not every lab has access to the tools to perform all three measurements; furthermore it is not always prudent to perform all three measurements. For these reasons, I wished to create a simple fit of the data so that the exciton fraction can quickly be estimated from just a measurement of the lower polariton's energy. The two-level model from Section 4.4 gives a functional form to try. I begin by manipulating equation 4.8 to solve for E_c:

$$E_c = E_{LP} + \frac{g^2}{E_x - E_{LP}}. \tag{4.27}$$

I am specifically working with $k_{\parallel} = 0$ so I dropped the k-dependence in the notation. The Rabi model gave an expression for the exciton fraction, written as

$$|X|^2 = \frac{1}{2}\left(1 - \frac{E_x - E_c}{\sqrt{4g^2 + \left(E_c - E_x\right)^2}}\right). \tag{4.28}$$

Equation (4.27) may be plugged into equation (4.28) in order to arrive at

$$|X|^2 = \frac{1}{2}\left(1 - \frac{E_{x,global} - E_{LP} - \frac{g^2_{global}}{E_{x,global} - E_{LP}}}{\sqrt{4g^2_{global} + \left(E_{LP} + \frac{g^2_{global}}{E_{x,global} - E_{LP}} - E_x\right)^2}}\right). \tag{4.29}$$

This may now be viewed as a fitting equation for the total exciton fraction at normal incidence as a function of the lower polariton energy, where the exciton energy $E_{x,global}$ and the coupling strength g_{global} are now fitting parameters. This functional form is a two-parameter fit that I found works quite well, as it automatically imposes limiting behavior that should be present in the fit. However, physical meaning should not be inferred from the values of the fitting parameters; for this reason I have attached the "global" subscript to help keep this in mind for the reader. I then tuned the fitting parameters and arrive at the best fit as shown by the solid line in Figure 53. The fit works quite well. The fitted parameter values are given in Table 4 along with the fitted parameters for many other samples on which I have performed similar measurements and fits. This table has been made available online for our collaborators [13]. Many of the samples are from our collaborators at Princeton, however our collaborators at Waterloo have also begun producing high-Q samples as well.

The reader should note that this fit applies to the polariton energy at normal incidence only. At a single location, the polaritons at higher in-plane momentum will also have higher energy, but the exciton fraction at that higher in-plane momentum is not quite governed by the fit here. These curves are shown in figure 53B. It is apparent that they are similar to the global fit in Figure 53B, but they do diverge slightly.

Table 4: The fitting parameters corresponding to equation 4.29 at cryogenic temperatures for some of the many different samples I have used and shared with collaborators over the years. Note that these parameters are just fitting parameters and do not have direct physical meaning. Most of the samples are high-Q factor, consisting of a 40-period bottom DBR and a 32 period top DBR. However, there are a few low-Q samples consisting of a 20 period bottom DBR and a 16 period top DBR. All the samples are 3/2 λ cavities, and most have 4 quantum wells per antinode for a total of 12 quantum wells. However, one sample had only one quantum well per antinode for a total of 3 quantum wells. Most of the samples were grown by our collaborators in Princeton, but in the last few years our collaborators at Waterloo have begun growing comparable samples. The high-Q samples are all similar designs, with slight modifications to thicknesses of the various layers.

Sample	$E_{x,global}$ (eV)	g_{global} (meV)	Grower	Q factor	# QWs
4-6-15.1	1.6066	7.8051	Princeton	High	12
P8-10-17.1	1.6091	12.8099	Princeton	High	12
P3-9-18.1	1.6134	18.8152	Princeton	High	12
P6-19-19.1	1.6099	9.5462	Princeton	High	12
P9-10-20.2	1.6060	11.7454	Princeton	High	12
P9-28-20.1	1.6089	7.7175	Princeton	High	12
P9-29-20.1	1.6072	12.5079	Princeton	High	12
P1-9-23.1	1.6102	12.5866	Princeton	High	12
P1-10-23.1	1.6138	11.6257	Princeton	High	12
4-20-05.1	1.6132	11.4051	Princeton	Low	12
10-5-10.1	1.6050	6.3032	Princeton	Low	3
G0792	1.6196	11.5415	Waterloo	High	12
G0921	1.6137	10.6059	Waterloo	High	12
G0985	1.5703	7.8086	Waterloo	High	12

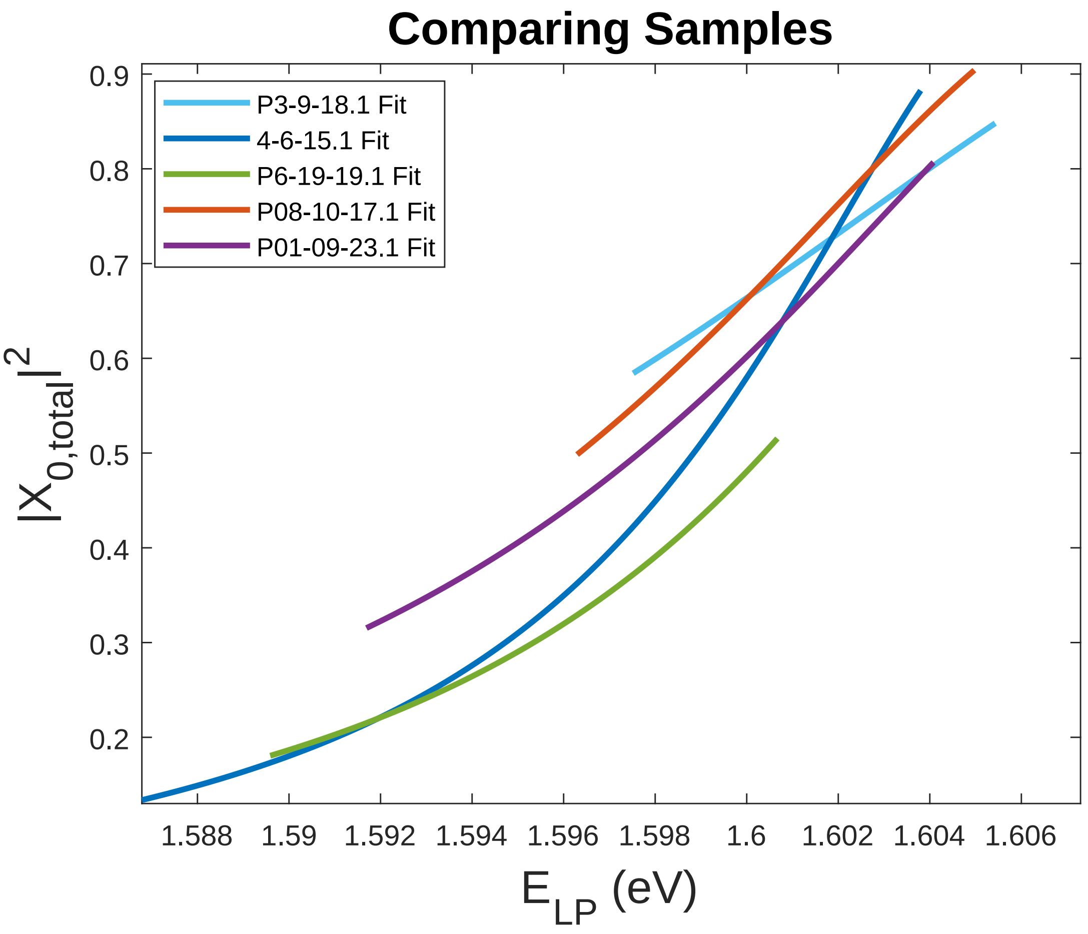

Figure 54: The fitted exciton fraction for several of our samples. We see that the characterization procedure must be performed on each individual sample. They do not share a global curve.

An important feature to summarize from the results in Table 4 is that the final characterization curves are not universal from sample to sample. This is somewhat surprising. Most of the samples have design specifications which are proportional scalings of one another. Additionally, it is believed that as the observation location on the sample moves radially outwards, all layers should proportionally scale down. Therefore, I would have expected all samples to lie on a single, global curve. But that is not the result I found. Figure 54 shows the fitted characterization curves of several such samples. The curves clearly intersect. This leads me to conclude that in the MBE growth process, not all layers of the sample are scaling proportionally with radial position. Specifically I suspect the quantum well thickness may be much flatter than the DBR layers.

Lastly, in Figure 40 and 44 I report the temperature dependence in the PL and PLE measurement at a few locations on Sample P01-10-23.1. I also performed the corresponding reflectivity measurements and TMM fitting in order to investigate the effect sample temperature has on these results. A summary of these results are shown in Figure 55.

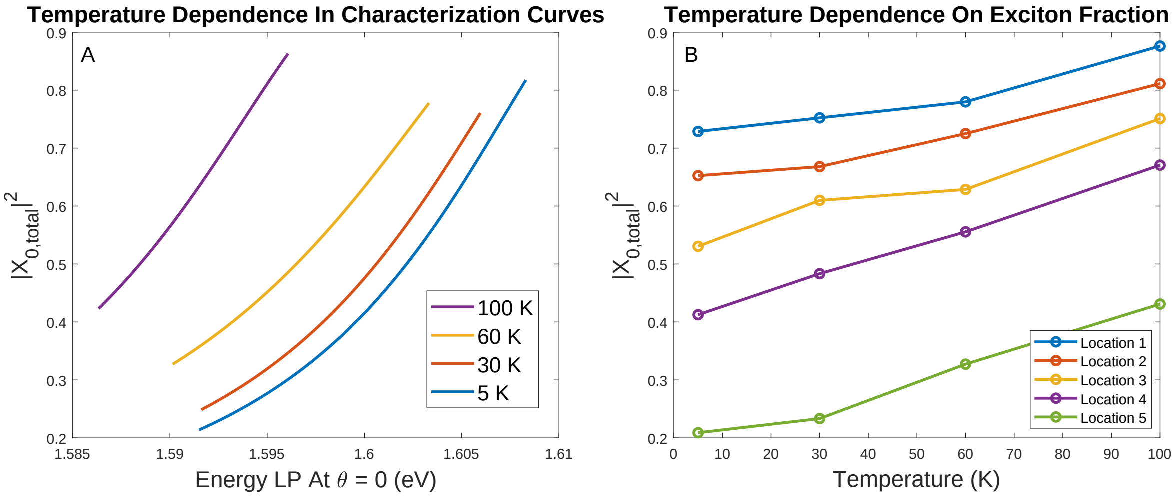

Figure 55: Image A: The fitted exciton fraction curves for a sample at several temperatures. Image B: The exciton fraction at five locations as a function of temperature.

Image A shows how temperature influences the characterization curve. In general, higher temperature causes the curve to move up and to the left. Image B shows the effect of temperature at a specific location. The data from five different locations is shown. The exciton fraction has a direct relationship with temperature, increasing as temperature increases. The limited data makes the exciton fraction look fairly linear with temperature.

Although I have stated above that the global fit parameters in Table 4 do not have direct physical meaning, it is still interesting to look at their temperature dependence. These plots are shown in Figure 56. The $E_{x,global}$ data in Image A has a clear functional form to it. Because exciton energy should track with bandgap energy, I tried fitting it with a Varshni equation [73], written as

$$E_x = E_o - \frac{\alpha T^2}{T + \beta}. \tag{4.30}$$

Here, E_o, α, and β are taken as fitting parameters. However, this means I have three fitting parameters and only four data points, so there is healthy skepticism of this fit. The fitted Varshni parameters for $E_{x,global}$ of sample P01-10-23.1 are $E_o = 1.6144\text{eV}$, $\alpha = 19.3\text{meV/K}$, and $\beta = 15000\text{K}$. These numbers are not particularly close to reported values for GaAs [37]. Additionally, the behavior of g_{global} does not seem to have any functional form. Overall, I conclude that the global fitting parameters are just fitting parameters, and no obvious physical significant should be inferred from them.

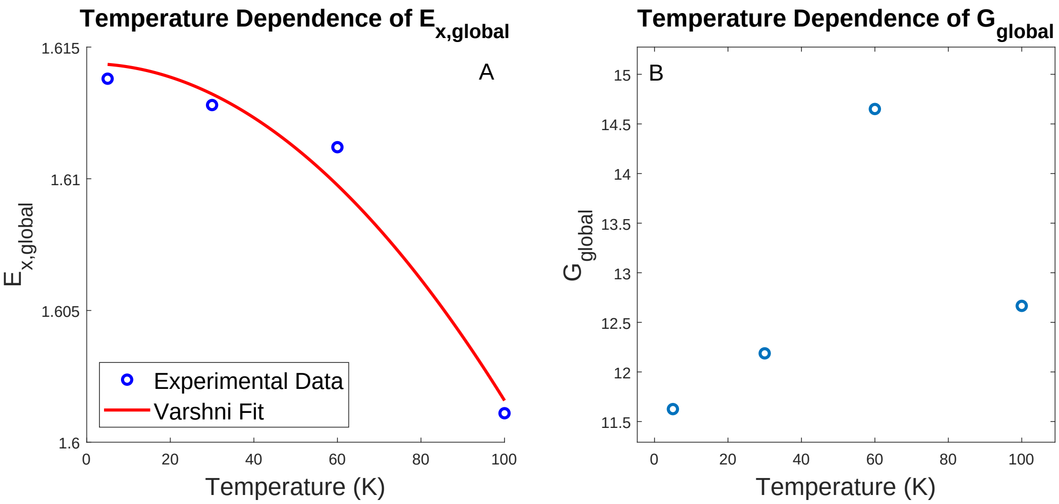

Figure 56: The temperature dependence for the global fitting parameters in sample P01-10-23.1

4.7 Predictive Power Of The TMM Simulation

One of my stated goals with the model is to be able to accurately predict how future samples will behave in order to guide future designs. Here I will briefly discuss how well the TMM model works for this purpose.

In order to test this, I started with the fitted simulation for the data sets in the middle of our range. That is, I started with the simulation that is not the most photonic, nor the most excitonic, but in the middle near $|X_{o,total}|^2 = 0.5$. I then uniformly scaled every layer of the simulated sample so that the simulated reflectivity curve would approximately match the measured reflectivity curve at the other locations on the sample.The reflectivity curves are very reliable means of measuring the sample thickness. Thus, this should be viewed as tuning the simulation to the thickness of the sample at other locations.

I then used the numerical derivative method outlined above to extract the total exciton fraction. This number is then compared to the result I obtained when I explicitly fit the simulation to the PL,PLE, and reflectivity data at the second location. The comparison between the simulation's predicted exciton fraction and the exciton fraction I extracted by explicitly fitting the data is shown in Figure 57.

In Figure 57, the location the simulation is tuned to is denoted as a sample thickness of 1.0, and all the other locations are measured relative to that. We see that the simulation's accuracy is reasonably good when scaling by a few percent. Notice that the exciton fraction ranges from 0.9 to 0.1 in only about 4% of relative thickness variation. So when designing and modifying samples, typically I am only looking to modify the thickness by a few percent.

Overall, this is a very useful result. If, for example, we have a sample that is 75% excitonic at the center and we wish to design a new sample that is 50 % excitonic at the center, we can use this to guide the redesign of the sample for future growth. Unfortunately, as indicated in Figure 54, different samples are not all on the same curve. I believe this ultimately is attributable to inconsistencies in the MBE growth; unfortunately, no model can account for that. However, this model does give a starting point that I have used successfully to modify sample designs for desired outcomes.

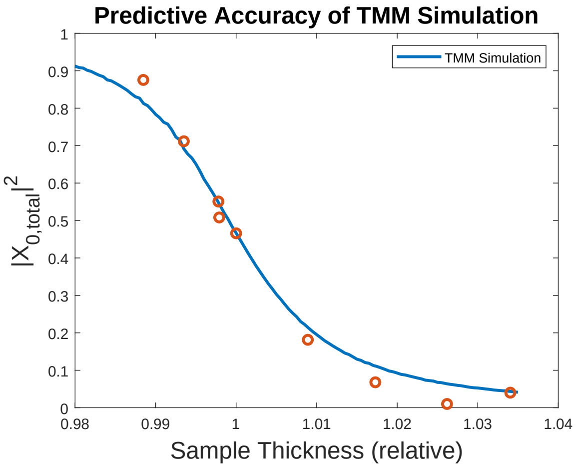

Figure 57: The nine data points are from our full TMM fitting of the reflectivity, PL, and PLE data as described in section 4.5.2. The solid line is generated by taking our fitted simulation at the center (denoted as a thickness of one) and uniformly scaling every layer of the simulated sample, and then extracting our exciton fraction. We see that once the simulation has been fit to a location on the sample, it may be used to predict the polaritons at other locations reasonably well. This may also be utilized to predict the polaritons produced by new sample designs.

5.0 Short-Lifetime Sample Characterization

In Chapter 4, I discussed the methods for characterization of long-lifetime samples, including the difficulties that arise in them. In particular, in long-lifetime samples the polariton branches may not be observed in reflectivity measurements. Additionally, the upper polariton's PL is essentially invisible. However, in short-lifetime samples, these issues are not present. As a result, short-lifetime samples provide a great opportunity to confirm that the methods I deployed on long-lifetime samples are valid.

In this chapter, I will give an overview of the PL, PLE, and reflectivity measurements I performed on short-lifetime samples. Then I will explore the results and discuss how they guided my decisions in the development of the long-lifetime methodology. The results presented in this chapter are from sample 4-20-05.1 from the Pfeiffer group. This sample was the sample used in Reference 7, and the design is similar to other samples from the era. This also serves as a testament to the remarkable stability of these samples, as nearly twenty years of room-temperature conditions have not significantly degraded or oxidized them.

5.1 Sample Structure

The short-lifetime samples are very similar to the long-lifetime samples. They key difference is just the number of periods of repetition for the top and bottom DBR. In the long-lifetime samples the bottom DBR had 40 periods and the top DBR had 32 periods. In the short-lifetime samples, the bottom DBR has 20 periods and the bottom DBR has 16 periods. A simple illustration showing the design of a short-lifetime sample is shown in Figure 58.

In this illustration, the widths of the bars are proportional to actual design specifications. These samples are grown on a two-inch GaAs wafer, not shown in this simple illustration. In the long-lifetime samples, a 3 nm capping layer of GaAs is present as the topmost layer. The short-lifetime samples are older, before that practice became standard, and so the

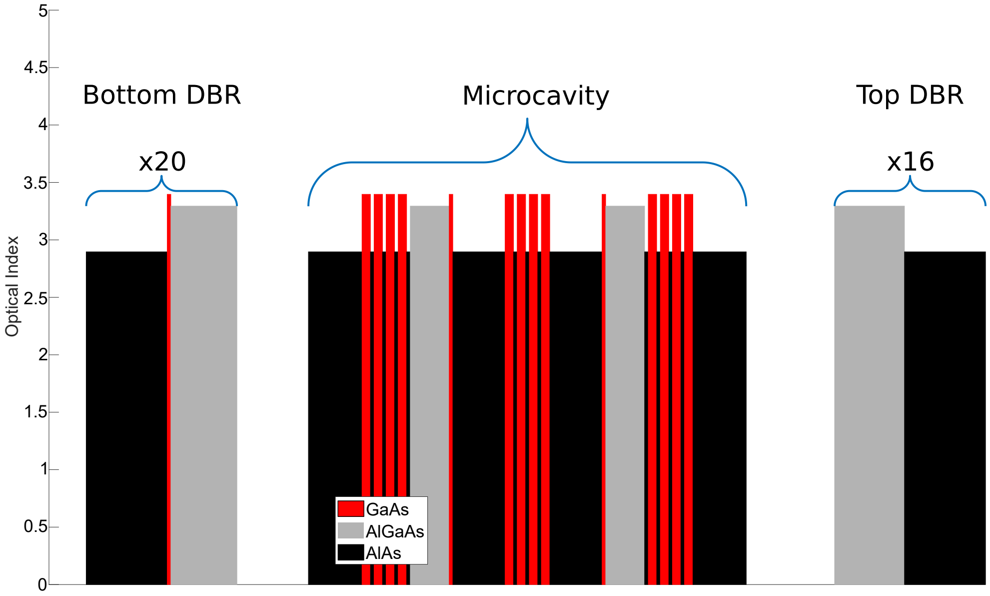

Figure 58: A simple illustration of a short-lifetime sample.

capping layer is not present. The capping layer's purpose is to prevent oxidation of the AlAs. However, all of the samples, both long and short lifetime, have AlGaAs as the final DBR layer. As such, even in the short-lifetime samples, pure AlAs is not directly exposed to the atmosphere. Oxidation is readily apparent when it occurs, and can easily be seen both in microscopic images and by human eyesight. The short-lifetime samples were inspected, and no oxidation had occurred.

Just as in the long-lifetime samples, there are excitons associated with the wetting layer of GaAs between the DBR layers, and once again they are at high enough energy to not be of consequence.

5.2 Measurements

In this section I will show typical data of the measurements performed on the short-lifetime samples. The key differences is the ability to observe polaritons in reflectivity measurements, and the upper polariton being observable in the PL measurements.

I begin with non-resonant pumping of the sample in order to measure PL. Care was taken to make sure that the samples were not heated up by the laser. Because the short-lifetime samples generally have lower reflectivity, heating by the laser is more of a concern. Lower laser powers or a mechanical chopper may be utilized to check and make sure red shifting is not occurring. Examples of these PL measurements are shown in Figure 59.

The PL of the upper polariton is faintly visible at the higher energy (shorter wavelength) side of these three images. Additionally, some PL from uncoupled excitons is present in between the two polariton branches, in the faint vertical line at constant wavelength; this line shows no dispersion because the effective mass of the bare excitons is orders of magnitude heavier than the mass of the polaritons.

In the long-lifetime samples, extraction of the dispersion curve $E_{lp}(\theta)$ was quite simple. A simple Lorentzian fit was utilized for each row independently, and this fit was straightforward at all angles. This method still works for the lower polariton in short-lifetime samples. However, the upper polariton is quite faint, and so extraction from one row at a time becomes

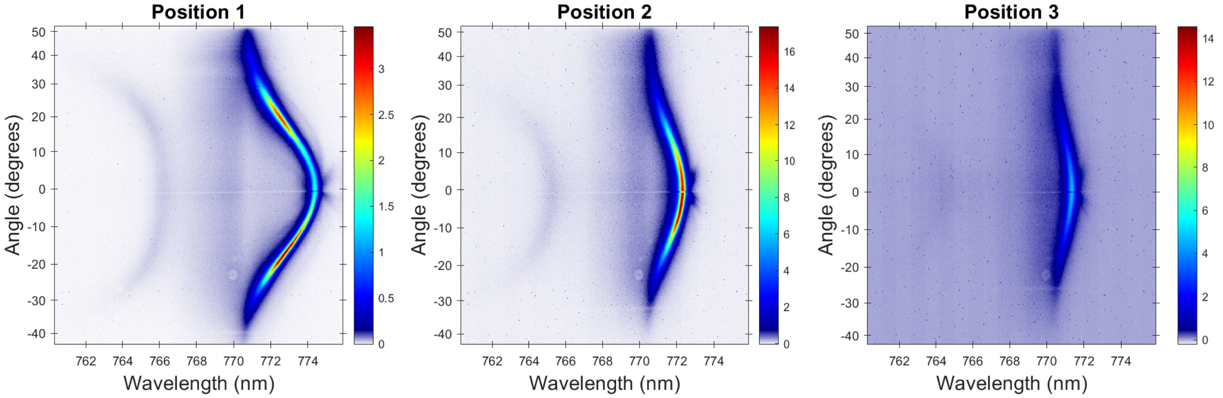

Figure 59: The PL measured from three locations on a short-lifetime sample.

quite difficult. I needed to develop additional data processing methods for the short-lifetime samples. The first method may be termed "equalization." This technique is mostly used to guide the human eye during further data processing methods. What equalization does is independently scale every row of the image to range between zero and one. This can be extremely helpful in helping to identify the curves in these images. An example of this method applied to the PL data from Position 3 is shown in Figure 60A.

After the image has been equalized, it is fairly easy to identify the upper polariton's location. I then used this knowledge to trim the original unequalized image down to isolate the upper polariton's PL, as shown in Figure 60B. A tight color scale is necessary to resolve the upper polariton properly. The task now is to extract the dispersion curve of the upper polariton from this image. Extraction for one row of pixels at a time is still not feasible. Why then is it possible for our eyes to see the upper polariton curve? It is because our eye and our brain do not process the image one row at a time. We process the entire image and see the pattern within the noise. In order to extract the curve properly, I needed to devise a two-dimensional fit of the data, rather than a row-by-row peak fit for $E_{up}(\theta)$.

In order to do this, I turned to the Rabi model, specifically with complex energies to get linewidths. The complex exciton energy is now two fitting parameters, the real and imaginary

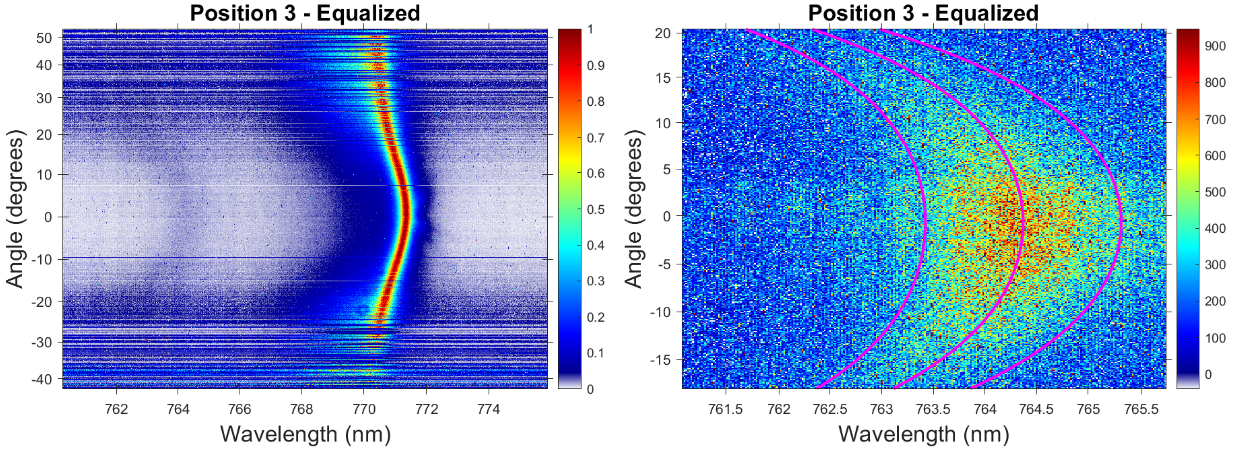

Figure 60: Image A: The PL data from Position 3 after being equalized. Image B: The original unequalized data, trimmed down to focus on the upper polariton, with a tight color scale. The pink lines are the fitted curve.

parts of $E_{ex,c} = E_{ex} + i\gamma_x$, where I use the subscript c to denote complex-valued quantities. The quantity γ is the full width at half maximum of the quantities it is attached to. I still approximate that the exciton has infinite mass, but now it has a linewidth independent of angle too. The complex cavity photon energy is then given by

$$E_{cav,c}(\theta) = \frac{E_{co}}{\sqrt{1 - \left(\frac{\sin\theta}{n_c}\right)^2}} + i\gamma_{cav}. \tag{5.1}$$

As an approximation, the cavity mode is given a linewidth independent of angle as well.

Now the upper polariton's dispersion relation is

$$E_{up}(\theta) + i\gamma_{up}(\theta) = \frac{1}{2}\left(E_{ex,c} + E_{cav,c} + \sqrt{4g^2 + (E_{ex,c} - E_{cav,c})^2}\right). \tag{5.2}$$

This dispersion must be turned into a two-dimensional function $Z_{fit}(E, \theta)$. The data is collected at set wavelengths, so converting back and forth between wavelength and energy is no problem. However, since the data is linearly spaced in wavelength due to the spectrometer, it is better to convert the fit from energy to wavelength, rather than to convert the data; although over such a small wavelength range it is not too much of a difference either way.

A Lorentzian distribution centered at $E_{up}(\theta)$ was placed in each row of the image. This may be expressed mathematically as

$$Z_{fit}(E,\theta) = A(\theta)\left(\frac{(\gamma/2)^2}{(E - E_{up}(\theta))^2 + (\gamma/2)^2} - B(\theta)\right). \tag{5.3}$$

This can be converted to wavelength dependence easily by just modifying the input argument as $Z_{fit}(hc/\lambda,\theta)$ where hc is 1239.8 ev-nm, which can be written simply as $Z_{fit}(\lambda,\theta)$. The $A(\theta)$ and $B(\theta)$ terms are a shift and scaling of the fitting function. This is necessary because the experimental data has a distribution scaling its rows of data relative to one another, as well as a DC background term. If the experimental data had a good signal-to-noise ratio, one could simply normalize each row of the image to scale between zero and one, and set $A(\theta) = 1$ and $B(\theta) = 0$. However, the signal to noise is quite weak, so this will not work well. Instead, $A(\theta)$ and $B(\theta)$ are chosen at each angle θ so that the difference between $Z_{fit}(\lambda,\theta)$ and the experimental data $Z_{exp}(E,\theta_i)$ is minimized in a least squares fit. Essentially, I scaling the fit function to have the same thermal distribution as the experimental data.

The problem of solving for $A(\theta)$ and $B(\theta)$ is actually solvable in closed form. A derivation of the solution is given in Section 5.2.1.Now, with the A and B coefficients defined, the total error between the fit and the simulation may be written as

$$\text{Error} = \sum_{\theta}\sum_{\lambda}(Z_{fit}(\lambda,\theta) - Z_{exp}(\lambda,\theta))^2. \tag{5.4}$$

The fitting parameters which lead to Equation (5.2) are now chosen to minimize this error. It is best to start the two-dimensional fit with linewidths that are too large, and allow the optimizer to narrow them down. The result is a fitted curve $E_{up}(\theta)$. The results of this fit are shown in Figure 60B as the pink lines plotted on top of the data. The central line is the dispersion relation $E_{up}(\theta)$; the other two pink lines represent the full width at half maximum of the two dimensional fitting function $E_{fit}(\lambda,\theta)$. One can see that the fit has done a pretty good job of pulling the dispersion relation out of the noise.

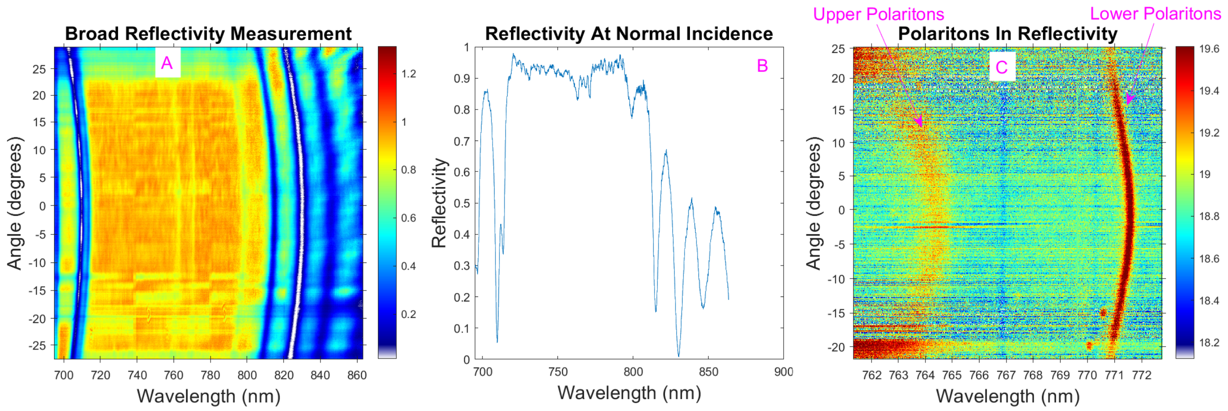

Figure 61: Image A: A broad reflectivity measurement in a short-lifetime sample. Image B: The profile across Image A at $\theta = 0$. Image C: A zoomed in image of the polaritons inside the stopband. The image intensity has been inverted with a tight color scale to make the polaritons more visible. The vertical lines between the two types of polaritons are from the substrate.

Next, we can turn to reflectivity measurements for the same sample, shown in Figure 61. In Image A, the polariton dispersion curves are faintly visible in the stopband. The profile at normal incidence shown in Image B also makes the dips visible.

In Image C I have zoomed in, inverted the image, and performed a few data processing techniques to make the polaritons more easily visible. The first thing I did was fit each row with a function in order to remove as much of the noise and variation between rows of the image as possible. In this case, I fit each row with two Lorentzians added to a second-order polynomial. The Lorentzians fit the polariton peaks, and the polynomial fits the underlying stopband. Constraining the range of the line centers of the two Lorentzians is very helpful. Usually one can see by eye what the lower and upper bounds of the polariton curves are, so I used that to restrict the bounds on the fitting parameters. After each fit was performed, the polynomial was subtracted from the row of pixels. One can think of this as putting the polaritons all on a flat plane, which makes them easier to see in the image. Afterwards, the equalization technique described above may also be useful. These techniques can be helpful in reflectivity data because sometimes the stopband is not so flat, or at least the measurement of it isn't. This technique would not have been as helpful in the PL data above, because in that data there is no stopband to remove. The image as I have prepared it is easily fit with the two-dimensional function $Z_{fit}(\lambda, \theta)$ in order to extract the dispersion curves of the polaritons.

The vertical lines between the upper and lower polaritons in Image C are caused by the substrate, not the quantum wells. I naturally first thought they may be exciton lines. However, as I move vast distances across the sample, they never move. That is to say, the thickness of the quantum wells does not seem to influence them. Further more, when moving across the sample, one can watch the polaritons curves move through them with no sort of coupling showing up. In long-lifetime samples, the top and bottom DBR are much thicker, making anything from the substrate harder to see.

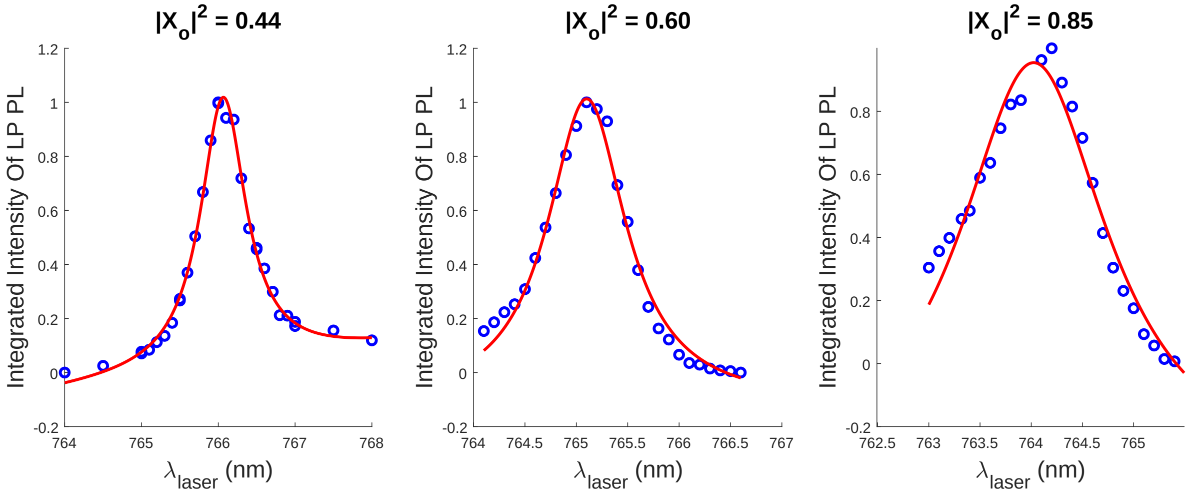

Figure 62: PLE results for the upper polariton in short-lifetime sample at three different locations. Fits to extract line center are shown as solid lines.

Finally, I performed PLE measurements in the short-lifetime samples. Results from PLE measurements on three different locations on the sample are shown in Figure 62. The titles give the exciton fraction at that location as measured through reflectivity measurements, discussed in Section 5.5. PLE works well in the short-lifetime samples. This is good news, as it will allow me to compare PLE measurements to reflectivity measurements in Section 5.3.

5.2.1 The Shift Scale Optimization Problem

Here I present a derivation of a result I used in Section 5.2 to help extract dispersion curves using two dimensional fits. Suppose that we have some experimental data $\vec{E}$ and some simulation data $\vec{S}$, and we wish to shift and scale the simulation data so as to minimize the difference between it and the experimental data in a least-squares fit. In this case, I write the data as a vector, but an array of any finite size and dimensions may be written in a one-dimensional form, so the method used here is quite general. We seek to find shift and scale coefficients A and B to minimize the quantity

$$E = \sum_{i}^{N} (A(S_i - B) - E_i)^2. \tag{5.5}$$

This quantity may be written in vector notation as

$$E = (A(\vec{S} - B\vec{U}) - \vec{E}) \cdot (A(\vec{S} - B\vec{U}) - \vec{E}), \tag{5.6}$$

where $\vec{U}$ is a unit vector, a vector with one for every component. Next, the dot product is expanded out, leaving

$$E = A^2(S^2 - 2B\vec{S} \cdot \vec{U} + B^2U^2) - 2A(\vec{S} \cdot \vec{E} - B\vec{U} \cdot \vec{E}) + E^2. \tag{5.7}$$

The goal is to choose A and B to minimize E, so one can utilize the partial derivatives. Taking the partial derivative with respect to B and setting it equal to zero results in

$$0 = A^2(-2\vec{S} \cdot \vec{U} + 2BU^2) + 2A\vec{U} \cdot \vec{E}. \tag{5.8}$$

One can see there are two solutions here. $A = 0$ corresponds to a maximum. The other solution is the useful one, given aby

$$A = \frac{\vec{U} \cdot \vec{E}}{\vec{S} \cdot \vec{U} - BU^2} = \frac{\vec{U} \cdot \vec{E}}{\vec{S} \cdot \vec{U} - BN}, \tag{5.9}$$

where I have utilized the fact that U^2 is equal to the length of the vectors N. Next, one can take the derivative of E with respect to A and set it equal to zero, which gives another equation for A,

$$A = \frac{\vec{S} \cdot \vec{E} - B\vec{U} \cdot \vec{E}}{S^2 - 2B\vec{S} \cdot \vec{U} + NB^2}. \tag{5.10}$$

Equations (5.9) and (5.10) may be set equal to one another in order to solve for B, which leads to

$$B = \frac{(\vec{U} \cdot \vec{S})(\vec{S} \cdot \vec{E}) - (\vec{S} \cdot \vec{S})(\vec{U} \cdot \vec{E})}{N(\vec{S} \cdot \vec{E}) - (\vec{U} \cdot \vec{E})(\vec{U} \cdot \vec{S})}. \tag{5.11}$$

The result from Equation (5.11) may be plugged into equation (5.9) in order to calculate A. Thus, our problem is solved.

5.3 Shifts Between Measurements

I have shown how to perform PL measurements of both polaritons, as well as reflectivity measurements, and PLE measurements of the upper polariton. My goal in this chapter has been to verify the methods used in the long-lifetime samples, where some of these measurements are not possible. Specifically, in long-lifetime samples I can not measure either polariton in reflectivity, and I can not measure the upper polariton in PL; as a result, in the long-lifetime samples I substitute them with PL measurements of the lower polariton and PLE measurements of the upper polariton instead. In this section, I compare the various measurements to gain some insight into how much error is present in the characterization scheme due to those substitutions.

A summary of the lower polariton data is shown in Figure 63. In A, I have plotted the extracted dispersion curves of the lower polariton at six different locations on the sample as measured in both reflectivity and PL measurements. The dispersion curve is extracted using the two-dimensional fitting procedure outlined in Section 5.2. Each PL curve is at the same sample location as the reflectivity curve directly below it in terms of energy. One can see that the PL measurement is consistently at higher energy than the reflectivity measurement. We expected to see a Stokes shift due to disorder, in which case the PL measurement would be at lower energy. Instead, we see the opposite. We believe that this is related to incomplete thermalization of the polaritons.

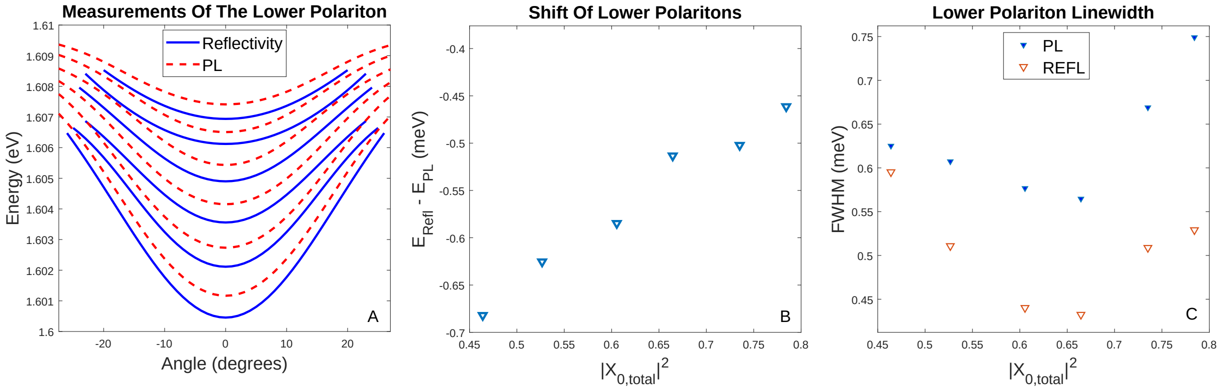

Figure 63: Image A: Curves of the lower polariton in short-lifetime samples as measured through both reflectivity and PL at six different detunings.Image B: The associated shifts between reflectivity and PL measurements at normal incidence. Image C: the linewidths of the of the polaritons at normal incidence.

Figure 63B shows the shifts from A specifically at normal incidence ($k_{|} = 0$). The horizontal axis is the total exciton fraction of the lower polaritons as measured in reflectivity measurements which I will discuss in section 5.5. One can see that as the polaritons become more excitonic, the magnitude of the discrepancy between reflectivity measurements and PL measurements decreases.

Figure 63C shows the linewidths at normal incidence from the two series of measurements. It is a common rule of thumb in the field that the lower polariton linewidth is minimized when the exciton fraction is at 50%, called the resonant position. However, the data shows that to be slightly off, and the minimum is closer to exciton fraction of 66%. I have not done a systematic study of this and do not recommend it as a means of finding the resonant position; as discussed previously, dynamic scattering of the particles can play a role in linewidth [72].

The shift between reflectivity measurements and PL measurements is a source of error. The reflectivity measurement should be viewed as the more pure measurement of the polaritons, since it does not depend on the dynamics of the particle thermalization. Since the lower polaritons may only be measured through PL measurements in the long-lifetime samples, this error should be accounted for in characterization of the long-lifetime samples, discussed in Section 5.5.

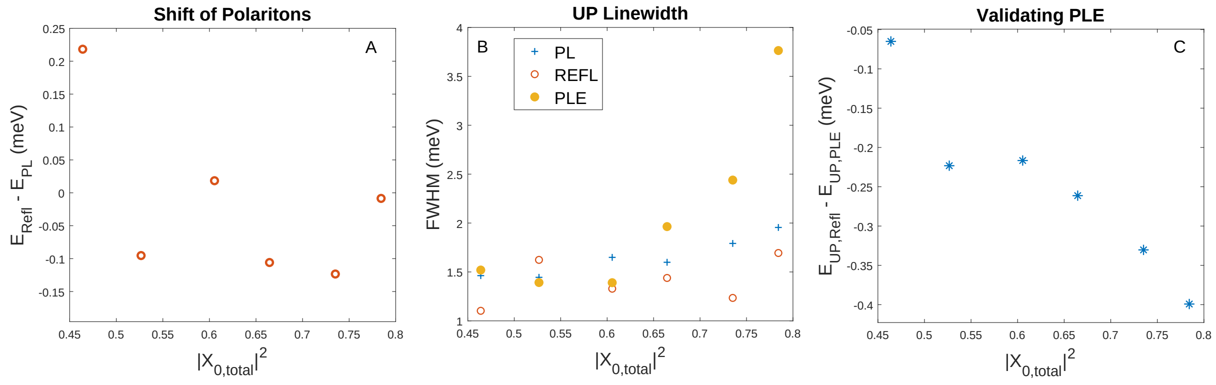

Figure 64: Image A: The discrepancy between PL and reflectivity measurements of the upper polariton at normal incidence. Image B: The linewidth of the upper polariton as measured by PL, reflectivity, and PLE. Image C: The comparison of the reflectivity measurement to the PLE measurement at normal incidence.

For the short lifetime samples, there are three measurements of the upper polariton energy: the PL, the reflectivity, and the PLE. A summary of the results is shown in Figure 64. Unlike the case of the lower polariton, the upper polariton had a very small difference between the dispersion curve as measured through PL versus reflectivity measurements. The curves lie almost directly on top of each other. The discrepancy between these two measurements at normal incidence is shown in Figure 64A. Figure 64B shows the linewidths of the upper polaritons as measured in all three measurements. Unlike the lower polariton, there does not seem to be a local minimum.

If we stayed with the Rabi model, we would expect the upper polariton linewidth to have a maximum when the lower polariton has a minimum. A simple way to see this is to think of the polariton Hamiltonian H_{pol} in terms of the exciton-cavity Hamiltonian $H_{ex,cav}$ and the transformation matrices A used to diagonalize it. Because the trace of a matrix obeys

$Tr(AB) = Tr(BA)$, we have

$$Tr(H_{pol}) = Tr(AH_{ex,cav}A^{\dagger}) = Tr(H_{ex,cav}A^{\dagger}A) = Tr(H_{ex,cav}). \tag{5.12}$$

This leads directly to

$$E_{ex} + E_{cav} = E_{lp} + E_{up}. \tag{5.13}$$

In this equation, the energies are complex. Separating the real and imaginary pieces, we are then left to conclude that the linewidths of the polaritons should add together. However, comparing the measured linewidths of the upper and lower portions to the exciton PL showed in Figure 6, we see that even if the cavity where to have a linewidth of zero (which it certainly doesn't in short-lifetime samples), it appears that the sum of the linewidths of the polaritons undershoot the linewidth of the excitons. The issue here is that the Rabi model does not account for motional narrowing. Motional narrowing is a phenomenon in quantum particles, where the linewidths of the particles is narrower than expected. Essentially, because the particle is not localized, it experiences an average of the potential it is within, which narrows its linewidth. See Reference 48 for a more thorough discussion within the context of polaritons. Furthermore the bulk of the linewidth of the exciton comes from disorder in the form of inhomogeneous broadening. The two-level Rabi model can only account for homogeneous disorder.

Figure 64C shows the discrepancy between PLE and reflectivity measurements of the upper polariton. This is a crucial piece of information as it will play a role in determining how much error there is in the long-lifetime characterization scheme in Section 4.6. One can see there is certainly a trend in the data, and PLE seems to be more accurate near the resonant position on the sample, as opposed to the very excitonic locations. The specifics of how to handle this are discussed in Section 5.5.

5.4 TMM Fitting Of Short-lifetime Samples

I will now discuss the process of fitting the transfer-matrix method (TMM) simulation to the measurements in short-lifetime samples. There are two types of fits to perform. In both types, first the simulation is fit to the broad reflectivity data, exactly as demonstrated in Section 4.5.1. Then, in the first type of fit, the simulation is fit to the PL measurement of the lower polariton, along with the PLE measurement of the upper polariton. This process is exactly the same as I have already demonstrated for long-lifetime samples in Section 4.5.2, and so I do not include it here, only the its results.

The second type of fit is only available in short-lifetime samples, and will be demonstrated in this section. After the broad reflectivity curve has been fit, the simulation is then fit to the dispersion curves of both the upper and lower polariton as observed through reflectivity measurements. One crucial difference here is that the single PLE data point is now replaced with an entire upper polariton dispersion curve. This means the curvature of the upper polariton in the simulation can be tuned into agreement with the measured curvature. Because the reflectivity measurement should be viewed as the more pure measurement of the upper polariton, this will enable us to estimate the errors in the characterization of the long-lifetime samples.

First, I fit the TMM simulation to the broad reflectivity data from Figure 61A data using the methods of section 4.5.1. All the tuning parameters are the same, the optimizer is allowed to slightly tune the thicknesses of the DBR layers as well as the index curves. The results of this fitting are shown in Figure 65.

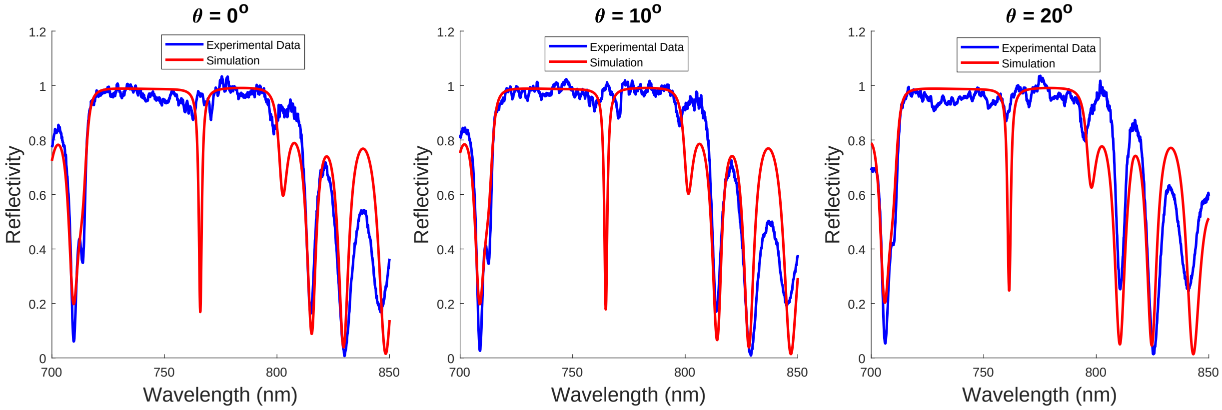

Figure 65: Fits of the TMM simulation to the experimental

As in the long-lifetime samples, the fits are good, although not perfect. Again, the region right around 805 nm is notably off, likely due to PL of the bulk excitons of GaAs not being incorporated into the model. In the center of the stopband is a large dip. This is the cavity mode. Unlike in the long-lifetime samples, this cavity mode is readily apparent without needing to zoom in or increase the mesh of the simulation. Once again, I include the excitons inside of the quantum wells using a simple charged-oscillator model [73], given as

$$\chi_{ex}(E) = A\frac{(E_{ex}^2 - E^2) + i\Gamma E}{(E_{ex}^2 - E^2)^2 + \Gamma^2 E^2}. \tag{5.14}$$

As previously mentioned, the results in short-lifetime samples ultimately lead me to include two excitons; the heavy-hole and light-hole excitons. For the purpose of demonstration, we can leave the light-hole exciton out of the simulation for now. The equation for the complex electronic susceptibility of the quantum wells is

$$\chi_{QW}(E) = \chi_{GaAs} + \chi_{ex,hh}. \tag{5.15}$$

This results in the cavity mode splitting into the upper and lower polaritons. From the simulation I can extract the dispersion curves $\lambda_{lp,sim}(theta)$ and $\lambda_{up,sim}(\theta)$. In order to fit the simulation to the data, I need to define a metric to minimize. Just as in the long-lifetime samples, first there is the error associated with the lower polariton:

$$e_{lp} = \frac{1}{n_\theta}\sum_{\theta} |\lambda(\theta)_{lp,sim} - \lambda(\theta)_{lp,refl}|, \tag{5.16}$$

where n_θ is the number of terms within the summation; this time we use the lower polariton dispersion measured through reflectivity $\lambda_{lp,refl}$. Additionally, I am now fitting a full reflectivity dispersion curve for the upper polariton $\lambda_{up,refl}$, so the error of the upper polariton is

$$e_{up} = \frac{1}{n_\theta}\sum_{\theta} |\lambda(\theta)_{up,sim} - \lambda(\theta)_{up,refl}|. \tag{5.17}$$

The total error is simply the sum of these two: $e_{total} = e_{up} + e_{lp}$. In this case I do not need to introduce a relative scaling parameter to add the two types of error together. I can now fit the TMM simulation to the measured data. The fitting parameters are the heavy-hole

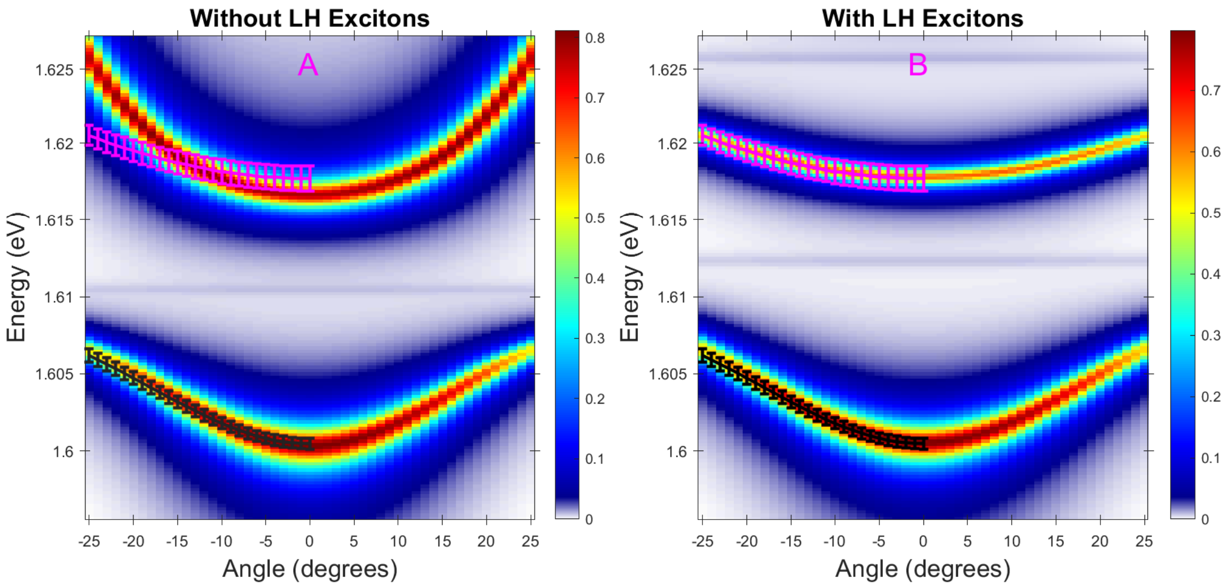

Figure 66: In the background, the inverted power Reflectivity $1 - R$ from the TMM simulation. The experimental data is plotted on top, with the error bars representing the FWHM of the data. In Image A the light-hole exciton is not included in the complex electronic susceptibility of the quantum wells. In Image B the light-hole excitons are included.

exciton energy and associated amplitude A_{hh} and a scale factor for the non-quantum well layers inside the cavity. The results of this fitting are shown in Figure 66A.

One can see that upper polariton in the simulation does not curve downwards in the same manner as the experimental data. I find this consistently be the case when I perform this fit at various locations on the sample without including the light-hole excitons in the model. I tried multiple alternative approaches to fix this discrepancy, but all of them were unsatisfactory, as they required unrealistic numbers to work. I tried convolving Equation the complex electronic susceptibility of the heavy-hole exciton oscillator with a Gaussian distribution in order to account for inhomogeneous disorder (i.e. a Brendel-Bormann oscillator). I also tried creating a finite series of exciton states to account for the 12 slightly different quantum wells in the sample. I also tried making E_{ex} dependent on angle in order to account for the non-infinite mass of the excitons. None of these approaches have been as broadly successful as the light-hole exciton.

Now that the TMM simulation has been fit to the data, I can calculate exciton fractions just as I did for the long-lifetime samples using the numerical derivative,

$$|X_{k_\parallel}|^2 = \frac{d}{dE_x} E_{lp}(k_\parallel) \approx \frac{E_{lp_+}(k_\parallel) - E_{lp_-}(k_\parallel)}{2\Delta E_{ex}}. \tag{5.18}$$

Specifically I am interested at the result at normal incidence. The results of this are discussed in Section 5.5.

5.5 The Errors Of The Characterization Curve

I have now shown how to extract the exciton fraction from the short-lifetime samples two different ways. The first involves fitting the TMM simulation to the lower polariton PL data and the upper polariton PLE data. The second is fitting the TMM simulation to the polaritons as measured in reflectivity. As I have shown, there is a notable shift between the PLE measurement and the reflectivity measurement of the upper polariton. Additionally, there is a shift between the PL measurement and the reflectivity measurement of the lower polariton. Here, I will show what those shifts ultimately lead to.

Once again I construct the characterization curve, allowing the exciton fraction to be obtained by just a measurement of the lower polariton's energy at normal incidence. This time I have two different curves to construct, namely the curve from fitting the PL and PLE measurements, and the curve from fitting the reflectivity measurements. The results are shown in Figure 67. The horizontal axis is energy, but it has different meaning for the two sets of data. For the PL+PLE data set, the horizontal axis is the energy of the lower polariton as measured through PL. For the reflectivity data set, the horizontal axis is the energy of the lower polariton as measured in reflectivity. One can see that plotted together like this, both sets of data lie on nearly the same exact curve.

This is a truly shocking result. I expected a result similar to Figure 54, in which the samples lie on different curves which can intersect; or perhaps Figure 55 where a change in

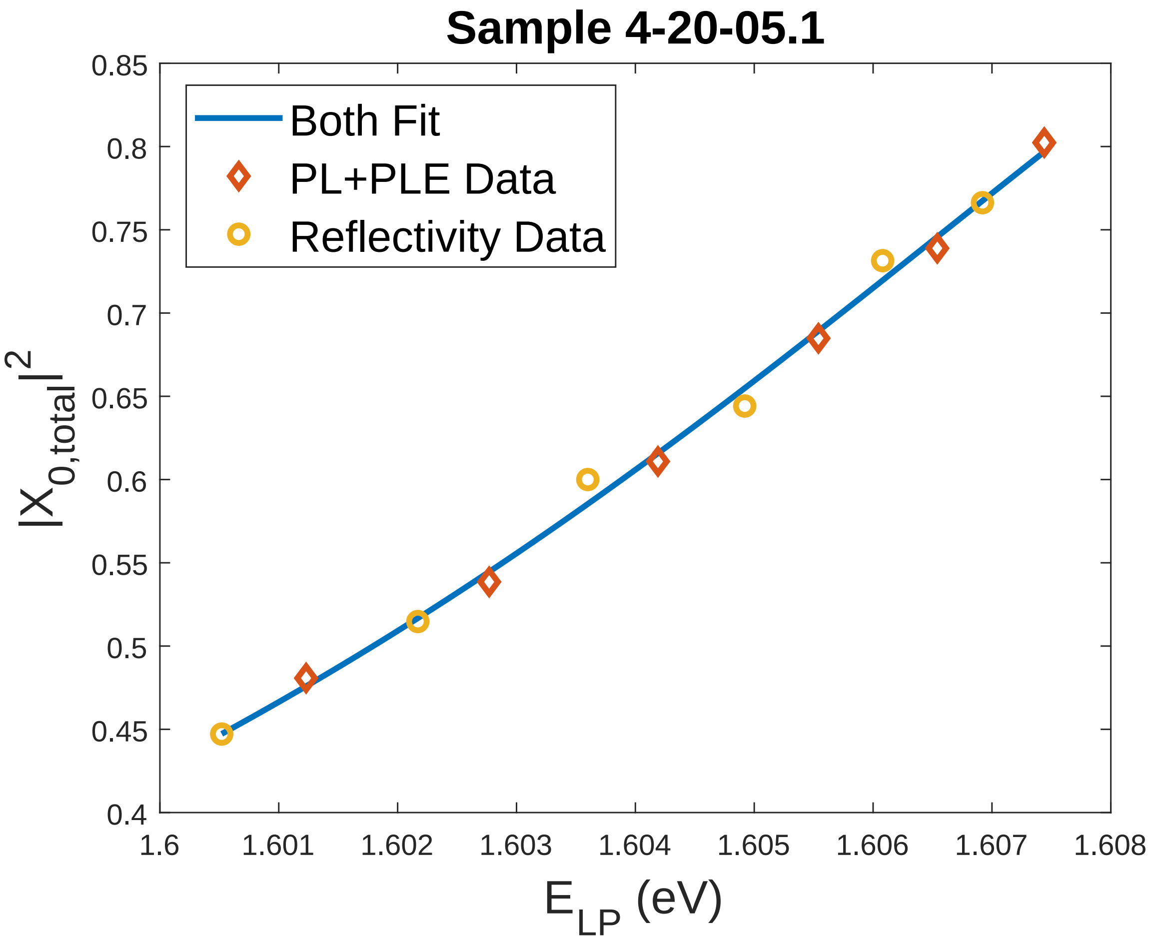

Figure 67: The characterization curve for the short-lifetime sample as measured through reflectivity measurements versus PL and PLE measurements. A simultaneous fit of both sets of data.

the temperature of the sample shifts the curve up and to the left. Instead, I find that the apparently the two shifts described above seem to offset one another and lead to translation along the characterization curve.

It would be really nice if we had a thicker sample with a detuning range extending further into the photonic regime; unfortunately we do not. Another student in the lab has plans to obtain some and hopefully they will perform this characterization to extend my result.

Although this seems to be a fortuitous coincidence, it may be utilized to help determine the error bars on the characterization curves. The main source of error is going to be the difference between the PL measurement and the reflectivity measurement for the lower polariton. Essentially, this means the main source of error causes translation along the horizontal axis in Figure 67.

The largest discrepancy seen at any one location for the total exciton fraction as measured using either of our two methods is $\Delta|X_{0,total}|^2 = .04$, while the average value was $\Delta|X_{0,total}|^2 = .025$. Additionally, when account for possible errors in the calibration of the equipment, another $\Delta|X_{0,total}|^2 = .02$ is reasonable, so ultimately I settled on the uncertainty of the total exciton fraction to be $\Delta|X_{0,total}|^2 = .05$, which I used for the error bars in Figure 53A.

So then, the main task is complete. In Chapter 4 I demonstrated how to characterize the long-lifetime samples and in this chapter I've now shown where the error bars came from.

I tried repeating these results in Sample 10-5-10.1 which is a short-lifetime sample with only three quantum wells instead of 12. However, I found it that the splitting between the upper and lower polariton was too small to work with. Essentially, the lower polariton would show up as a shoulder in the PL of the lower polariton. Example data of this is shown in Figure 68. This is a problem for my TMM simulation, as it could not reliably extract and differentiate one branch from the other. A human fitting the curves manually is able to do so. But the optimization required for fitting the simulation to the data requires a fit to be performed hundreds to thousands of times.

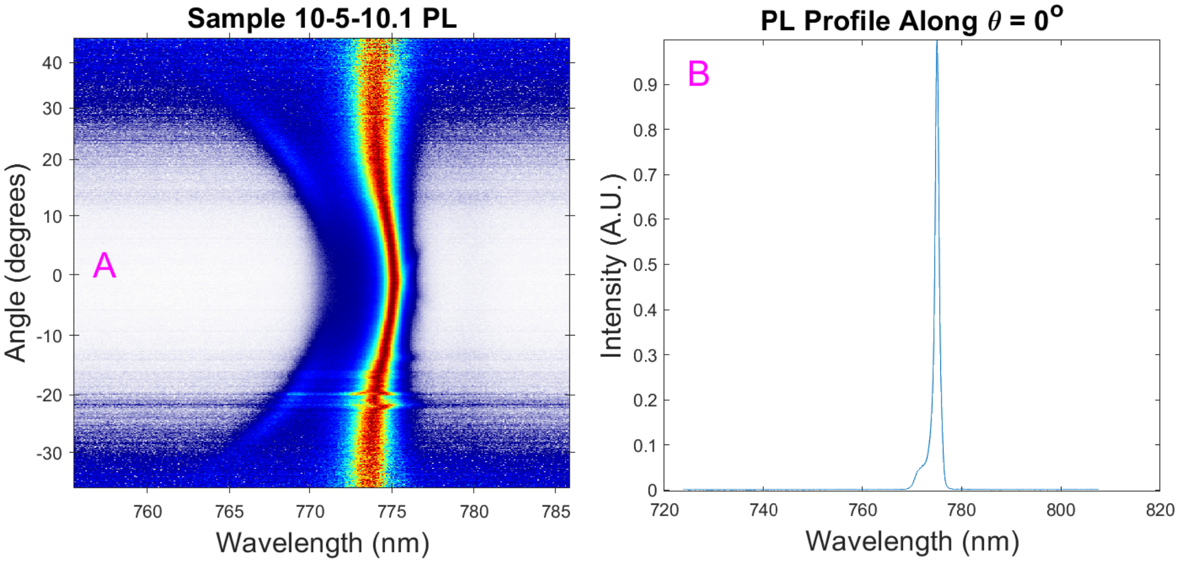

Figure 68: Image A: The PL of sample 10-5-10.1 after using the equalization algorithm to make it more visible. Image B: the profile across $\theta = 0$.

6.0 TMD Work

As discussed previously, exciton-polaritons formed out of GaAs quantum well excitons are mostly studied at cryogenic temperatures due to their low binding energy of about 8 meV [10, 14]. However, excitons may be formed in other materials with a higher binding energy in the pursuit of room temperature polariton systems. One such family of alternative materials which I've worked with is monolayer flakes of transition metal dichalcogenides (TMDs). In TMDs, binding energies of several hundred meV are commonly observed [44]. At room temperature, $k_B T$ is about 26 meV. Thus, we would expect TMD monolayer excitons to survive quite well at room temperatures, and indeed we've found this to be the case.

However, there are also additional difficulties in TMD systems. Specifically, the creation of samples is more difficult because it is always a multistage process. By contrast, our GaAs samples are made in one continuous MBE growth. As a result, more noticeable impurities show up in the PL of TMD samples. Additionally, TMD systems boast very impressive light absorption efficiency. However, they are still just one atomic layer thick. Thus, while TMD samples may be able to produce more PL per unit thickness, GaAs samples typically have much brighter total PL.

In this chapter I briefly outline work towards room temperature TMD monolayer exciton-polaritons, as well as some side projects worked on along the way.

6.1 Sample Design

There are two main options available to create single atomic layers of TMDs: exfoliation and large-area growth through chemical vapor deposition (CVD). Wide-area growths were produced by our collaborators; the Banerjee group at the University of Texas at Austin [36]. The exfoliation technique can be performed in our lab, and in many other labs, using commonly available equipment. The basic premise of the exfoliation technique is to use tape to peel thin layers of TMD off of bulk crystals. Bulk crystals of various TMDs may be purchased commercially. Some of these materials of particular interest to us are MoS_2, $MoSe_2$, and WSe_2; which may be purchased in undoped, p-type, and n-type variants. These two approaches have both pros and cons. In this section, I discuss the wide-area growths.

The wide-area growths tend to produce large amounts of usable sample. An example of a wide-area sample is shown in Figure 69. The CVD growths tend to not be completely uniform across the sample, so there tends to be a gradient where some regions have relatively few monolayer flakes, some have many flakes, and then eventually some of the sample has bulk material growth. The monolayer flakes are easily distinguishable by their triangle- shaped domains. The triangle domain pictured is approximately 50 micrometers across. There is a region several hundred micrometers across where such flakes are common.

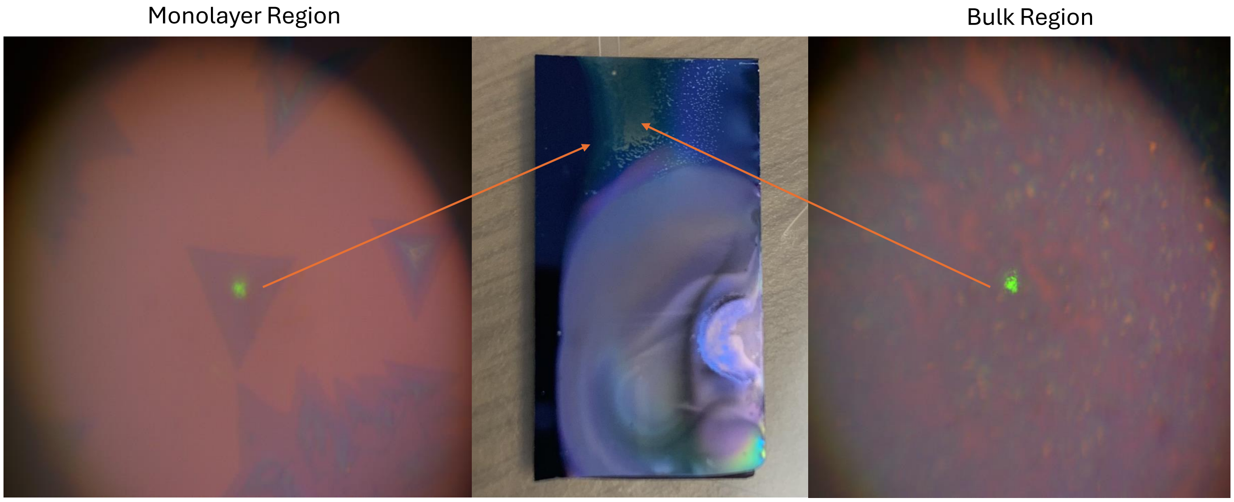

Figure 69: A wide-area TMD sample grown by the Banerjee group. The green spot is laser excitation.

The wide-area sample shown was grown on a Si substrate. However, the ultimate goal is growth on top of a bottom DBR. The Banerjee group felt they could grown effectively on SiO_2. DBRs can be grown effectively via CVD using SiO_2 as the low-index material, and Si_3N_4 as the high-index material. However, they could not be grown in the same CVD machine, and DBR growth is not a specialty of the Banerjee group. Therefore, the DBRs were grown by our collaborators at the Forrest group at University of Michigan. Typically the design was eight periods. The topmost layer of the DBR was SiO_2, which also will act as the spacer for the microcavity. These DBRs could be grown on an opaque Si substrate or a transparent quartz substrate.

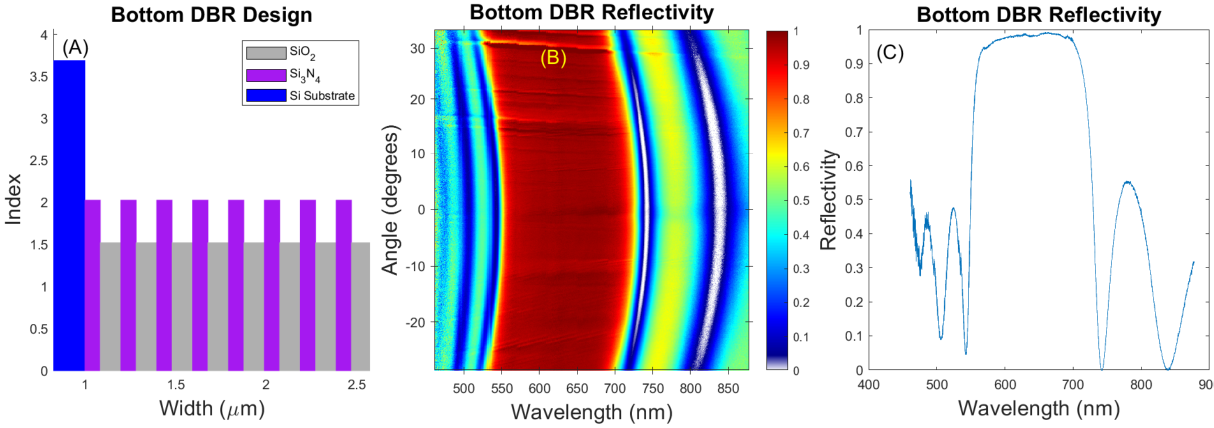

Figure 70: A: Design of the bottom DBR in our TMD samples. B: Angle resolved reflectivity measurement of this sample (grown by the Forrest group). C: The profile through image B at $\theta = 0$.

The design of the bottom DBR is shown in Figure 70A. Reflectivity data of the bottom DBR grown by the Forrest group is shown in Images B and C. We see that despite being only eight periods, the reflectivity looks very nice, and the stopband quite flat. After optical characterization in our lab, this sample was shipped to the Banerjee group for MoS_2 growth, and then shipped back to us.

Unfortunately, there was a noticeable degradation in quality in MoS_2 grown on these DBRs, compared to MoS_2 grown on Si substrates. An example of this data is shown in Figure 71A. We see that the exciton lines for the MoS_2 monolayers grown directly on the bottom DBR are much broader than on pure Si. For the sample grown on Si substrate, the lower-energy peak is the A exciton and the higher energy peak is the B exciton (a splitting of the excitons states in two levels, A and B, is typical in TMD monolayers [28, 44]). Overall, it isn't clear that the growth on SiO_2 is working very well. It is possible that multiple layers are forming, rather than monolayers. However, in Fig. 71B, one can see the reflectivity

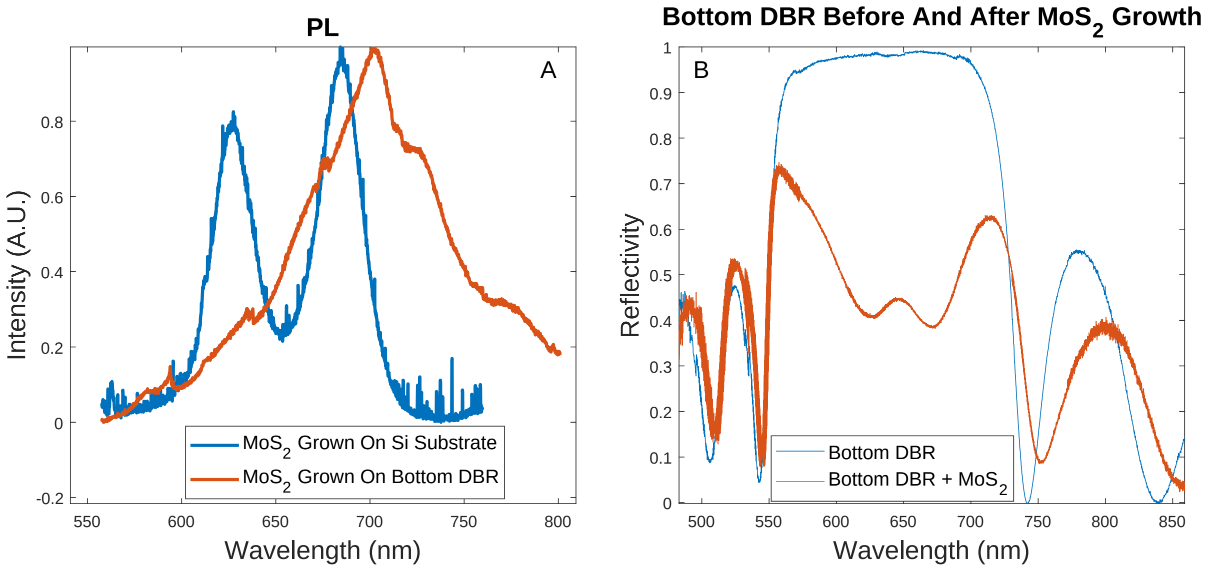

Figure 71: A: PL of MoS_2 grown on pure Si substrate. B: PL of MoS_2 grown on the top SiO_2 layer of the bottom DBR.

measurement of the bottom DBR before and after MoS_2 growth. Two dips at wavelengths very similar to the two peaks in Figure 71A are seen. Even though the PL is discouraging, the reflectivity is encouraging.

I tried putting a top DBR onto this and similar structures a few different ways. One method was simply laying a top DBR grown on a transparent substrate on top, and imaging the structure through the transparent substrate. However, this resulted in a lot of interference fringes due to air gaps. After this, I constructed a simple clamp system with springs and rubber pads to firmly squeeze the top DBR down onto the bottom DBR plus MoS_2 sample. This helped reduce the interference fringes. However, polariton modes were not visible in either the reflectivity nor the PL measurements.

My TMM simulator was very helpful in designing this top DBR. Overall, I've found very good agreement between the simulated sample and the measured reflectivity using commonly available references for the optical indexes of the materials. The index of Si_3N_4, SiO_2, Si substrates, and glass substrates were taken from References 52, 20, 52, and the Bk7 page of Reference 66 respectively. The process occurs in steps. First, a bottom DBR is measured and fit. Then, those fitted values are assumed to apply to the top DBR to be added to it. The simulation then allows us to fine tune the top DBR to hopefully put the cavity into resonance with the excitons.

Examples of the fitted bottom DBR reflectivity are shown in figure 72. In image A, the experimental data is shown with the simulation using the (mostly) unfitted simulation. The only parameter which as been modified for the simulation is the total design thickness, which is just accounting for a total systematic undergrowth due to location within the CVD machine. We see that even without fitting, the width of the stopband is already quite close to the experimental data. In Image B I have done the standard fitting procedure described in section 4.5.1. This process was also done at multiple angles, same as I did in the GaAs samples.

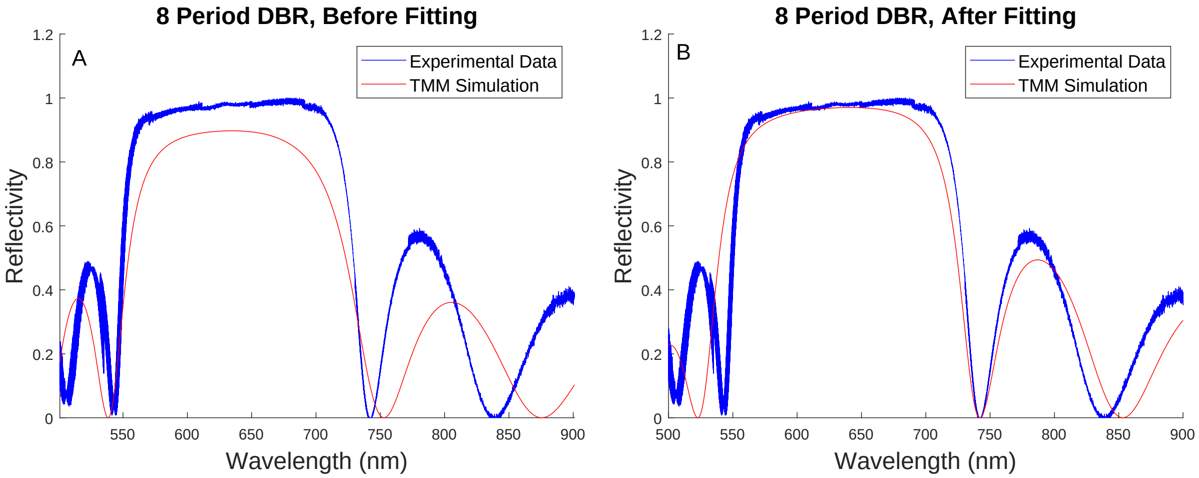

Figure 72: Reflectivity of 8 period SiO_2/Si_3N_4 DBR.

In samples with more periods, for example 16, I find better agreement between the simulation and the data. I believe this is because proper treatment of the substrate is tricky. In samples with more DBR layers, the substrate plays less of a role.

Once the bottom DBR is fit, I then included TMD monolayer, the top spacer, and the top DBR in the simulation. I assumed the scalings in the bottom DBR would also apply to the top DBR (as they are grown in the same machine). The main problem here is how to create a cavity mode resonant with the monolayer (flake or wide area). One method I tried, was after the above measuring and fitting of the bottom DBR, use the fitted simulation to determine the necessary thickness of the top spacer layer. One variation of this was to use spin-coated polymethyl methacrylate (PMMA) as the top spacer layer. This was chosen because it has an index similar to SiO_2. Additionally, we thought that perhaps by squeezing the sample this PMMA layer would compress a little bit, giving us tunability of the cavity mode. PMMA is generally more squishy than SiO_2.

The excitons inside the TMD were approximated as a simple oscillator just as I did in section 4.5.1. However, there is now a key difference. In GaAs we had working polariton samples, and we wanted to fit a model to it. Here, we are using a model to predict what samples will work. So, it is necessary to have some accurate initial guess of the index of our monolayer. Reference 89 gave us such a starting point. Thus, the thickness of the PMMA layer could be tested in the simulation before we spin-coated it onto the real sample. Figure 73 shows this process. The thickness of the PMMA layer can be varied to make the lower polaritons more or less excitonic. The time-averaged magnitude of the E field for an incident beam at wavelength equal to the lower polariton's wavelength is plotted on top of the sample in Image A. The simulated reflectivity curve is plotted in Image B. We see that unlike in GaAs, the splitting here is much less. This can be corrected by utilizing multiple monolayers stacked on top of one another.

Ultimately, I was not able to produce TMD monolayer exciton-polariton samples. However, other students in our lab and our collaborators eventually did succeed in producing such samples using very similar methods [30]. Individual flakes of WS_2 were exfoliated and stacked on top of each other on a bottom DBR. Three such flakes were used in order to increase the coupling strength, with hexagonal boron nitride (hBN) separating them. More flakes results in more coupling and brighter signals, and the hBN separates the flakes so that they remain monolayer rather than multiplayer. Additionally, hBN encapsulation generally increases the intensity of the PL from these monolayer samples. In total there are 4 layers of hBN and 3 layers of WS_2. A PMMA spacer layer on top is used to fine tune the cavity

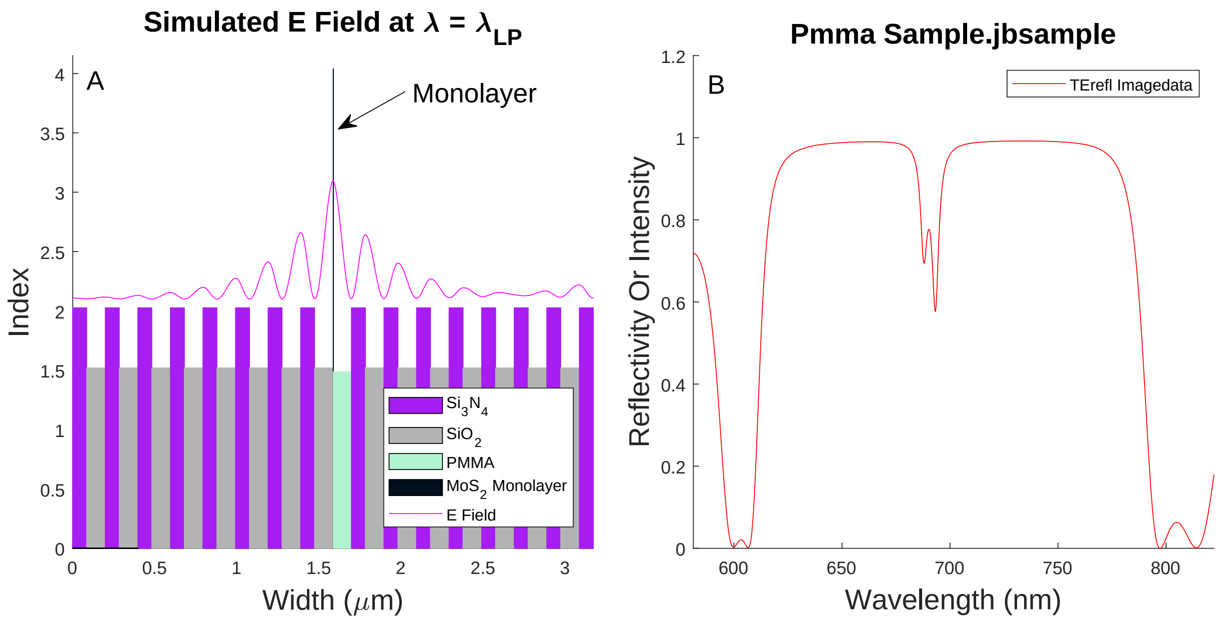

Figure 73: Image A: The design of our asymmetric TMD sample. PMMA is the top spacer layer of the cavity. Image B: Simulated reflectivity at normal incidence.

into resonance with the excitons. Finally, the top DBR was grown directly on top of the PMMA layer via CVD. The PMMA layer also serves to protect the monolayers from being destroyed by the CVD growth.

Unfortunately, a wide area tunable TMD exciton-polariton sample still has still not been produced.

6.2 Methods Of Tunability

One of the advantages of the wide-area TMD samples is that they could be grown on a bottom DBR with a thickness wedge across the sample. This would then allow for varying the detuning by moving along the sample, just as we did in GaAs samples in previous chapters. However, the PL from monolayer samples that were transferred onto the bottom DBR through exfoliation tended to produce much more pristine PL.

By contrast, exfoliation only results in a single usable flake, so bringing a cavity into resonance with the excitons in that flake poses a problem. In order to prepare for that possibility, Zheng Sun (a former postdoc in our lab) and I investigated other methods of putting the two into resonance.

One method we explored was placing electrical contacts on the sides of a WSe_2 flake and varying the voltage applied to it [82]. We found that applying a voltage of 40 V would result in a shift of the central energy of the excitons of 20 meV. The voltage could also be negative to shift 20 meV in the other direction, giving a total tunable range of 40 meV. We also saw hysteresis in the PL as we cycled the voltage. However, the hysteresis would not affect this method's viability for putting the exciton and cavity into resonance.

Another method we tried was stressing the sample [80]. I constructed a sample stage which allowed application of stress to the back side of the sample using a needle. The needle was pushed forward with a spring attached to a micrometer, which I calibrated as discussed in section 2.7. Because the needle pushes the sample from behind, this has the effect of stretching the TMD. Similar methods have been used by our lab in the past [8, 60]. Overall, we were able to shift the central energy of the sample approximately 9 meV.

6.3 Stacked Multilayer Structures

Zheng and I also explored multilayered structures, in which we used exfoliation to stack multiple monolayer samples on top of one another. We did this to construct interlayer excitons [79], which generally have a longer lifetime than the standard intralayer excitons. Additionally, we created a four-particle complex named the quaternion [81]. The quaternion is made from one positive or negative charge, surrounded by three negative or positive charges respectively (one of these surrounding charges is in the other layer, similar to an interlayer exciton). As such, quaternions are charged bosons. That makes them interesting as they could provide a path towards superconductivity without Cooper pairing.

For these experiments, we wanted to perform various measurements on our samples which involved data collection for long periods of time. However, stacked samples tend to have very small regions where all the layers are properly stacked. Thus, sample positioning is very important. To facilitate these measurements, I designed and built the dual imaging setup and active stabilization described in Section 2.6. That system effectively uses a second camera to monitor the samples position and automatically moves it back into position if some random vibration knocked it out of position.

7.0 Closing Remarks

A core goal of my work has been to finally solve the characterization problem in long-lifetime samples. Other students who have passed through our lab have tried their own approach at this problem, but each of those approaches had problems. The exciton fraction is a critical parameter for various calculations, yet the problem of measuring it accurately for our long-lifetime samples has remained a problem for over a decade. However, the work presented here is a thorough and successful solution to that problem.

This work required the development of angle-resolved reflectivity and PLE as practical tools. This involved automation of many pieces of equipment in the lab, and a central software GUI to coordinate and control them. This GUI eventually grew to cover nearly every piece of hardware in the lab, and has been broadly useful to every student in the lab.

At the outset, PLE was an unreliable measurement. However, the turning point was the development of the dynamic background subtraction technique. Utilizing this technique, PLE is now a reliable measurement at all detunings on the samples. Similarly, angle resolved reflectivity required the development of proper alignment procedures, as well as a more thorough calibration of several pieces of lab equipment.

My TMM simulation utilizes this data, along with the data from standard PL measurements to create a good working model of the polaritons. This working model may be used to effectively design future samples. I've tested the fitting procedure in short-lifetime samples and found that the inclusion of the light-hole exciton is necessary in order to fit the curvature of the upper polariton, which is now present in the fits of all the long-lifetime samples. With this addition, there is complete consistency of the same model fitting the two different types of samples. I've provided a table with the corresponding fits of over a dozen different samples we've shared with our collaborators over the years.

Furthermore, I've helped our collaborators at Waterloo begin to grow long-lifetime samples of their own; Mirroring the samples grown by our collaborators at Princeton. There are interesting new projects available utilizing these samples, as they are very flat.

Additionally, there are some interesting ideas that could be explored in future work. For example, a more sophisticated approach to the statistics and fitting presented in section 4.4.1 could be interesting and more convincing. Additionally, searching for some sort of relationship between the fitted parameters from section 4.5.2 and the thickness of the sample could be interesting.

Lastly, all of my techniques have been implemented in various GUIs for other students to utilize, as well as posted online with videos demonstrating how to use them. I hope that this work is useful to others.

Appendix A Software Manual For Main GUI

This appendix shall serve as a software manual and documentation for "Main GUI". The goal being to help others use the software, but also to help them expand upon it as new equipment and situations call for. Main GUI is the central graphical user interface (GUI) which serves to control all the various instruments in our lab. It is capable of performing various sweeps and algorithms while collecting and storing data. Essentially, every experiment currently ran in the lab is ran through Main GUI. For example, the PL, PLE, and reflectivity measurements described in this work as well as many others.

In addition to Main Gui, this appendix will also discuss "Quick Movie GUI", "Data Analyzer", and "Spec Fit". Which are separate software packages for displaying data, processing data, and fitting data respectively. The "Device Initialization GUI" will also be discussed, which is a secondary GUI that Main GUI utilizes when the user is first connecting to various instruments. A separate set of GUIs and software packages pertaining to transfer matrix method (TMM) simulation of our samples are discussed in appendix B.

A.1 Overview of Devices

When first beginning an experiment nearly all the buttons in Main Gui are grayed out to indicate they are disabled. They will be enabled once the user has initialized communications with devices which utilize the various buttons. The Initialize Devices button is available upon initial loading of Main Gui. Clicking on it will bring up the Device Initialization Gui, shown in figure 74.

The Device Initialization GUI will guide the user through the steps necessary to initialize communications with the various instruments through the use of popup text boxes. This GUI is unfortunately visually cluttered, and if I were rewriting this GUI from scratch I would recommend utilizing something such as a scrollable table in which the various devices are selected for initialization. When I first made this GUI, the devices of interest were

ready

Primary Camera: roper

Secondary Camera: none

Spectrapro Spectrometer: COM 3; Grating 1 300 g/mm ...; Mirror Front (Streak ...

Oriel Spectrometer: COM 3; Mirror Front (C port -...

Active laser: none

Transmission Power Meter: Model: pm100; COM 19; Wavelength free (wont ch...; Counts 1

Reflection Power Meter: Model: pm100; COM 9; Wavelength free (wont ch...; Counts 1

Beam Splitter calibration: none

Filterwheel calibration: none

○ synchronize power meters

Thermometer: Model: Thermometers; COM 8

fw102c Filter: COM 18

All apt devices: Serial Number ☐ V micrometer; Serial Number ☐ H micrometer; Serial Number ☐ f micrometer; Serial Number ☐ pump filterwheel; Serial Number ☐ image1 polwheel

Oriel whitelight filter: COM 20; Mirror Front (C port -...

Figure 74: The Device Initialization Gui. This GUI initializes communications with the various devices and has some control over basic settings.

two cameras, a spectrometer, and a laser, but it has feature crept over time. The buttons corresponding to instruments which have already been initialized are grayed out.

Below I provide an overview of the buttons, their corresponding devices, and the functions available in the respective software libraries.

Initialize Roper Camera Initializes communication with either the Roper Cascade 512B CCD camera or the Roper 1k CCD Camera. Communications are handled via a proprietary PCI card. Originally communications were handled via MEX files to directly call the Static library drivers for the camera. This became unstable when we transitioned to Windows 10 and a 64 bit computer. Communications are now handled by calling the Java libraries available from Micro Manager [29]. Controls implemented in Main GUI: set integration time, set gain, get image, open and close shutter, start and stop live imaging mode

Initialize Streak Camera Initialize communication with the Hamamatsu C6878 Delay Box, C9100 CCD camera, and M5676 SynchroScan Unit. Communications with the C6878 and M5676 are handled via Texas Instruments GPIB card, C9100 communications are through a proprietary PCI card. MEX files written in C are utilized for the GPIB communications. C9100 communications are handled directly in Matlab. Controls implemented in Main Gui: set c9100 integration time, start and stop live imaging mode, set c9100 gain, set M5676 MCP gain, set sweep mode, set photo mode, set lock mode, get image. Time axis calibrations are automatically displayed in Main Gui based on the sweep mode and photo mode.

Initialize Davis Camera Initialize communication with the LaVision DaVis camera. Communication is handled via a fiber optic cable connecting to a proprietary PCI card. Communications are handled by calling the Sen Cam .dll library. Controls implemented in Main Gui: set integration time, set gain, get image, start and stop live imaging mode.

Initialize Thorlabs Camera Despite its name, this button can initialize communications with any camera using the uc480 drivers. This includes the Thorlabs DCC camera as well as the Zeiss Axiocam 105 color. Communications are handled via direct calls from MATLAB to the uc480DotNet.dll library, over a USB cable. USB 3.0 specifically is required for the Zeiss camera. Zeiss camera is shipped with a generic USB 3.0 board which is installed in the various computers in our lab. Controls implemented in Main Gui:set integration time, set gain, get image, start and stop live imaging mode. These cameras may also be ran as secondary cameras, which are opened in a second instance of MATLAB with their own GUI for dual imaging setups. The second instance of MATLAB is controlled by the first instance of MATLAB via network communications over the IP 0.0.0.0 port 0.

Initialize Logitech 120 Camera Initializes communication with logitech 120 usb camera. Might work with a broader family of usb cameras. This was an alternative camera used for dual imaging and active stabilization, but is not used much today. Calls the built in MATLAB usb cameras library. Controls implemented in Main GUI: set integration time, set gain, get image, start and stop live imaging mode.

Primary Camera (dropdown menu) Select which camera is currently being displayed in Main GUI. Common usage is swapping between streak camera and roper camera in bay 1. Updates automatically when certain cameras are initialized.

Secondary Camera (dropdown menu) Select a secondary camera, utilized in dual-imaging setups and active stabilization. Typically this is the Thorlabs camera, Zeiss camera, or logitech usb camera. Controlled in a second GUI, often in a separate instance of MATLAB.

Initialize Spectrapro Spectrometer Initializes communication with either the SpectraPro 2300i or 2500i imaging spectrometers. Communication is handled through an RS-232 cable. The COM port can be entered in the COM box to the right. Typically we set the COM port to be 3 in the Windows device manager window. The Grating drop down box is used to select between currently installed gratings, and the mirror drop down selects between the exit mirrors, switching between roper camera and streak camera in bay 1 currently. When new gratings are installed they must first be installed using the SpectraPro Monochromator Control software. This software gives approximate calibration so that the entered wavelength will roughly be centered on the CCD camera. Further calibration of the wavelength axis is discussed in section A.2.1. Implemented Main GUI controls: go to wavelength.

Initialize Oriel Spectrometer Initializes communication with the Oriel MS-257 spectrometer. Communication is handled through an RS-232 cable. This spectrometer is not currently utilized as an imaging spectrometer, however it can be. Only one grating is available, and it would be prohibitively difficult to change. The exit mirror drop down box allows exit mirror selection. Implemented main gui controls: go to wavelength.

Initialize Msquared Laser Initializes communication with the Msquared solstis laser. Communication is handled through sending JSON strings back and forth to the Msquared ICE-BLOC (control server). In order to facilitate this, the ICE-BLOC must have a static IP address which the various computers in the lab can communicate with. This static IP is hard coded into the solstis MATLAB library. This is handled by placing the ICE-BLOC behind its own router with a static IP address. The router has port forwarding turned on ports 39933, 8088 and 80. All traffic on those ports is forwarded to IP 192.168.1.3 (the ICE-BLOC). A printout of all the port forwarding is saved in the solstis MATLAB library. The appropriate IP address of the laser is selected based on which computer is attempting communication, as determined by the determine-computer-function. This function works by examining the computer's static IP address. If anything concerning the network is changed, this function will need to be updated as well as the solstis-find-instrument function. Implemented Main GUI controls: get wavelength, go to wavelength.

Initialize Manual Laser This opens the Manual Laser GUI (MLG). The MLG is a convenient way to manually record laser information into Main GUI. Main GUI sees it as another electronically controlled laser. The numbers in its text box will be stored in the data files the same was as all the other electronically controlled lasers, so that data processing algorithms function the same. For manual laser sweeps (early versions of PLE measurements), the MLG will query the user for the current laser wavelength after every data call. If the laser is set to a single wavelength the query options may be disabled. Responses to the queries will automatically be displayed in the MLG and stored inside the data files saved by Main GUI.

Initialize Verdi-G Laser This initializes communication with the Coherent Verdi-G laser. This laser was originally used to pump the homemade laser (built by Gangqiang Liu). Currently it is used as a pump laser in bay four, typically for TMD experiments. It operates at a singe wavelength, 532 nm. However, its power may be tuned. Obtaining programmatic control of this laser in MATLAB was difficult. However, the laser may be controlled easily using the Coherent OPSL GUI. A MATLAB solution was created which works OK for performing power sweeps. The solution utilizes the JAVA Robot class to execute certain key combinations. Those combinations are then detected by an executable made in the AutoHotKey environment. In this was, power settings from the Microsoft clipboard may be copy and pasted back and forth from MATLAB to the OPSL GUI. Although a bit odd, this solution works well enough for power series sweeps. Caution must be utlized with this laser, as it is a very powerful green laser, and therefore very dangerous to human eyes. Implemented Main GUI controls: get power, set power.

Initialize Opera HP Laser Initializes communication with the Monaco laser. This communication is handled through text strings sent over the network. The strings are sent from the Matlab instance running on the user's computer to the Matlab instance running on the Monaco's dedicated laptop computer. In the code, the user's computer is referred to as the client, and the Monaco's laptop is referred to as the server. This communication is done on ports 32000 through 32005, which are all set to forward to IP 192.168.1.5 in the router that the Monaco sits behind. Implemented Main GUI controls: go to wavelength, open and close shutter (right click laser wavelength button).

Initialize Transmission Power Meter (Reference Power) Initializes communication with one power meter. The Transmission name is just a name in the code, it doesn't necessarily have to measure a transmitted beam. The lab currently has two models of Thorlabs power meters. The PM100 and the PM100D. The PM100 are the older analog models with gray boxes. The plastic housing of these boxes have begun to degrade and have been repaired with super glue. The PM100D are the newer digital models with red boxes. Each model has multiple power heads which are compatible with it, however our MATLAB code automatically detects and handles differences between the heads. The model of the power meter must be selected from the dropdown box. The PM100 communicates using an RS-232 cable connected to the computer, and a COM port must be specified. The PM100D communicates using a usb cable. The USB board number must be specified. In order to make this more convenient for the user blue tape has been placed on one of the PM100D red boxes, and the user is prompted to selected between the taped and non-taped power meters.

After a power meter has been initialized, behavior may be specified. The powerheads often require the wavelength of the measured laser for accurate measurement, which must be input into the control box somehow. This behavior may be set to the free mode, in which case MATLAB will not adjust the wavelength the power meter is currently set to. It may also be set to custom, in which case the user will be asked what wavelength the laser is set to, and this information will be entered into the control box one time. Finally, it may be set to lock mode. In lock mode, the wavelength of the power meter's control box will automatically be set to match the wavelength of the current laser. Lock mode is typically used with the solstis and Monaco lasers, which frequently change their wavelength.

The counts box is how many measurements the software will perform, which will then be averaged together and saved with the data. Generally 100 to 1000 counts is good enough. When calibrating a beamsplitter or filter wheel, 1000 counts or more is recommended.

Implemented Main GUI controls: Get power. This variable is stored in the data as trans-power, named after the transmission power in the beamsplitter calibrations discussed in sections 2.3 and A.2.2.

Initialize Reflection Power Meter (Samples Power) This is essentially identical to the transmission power button, except the data is stored as refl-power. The reflection name is just a name in the code, it doesn't have to necessarily measure a reflected beam. Additionally, through a beamsplitter or filter wheel calibration, the value on this meter may be predicted based on the currently measured transmission power or angle of the filter wheel.

Synchronize power meters When enabled, this puts control of the two power meters into two separate parallel workers in MATLAB. This allows both power measurements to be performed simultaneously, rather than consecutively. This is helpful if the laser's power is fluctuating over time. It can also save time when performing a beamsplitter calibration. However, there is significant overhead time in turning on MATLAB's parallel computing packages. So, if calibrating only one beamsplitter, this option is probably not worth utilizing.

Initialize Beamsplitter This checks if a Reference power meter has been initialized, and if so turns on the Load Calibration button.

Load Calibration This allows the user to load a calibration file for a beamsplitter in the experiment. It is only available if a Reference power meter has been loaded. Take note Transmission and Reflection are just names. In different optics geometries the beam reaching the sample could be the transmitted beam and the reference beam could be the reflected beam. The beamsplitter power predictions in Main GUI will be accurate all the same. For this reason, it may be clearer to refer to them as Sample and Reference power meters.

Load Filter Wheel Calibration Allows the user to load in a filter wheel calibration. This is a calibration relating filter wheel angle to the power delivered to the sample. See section A.2.3.

Initialize Thermometer Initializes communication with one of the three cryostat temperature controller boxes in the lab. The first box is the cryocon, which has two temperature probes A and B, which are both recorded and stored with data. Communications with the cryocon are handled through an RS-232 cable and the COM port number must be entered in the box before initialization. The second box is the Oxford ITC503. This box has one temperature sensor and a heater, and communications are handled using a GPIB cable. The third box is the LakeShore 321, which has one temperature sensor and a heater, and communicates via an RS-232 cable. All three boxes must have a proper calibration curve selected for their respective temperature probes. The Oxford box simply works when turned on. The LakeShore box however requires the user to select the correct curve when turning the box on for the first time each day. Notes have been written on these boxes to guide users when turning them on. Implemented Main GUI controls: set (heater) temperature when applicable, get temperature (A and B when applicable).

Initialize fw102c Filter Initializes communication with the Thorlabs FW102C six-position motorized filter wheel. Each optics table has one of these filter wheels. Five neutral density filters have been installed on these filterwheels, OD5, OD4, OD3, OD2, and OD1. In the sixth slot, no filter is installed (OD 0). This allows the intensity of captured images to be filtered up or down based on signal. For reasons not discussed here, the filters are arranged in a staggered order, 5, 0, 1, 2, 3, 4. Meaning, the first slot in the wheel corresponds to OD5 filter. And the last slot corresponds to OD4. Main GUI features a drop down panel for selecting filters, and the correspondence between OD filter and wheel position is stored in the software, so the user simply requests which OD filter they want, and the code goes to the correct position for that filter.

Initialize all APT Devices Initializes all Thorlabs APT devices. Because of the nature of the APT devices, all devices which will be used in an experiment must be initialized at the same time. The devices must first be configured to work in the Thorlabs APT software, outside of MATLAB. Then, that software should be closed.

Stepper motors which move the sample within the imaging plane of the optical system are the horizontal (H) and vertical (V) micrometers. The focal (f) micrometer is reserved for a possible future stepper motor that adjusts the focus of the image by moving the sample along the optical axis. The serial number of the corresponding stepper motor's control box must be selected from the drop down menu. Main Gui controls are: go to position.

The pump filter wheel is reserved for rotating the Thorlabs continuously variable neutral density filter wheel. Typically this is for adjusting the power of a pump laser, not the intensity of the image. Main GUI controls are: go to angle (degrees).

The Polwheel box is for controlling the Thorlabs PRM1Z8 motorized precision rotation stage. This stage has a linear polarizer in it and is utilized for quick Stokes Vector measurements. Main GUI controls are: go to angle (degrees).

Initialize Keithley Initializes communication with the Keithley 2600B System SourceMeter. This communication is handled through sending text strings back and forth over the network between the user's MATLAB instance and the Keithley box. This is done on port 5025, which is automatically forwarded to IP 192.168.1.6 inside the router that the Keithley sits behind. There are two separate voltage/current sources designated A and B. Not all of the Keithley box's functionality is implemented in Main GUI, consult David Myer's code for more advanced functionality. Main GUI essentially treats the Keithley box as a controllable voltage source, with current limit options. Implemented Main GUI controls are: set voltage of A and B (volts), Get current of A and B (amps), set current limits of A and B (amps).

Initialize Oriel white light filter This initializes communication with the Oriel spectrometer. However, rather than treating it as an imaging spectrometer, it is now treated as a tunable filter. White light may be sent into the front port, and then with proper wavelength selection, a narrow band of light may exit the back port of the spectrometer. This feature turned out to not be so useful. The idea was to do reflectivity measurements without a narrow band of white light so that we would not have to worry about absorption at higher energies causing emission at lower energies, which can skew our stopband. Main GUI controls are: go to wavelength.

A.2 Main GUI

The aptly named Main Gui is the GUI which controls all the various devices when collecting experimental data. An image of Main Gui is shown in figure 75. In this image, a camera has been connected as well as a spectrometer. An angle resolved spectral resolved image has been captured from a bottom DBR. The crosshair feature is enabled, which is frequently used for optics alignment. The profiles above and to the right of the main image are the intensity profiles along the crosshair.

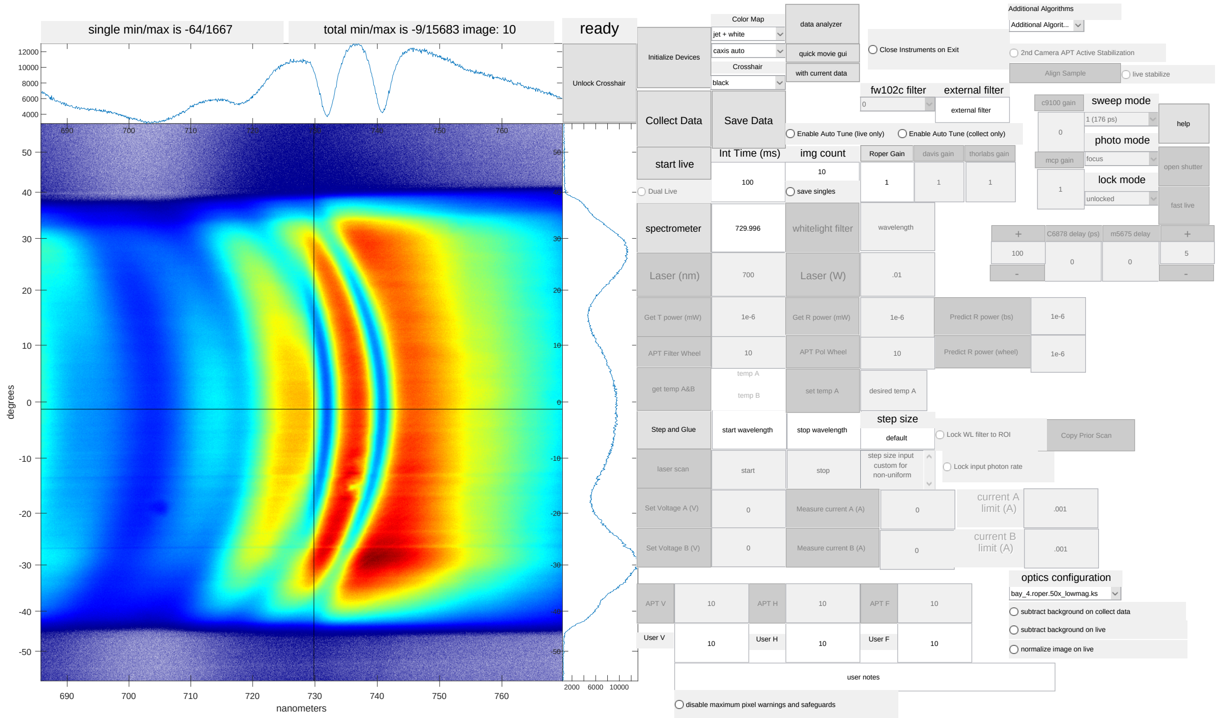
single min/max is -64/1667

total min/max is -9/15683 image: 10

ready

Color Map: jet + white

caxis auto

Crosshair: black

fw102c filter: 0

external filter: external filter

Int Time (ms): 100

img count: 10

Roper Gain: 1

davis gain: 1

thorlabs gain: 1

c9100 gain: 0

mcp gain: 1

sweep mode: 1 (176 ps)

photo mode: focus

lock mode: unlocked

spectrometer: 729.996

whitelight filter: wavelength

Laser (nm): 700

Laser (W): .01

Get T power (mW): 1e-6

Get R power (mW): 1e-6

Predict R power (bs): 1e-6

Predict R power (wheel): 1e-6

APT Filter Wheel: 10

APT Pol Wheel: 10

C6878 delay (ps): 0

m5675 delay: 0

step size: default

Set Voltage A (V): 0

Measure current A (A): 0

current A limit (A): .001

Set Voltage B (V): 0

Measure current B (A): 0

current B limit (A): .001

APT V: 10

APT H: 10

APT F: 10

User V: 10

User H: 10

User F: 10

optics configuration: bay_4.roper.50x_lowmag.ks

Figure 75: The aptly named Main Gui, responsible for controlling all the various devices in our lab.

I won't go through every button shown in figure 75. Most of the buttons are setup as input-set pairs. For example, the spectrometer button has an input box to the right of it. By clicking on the spectrometer button, the software will tune the spectrometer to that setting. In this case, I initially input 730 into the box, then clicked the spectrometer button. The software tuned as close as possible (729.996) and then updated the display to show that number in the input box. Many of the buttons work like this. Examples include but are not limited to the laser, APT filter wheel, and voltage buttons. However, I will draw attention to a few specific buttons and features. The most important features are given their own section inside this appendix.

Additional Algorithms This is a dropdown menu which allows future users to easily add new algorithms to Main GUI and execute them. Discussed in more depth in section A.2.7.

Collect Data This is the main data collection call button in Main GUI. It performs full data collections from every initialized device, updates the displayed data in the GUI, and stores the data in the data structure (structure in the MATLAB sense). The data variable can then be saved using the Save Data button. This includes things such as current filter wheel settings, current camera image, current date and time, and many other useful things. I've attempted to make this all encompassing, every detail I could thing of for every instrument is saved inside the data file. This button is also the basic building block for all the other more advanced algorithms. For example, the laser scan button changes the wavelength of the laser, and calls this button every step of the scan (and also saves the data along the way).

Enable Auto Tune (live only) This turns on the autotune feature. This adjusts the camera integration time, camera gain, and fw102c filter wheel settings to keep the maximum intensity of the displayed image at reasonable levels. If the image is too bright, it will lower intensity, and if the image is too dim it will raise the intensity. These automatic adjustments are only applied when the camera is operating in live mode, typically reserved for optical alignment and exploring the sample.

Enable Auto Tune (collect Only) This turns on the auto tuning feature for Collect Data calls. Specifically, this applies to more advanced algorithms that call the Collect Data function. For example, the Step and Glue algorithm calls the Collect Data algorithm as it sweeps the wavelength setting of the spectrometer. With this Auto Tune enabled, the camera settings are adjusted to maintain strong signal to noise ratio. This is useful in taking the N images for reflectivity as discussed in section 2.5.

lock mode This button, and all the buttons surrounding it are specifically related to the Hamamatsu streak camera. The Hamamatsu camera is a very delicate instrument, and frequently unlocks without any direct user command to do so. Every time the Collect Data button is called, this dropdown menu will be updated to display the current status. Care should be taken to observe this state, especially in sweep mode one, which is the most delicate.

Predict R Power (bs) If a beamsplitter calibration file has been created as in section A.2.2, and loaded into the Initialize Devices GUI, this button will be turned on. When clicked, it will use the calibration along with the current measurement of the power on the reference power meter to predict how much power would be measured on the sample power meter. This is updated and recorded into the data structure every time the Collect Data function is called.

Predict R Power (wheel) If a filter wheel calibration file has been created as in section A.2.3 and loaded into the Initialize Devices GUI, this button will be turned on. When clicked, it will use the calibration along with the current measurement of the APT filter wheel to predict what power would be measured on the sample power meter.

Step and Glue This button performs the step and glue algorithm as described in section 2.5.1. The starting and stopping wavelengths are entered in the boxes to the right of it in the GUI. The step size may be left as default to use the default behavior, which is usually the correct choice. However, if you wanted to force the step and glue algorithm to take step sizes of a specific size, you could input the number into this box to force that behavior. At every spectrometer wavelength setting the Collect Data function will be executed and the data stored. A good signal to noise ratio may be enforced by enabling the Auto Tune feature. The N images described in section 2.5 are taken with this button with the Auto Tune feature enabled.

Copy Prior Scan This button asks the user to load in the data file of a previously saved Step and Glue run. Then, it mimics all conditions and stores all data into the original file. This B_n, S, and B_S images described in section 2.5 are collected with this button. This button ignores the Auto Tune feature, and instead exactly mimics the conditions of the data file it is given at the start.

laser scan The starting and stopping wavelength for a laser sweep may be entered in the boxes to the right. The step size for the sweep may be entered into the step size box. The word, "custom" may be entered into the step size box, in which case upon execution the user will be prompted for a vector of laser wavelengths to use. This allows the user to specify non-uniform distribution of laser wavelengths. This button is used in PLE measurements as described in section 4.3. The custom step size is helpful for having higher density of data points near the upper polariton peak.

optics configuration This allows the user to specify the optics configuration, which is used in creating the vertical axis of the image. This is either micrometers (for real space imaging) or angle (for Fourier imaging). The vertical axis for streak camera usage is not set by this dropdown box.

subtract background on collect data If enabled, this will prompt the user to load in a background image. Whenever the Collect Data function is executed this background image will be automatically subtracted from the collected image. There is a similar button for live imaging which can be nice when aligning systems with a lot of noise.

normalize image on life If enabled, this will prompt the user to load in a reference image. During live imaging the collected image will be divided by the reference image (pixel by pixel) before being displayed. This feature can be helpful attempting to align a system for reflectivity measurements.

A.2.1 Wavelength Axis Calibration

As discussed in section 2.2, it is important to calibrate the spectrometer. By calibration we mean a correspondence between horizontal pixel number and wavelength of light after passing through the spectrometer. We do this by sweeping the laser across the screen of the camera. A convenient algorithm is present inside the Additional Algorithms dropdown menu named, “Create Grating Calibration”. This algorithm will guide the user through taking all the data necessary to create the spectrometer calibration. Afterwards, the Data Analyzer button near the top of Main GUI may be pressed to bring up a list of data processing routines. One such routine is, “Create Spectrometer Calibration”. This routine will guide the user through processing the data, and ultimately end with allowing the user to save the calibration as a data file. This file must be saved in a specific location for the calibration to be usable by Main GUI.

Each computer in the lab is a little bit different, but all of them contain a folder named, “`software_by_our_lab`”. Inside that folder is a folder named, “working”. The working folder is added to the MATLAB path during the standard software installation. The spectrometer calibration file should then be saved in the folder:

`\working\main_gui\basic_libaries\spectrapro_2500i_Aug_11_2016`

Despite its name, this folder applies to both 2500i and 2300i spectrometers. The date attached to this folder was a poorly thought out implementation of version control. It no longer has any meaning and should be ignored. Once the calibration data has been saved, the code needs to be modified so that the calibration may be automatically utilized by Main GUI when collecting data. The only code that needs to be modified is located inside the function:

`\working\main_gui\basic_libaries\camera_axis\create_camera_axis.m`

Specifically the sub-function:

`spectrapro_2500i_getwavelengthaxis` must be modified to use our new calibration file. Inside that function are a series of If-statements, based on the computer currently in use. The computer bay string is determined by the function `determine_computer_bay` which essentially looks at the static IP of the computer and outputs a string which I've hard coded into various functions used in my code. This string tells me which computer is using the code, and also tells me which devices are in use. Built into this assumption is that the computers are not being moved around the lab (usually true).

So for example, if we calibrated the bay 1 spectrometer's 50 line grating (which is in slot three) while using the streak camera, and we named the saved data file `bay1_50g_streak_cal_new`, we would need to go into the code and modify as shown below. I've removed large sections of the code to highlight the only part that needs to be changed.

```
function wavelength_axis = spectrapro_2500i_getwavelengthaxis
if strcmp(bay_string,'bay_1')
    if strcmp(camera,'streak')
        if current_grating_setting == 3
            %file = 'bay1_50g_streak_cal';
            %COMMENTED OUT THE OLD FILE
            file = 'bay1_50g_streak_cal_new';
        end
```

```
    end
end
```

With this change, data collected in the bay one computer when the spectrometer is set to grating three on the streak camera will automatically use our new calibration for the wavelength axis.

A.2.2 Beamsplitter Calibration

As discussed in section 2.3, a calibration file may be made to quickly calculate how much power is being delivered to the sample based on a reference power measurement. The code's variable names are given assuming the geometry in Figure 20. However, a reversed geometry is perfectly fine too as the math works out exactly the same. For this reason, the power meters are also called reference and sample, rather than transmission and reflection. This calibration may be created quickly using Main GUI.

First, set up the power meters so that the sample power meter measures the power delivered to where the sample will be, and the reference power meter to measure the reference power. Initialize the sweeping laser, and the two power meters in the Device Initialization GUI. Set both power meters' wavelength mode to lock to the laser source. Set the counts to be at least 1000.

Next, in Main GUI perform a laser scan covering the range wavelength range of interest. Once the sweep is done and the data is saved, click the data analyzer button and select the, "Create beamsplitter calibration" algorithm. The software will guide the user through creating the calibration file, which may then be saved. Typically, a beamsplitter calibration is a more temporary thing (opposed to a spectrometer calibration), and so often the calibration file is just saved along with the experiment's data. The file is then loaded into Main GUI through the Device Initialization GUI.

Once complete, the system is prepared to predict power delivered to the sample based on measurements of the reference beam.

A.2.3 Filter Wheel Calibration

When working at a single wavelength, a calibration between the filter wheel's angle and the power delivered to the sample may be created for easy reference. To do this, place a single power meter in the sample plane of the optics setup, and place a filter wheel in the path of the beam before it reaches the power meter. After Initializing the filter wheel in the Device Initialization GUI, the APT Rotation Scan algorithm will be available inside the Additional Algorithms dropdown menu.

This algorithm will sweep the angle of the filter wheel while recording the power of the power meter. Once the data has been taken and saved, the "Create filter wheel calibration" algorithm inside of the Data Analyzer tool may be used to create the calibration file. This file may then be loaded into Main GUI through the Device Initialization GUI.

Once complete, the power delivered to the sample may be approximated based off the angle of the filter wheel.

A.2.4 Engaging Active Stabilization

As discussed in section 2.6 and 6.3, active stabilization is a method of moving samples back into their correct position after a random vibration has moved them out of position. Here I will outline how to engage the active stabilization inside of Main GUI. First, a dual imaging setup such as the one shown in Figure 28 must be utilized.

The process is somewhat complicated, not intuitive, and involves four separate instances of MATLAB running at the same time. The reason for this is because a single instance of MATLAB can not support two different GUIs running at the same time. This is a problem, because we wish to collect data from our CCD camera while also monitoring sample position on our USB camera. Preferably we want the user to be able to look at images on the screen and monitor the experiment as it progresses. Therefore it was necessary to expand beyond a single MATLAB instance. One instance will run the CCD camera and collect our main experimental data, and the other will monitor sample position. However, we would also like information to be shared among these two MATLAB instances. For example, the APT positioning is constantly being tuned for the active stabilization scheme, but we would also like to record that information with our experimental data so that it is available when processing the data after the experiment. The simplest way to do this is to run certain instruments in their own separate MATLAB instance, acting as a server, and providing their measurements to any other MATLAB instance which communicates with them via TCPIP calls on port zero.

The four MATLAB instances are listed below.

1. This instance has Main GUI one running on it, and coordinates the active stabilization scheme.
2. This instance runs the APT stepper motors and provides data when queried to instance one and four.
3. This instance runs the USB camera and provides data when queried to instance one and four.
4. This instance has Main GUI two running on it, and collects all other experimental data such as CCD camera images, power meter readings, temperature readings, etc.

Now I will list out the steps required to turn on active stabilization.

- Start from a clean slate. All instances of MATLAB should be turned off.
- Open instance one of MATLAB and open Main GUI one on it.
- In Main GUI one, initialize the APT devices. Read and follow the popup dialogue boxes, which will guide the user through opening MATLAB instance two, with the APT devices running as a server. It is recommended to zero the APT devices before proceeding.
- In Main GUI two, initialize the USB camera. Read and follow the popup dialogue boxes, which will guide the user through opening MATLAB instance three, with the USB camera running as a server. This should have opened the USB camera inside of the Thorlabs DCC GUI. Recall this GUI also is used for the Zeiss USB cameras.
- Open MATLAB instance four, and open Main GUI two on it. This GUI will be the one collecting the experimental data. Internalize all necessary devices on it such as the CCD camera, spectrometer, power meters, etc.

- Move the sample into the position you wish to lock to. The CCD camera in Main GUI two and the USB camera in the Thorlabs DCC GUI may be put into live imaging at the same time to aid with this. Once in position, turn off all live imaging.
- Make sure the image collection settings in the Thorlabs DCC GUI are set the way you want them for the entirety of the experiment. Use the Snap Image button in the Thorlabs DCC GUI to take the reference image.
- Use the Load Reference button inside the Thorlabs DCC GUI to load in the reference image. This tells the computer what the desired sample location is. Turn on the Overlap With Reference option and the Output Affine Transform option.
- In Main GUI one, in the top right corner, turn on the 2nd Camera APT Active Stabilization feature. This will guide the user through the setup procedure, including the creation of a calibration between pixels on screen and micrometers in the APT stepper motors. The sample will be moved a set amount in order to create this calibration. A minimum of 10 micrometers for this movement is recommend, and a larger movement will result in more accurate calibrations.
- In Main GUI one, turn on the Live Stabilize feature. Main GUI one will now periodically collect data from the USB camera and reposition the sample if it has moved out of location.
- Main GUI two may now be used to perform the experiment (examples would be voltage series, temperature series, etc.). Additionally, Main GUI two may initialize the USB camera and APT stepper motors (accessing their servers as a second client), in order to record their measurements.

A.2.5 Adding A New Laser

Main GUI is programmed to be able to easily accommodate new lasers being added to it. First, a new button must be added to the Initialize Devices GUI. Most of the new code for that button could be copied from the MSquared laser's button, with small changes utilizing the name of the new laser. Inside of Main GUI, the laser functions are defined using function handles inside of the Initialize Devices Callback function, new code would be added to the related switch-case statement. A simplified example is shown below.

```
switch handles.data.lasername
        case 'msquared'
            handles.laser_functions.getwavelength = @
               solstis_getwavelength;
            handles.laser_functions.gotowavelength = @
               solstis_gotowavelength;
        case 'new_laser'
            handles.laser_functions.getwavelength = @
               newlaser_getwavelength;
            handles.laser_functions.gotowavelength = @
               newlaser_gotowavelength;
```

Here a new laser named newlaser has been added to the code. All calls to laser functions inside of Main GUI are handled by calling the functions inside the handles structure. Therefore, by defining these functions using the functions of the new laser, all existing algorithms in Main GUI will automatically be able to utilize the new laser.

A.2.6 Adding A New Camera

Main GUI is programmed to be able to easily accommodate new cameras being added to it. First, a new button must be added to the Initialize Devices GUI. Most of the new code for that button could be copied from the Roper Camera's button, with small changes utilizing the name of the new camera. Inside of Main GUI, the camera functions are defined using function handles inside of the Initialize Devices Callback function, new code would be added to the if statement. A simplified example is shown below.

```
if strcmp(handles.devices.primary_camera.name,'roper')
    handles.camera_functions.set_int_time = @
       roper_ccd_setinttime;
```

```
    handles.camera_functions.start_live = @roper_ccd_start_live
       ;
    handles.camera_functions.stop_live = @roper_ccd_stop_live;
    handles.camera_functions.get_image_live = @
       roper_ccd_getimagefast;
    handles.camera_functions.get_image = @roper_ccd_getimage;
end
if strcmp(handles.devices.primary_camera.name,'new_camera')
    handles.camera_functions.set_int_time = @
       new_camera_ccd_setinttime;
    handles.camera_functions.start_live = @
       new_camera_ccd_start_live;
    handles.camera_functions.stop_live = @
       new_camera_ccd_stop_live;
    handles.camera_functions.get_image_live = @
       new_camera_ccd_getimagefast;
    handles.camera_functions.get_image = @
       new_camera_ccd_getimage;
end
```

Here a new camera named has been added to the code. All calls to camera functions inside of Main GUI are handled by calling the functions inside the handles structure. Therefore, by defining these functions using the functions of the new camera, all existing algorithms in Main GUI will automatically be able to utilize the new camera.

A.2.7 Adding A New Algorithm

The Additional Algorithms dropdown menu in the top right corner of Main GUI is designed to be an easy way for users to add their own custom tailored algorithms to Main GUI. Before adding a new algorithm, make sure an existing algorithm can not already accomplish the desired outcome. Specifically the Generalized Scan algorithm is very versatile and can handle many new situations.

If no existing algorithms are viable, new algorithms may be easily added. First, the additional algorithms dropdown menu is populated inside of the Initialize Devices Callback function. Inside this Callback function the necessary devices are checked for, and if present, the algorithm is added to the dropdown menu. Then, inside the Additional Algorithms Button Callback the switch-case statement must be modified so that if the new algorithm is selected by name, the associated function is called.

Finally, the new function must be defined and added into the Main GUI code. For clarity in the code, please add new algorithms into the section of the code containing the other additional algorithms. These functions are connected together with comments to indicate they are all part of the Addition Algorithms dropdown menu.

A.2.8 Adding A New Optics Configuration

Optics configurations are how Main GUI may choose the calibration for real space and Fourier space axes. To add a calibration for a new optics configuration, first the follow function must be modified.

`\working\main_gui\basic_libaries`
    `\camera_axis\get_possible_optics_configurations`

This function is where the name for all possible optics configurations is stored as a cell array. The Get Plausible Optics Configurations then narrows this array down to only optics configurations related to the current computer bay. Take care to name your optics configuration in a way that conforms with the pattern shown. The name is a series of words with periods in between them. The first word should be a possible output string from the Determine Computer Bay function, indicating which optics table this configuration exists on. The last word should be rs or ks, to indicate real space (units of microns) or ks (units of degrees). The middle words may be freely chosen by the user for convenience and reference.

```
output{end+1} = 'bay_1.short.50x.ks';
output{end+1} = 'bay_1.new_setup.5x.ks';  %the new
   configuration
```

Here, I have shown how to add a new angle resolved setup for a 5x microscope objective in bay one. Now some code and the calibration must be added to the function:

`\working\main_gui\basic_libaries\camera_axis\create_camera_axis`

In this case because we are adding a non-spectral non-temporal axis, we must utilize the specific sub-function:

`create_non_spectral_non_temporal_camera_axis_bay1`

Simplified example code is shown below.

```
function create_non_spectral_non_temporal_camera_axis_bay1
if strcmp(optics_string,'bay_1.short.50x.ks') && strcmp(camera,
   'roper')
     p(1) = 0.32;
     p(2) = -65;
     output_axis = polyval(p,1:axis_length);
     output_units = 'degrees';
end
if strcmp(optics_string,'bay_1.new_setup.5x.ks') && strcmp(
   camera,'roper')
     p(1) = .6;
     p(2) = -15;
     output_axis = polyval(p,1:axis_length);
     output_units = 'degrees';
end
end
```

The values of the P vector determine the calibration. They are obtained from the calibration procedure outlined in section 2.1. Once this is accomplished, the calibration will be used in Main GUI when this new optics configuration is selected from the Optics Configuration dropdown menu.

A.3 Quick Movie Gui

Next I will discuss the Quick Movie GUI. This is a GUI that specializes in quickly displaying data saved by Main GUI, and also processing it. An example image of Quick Movie GUI is shown in figure 76.

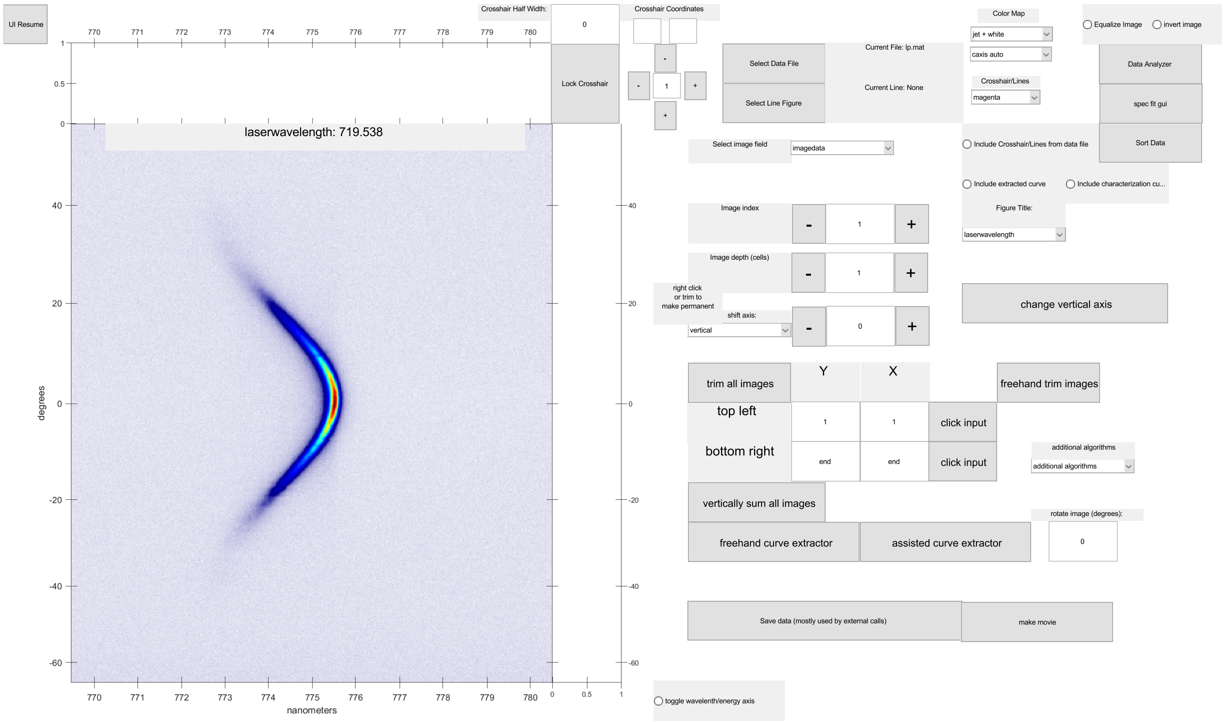

Crosshair Half Width: 0

Current File: lp.mat

Current Line: None

laserwavelength: 719.538

Select image field: imagedata

Color Map: jet + white

caxis auto

Crosshair/Lines: magenta

Figure Title: laserwavelength

Image index: 1

Image depth (cells): 1

shift axis: vertical 0

top left: Y 1, X 1

bottom right: Y end, X end

additional algorithms: additional algorithms

rotate image (degrees): 0

Figure 76: Quick Movie Gui displaying a set of PLE data saved by Main GUI.

I won't discuss every button in Quick Movie GUI, the uploaded videos [11] show many of its functions, but I will discuss a few of its buttons.

Select Data File This loads the data saved by Main GUI, and begins displaying it.

Select Line Figure This loads a standard MATLAB figure and plots it on top of the image. All lines in the figure are plotted on top of the image.

Image Index If the loaded data set was created through some sort of sweep, the plus and minus buttons here march through the sweep.

Figure Title This dropdown allows the user to select from various data variables and plot them on top of the image as a title. In Figure 76 the data was taken as a PLE laser sweep, and the laser's current wavelength setting is displayed and updated as the Image Index value is marched along.

Shift Axis This may be used to shift the angle axis if the $\theta = 0$ location has been shifted.

Trim All Images These buttons may be used to define a box's top left and bottom right corners. When executed, the trimmed out portion of the image will be kept and the scraps will be discarded. This trimming will be performed on all images in the data set. This is helpful for creating zoomed-in images of the data.

Freehand Trim Images This allows the user to draw an irregular shape on the image and keep all pixels inside or outside of that shape. The discarded pixels may be set to any value, including zero or NaN depending on what is needed. This function is useful for quickly removing noise from an image before performing other routines on it.

Freehand Curve Extractor This allows the user to draw on an image and then export the corresponding curve in the units displayed on the image. This is helpful for quick approximate curve extractions, but usually the Assisted Curve Extractor button is better.

Assisted Curve Extractor This takes the current image and opens it in the Polariton Fitter GUI. Discussed in section A.3.1.

Equalize Image This turns on image equalization, which independently normalizes each row of the image to scale between zero and one. This function is great at removing thermal distribution effects and making dispersion curves more visible.

Additional Algorithms Similar to Main GUI's Additional Algorithms dropdown box, this is an area for users to easily add new algorithms to Quick Movie GUI. The Recenter Horizontal Axis algorithm is useful for extracting the dispersion curve of the polaritons when the signal is very strong such as in figure 76. This algorithm requires the Rotate All Images algorithm to be called once first. The Rotate All Images (x3) algorithm may be called afterwards to put the image back into the standard orientation.

A.3.1 Polariton Fitter

The Polariton Fitter GUI is used to extract curves with faint signals, as discussed in section 5.2. An example of this GUI is shown in Figure 77.

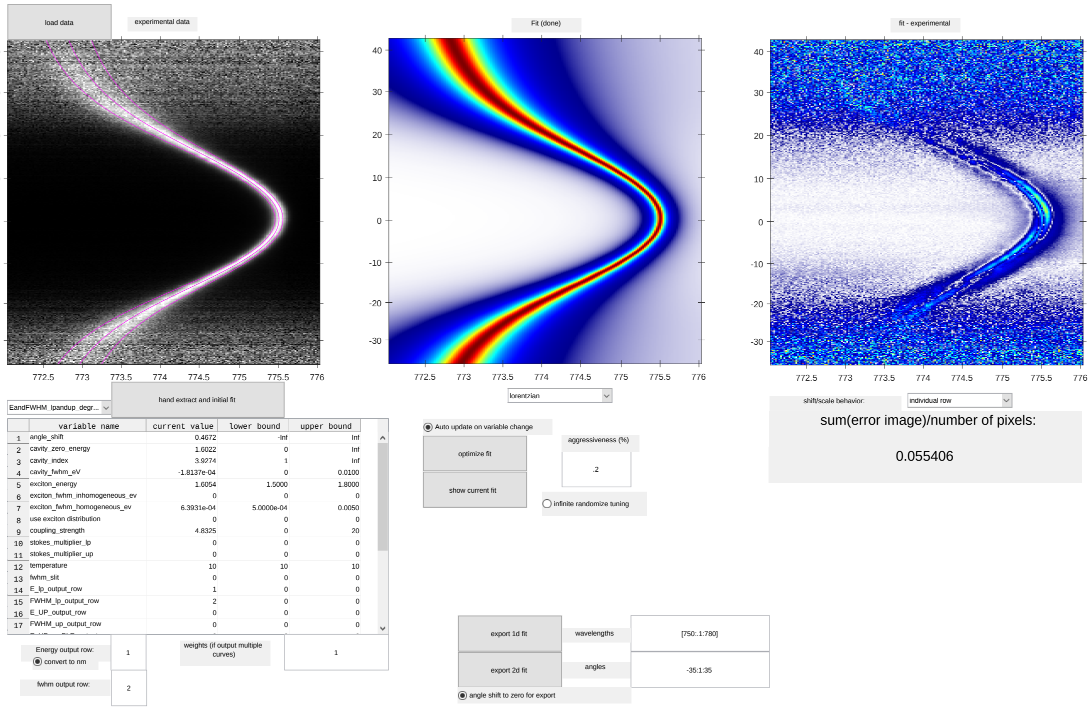

EandFWHM_lpandup_degr...

	variable name	current value	lower bound	upper bound
1	angle_shift	0.4672	-Inf	Inf
2	cavity_zero_energy	1.6022	0	Inf
3	cavity_index	3.9274	1	Inf
4	cavity_fwhm_eV	-1.8137e-04	0	0.0100
5	exciton_energy	1.6054	1.5000	1.8000
6	exciton_fwhm_inhomogeneous_ev	0	0	0
7	exciton_fwhm_homogeneous_ev	6.3931e-04	5.0000e-04	0.0050
8	use exciton distribution	0	0	0
9	coupling_strength	4.8325	0	20
10	stokes_multiplier_lp	0	0	0
11	stokes_multiplier_up	0	0	0
12	temperature	10	10	10
13	fwhm_slit	0	0	0
14	E_lp_output_row	1	0	0
15	FWHM_lp_output_row	2	0	0
16	E_UP_output_row	0	0	0
17	FWHM_up_output_row	0	0	0

Energy output row: 1

● convert to nm

fwhm output row: 2

weights (if output multiple curves): 1

● Auto update on variable change

lorentzian

aggressiveness (%): .2

○ infinite randomize tuning

wavelengths: [750:.1:780]

angles: -35:1:35

● angle shift to zero for export

shift/scale behavior: individual row

sum(error image)/number of pixels:

0.055406

Figure 77: The Polariton Fitter GUI, after fitting the lower polariton PL data from a PLE sweep.

The first thing to do when loading up this GUI is to select a fitting function from the dropdown menu. Typical the function EandFWHM_lpandup _degrees is utilized. The output of this function may be tailored to upper or lower polaritons, and the linewidths may be set as a second output.

Once the fitting function is selected, the Hand Extract And Initial Fit function is utilized. The user then draws an approximate curve on the image indicating where the curve to extract is. This hand drawn curve is then exported to the Spec Fit GUI, and the user should fit the hand drawn curve using the fitting function. Once that fit is complete, pressing the UI Resume button on Spec Fit GUI will close Spec Fit GUI and return the fit to Polariton

Fitter.

Then, the two dimensional fit described in section 5.2 may be performed by pressing the Optimize Fit button. This should result in a fit of the two dimensional curve to the data. One such fit is shown in the left most plot of the GUI in Figure 77. The fitted curve is shown with the pink lines plotted on top of the data. The two dimensional fitting function is plotted in the middle image, and the error is plotted in the third image.

Once these fits are done, a dispersion curve may be exported as a figure by the Export 1D Fit button. This exported figure may be fed into the Reflectivity GUI for simulation fitting as described in appendix B.

A.4 Data Analyzer

Data Analyzer is an incredibly simple GUI just consisting of a table from which the user may selected a desired algorithm to perform. A few of the algorithms of interest will be highlighted here.

Specfit - Extract Curves From Image And Store It This is used to do a row by row fitting of an image to extract the dispersion curve, as discussed in section 4.2. This works well when signal is very bright, but poorly if the signal is dim.

PLE Measurement For Single Spatial Point This is the basic PLE data processing algorithm, used to implement the dynamic background subtraction scheme described in section 4.3. This algorithm will walk the user through its use.

Reflectivity Calculator This is used to process the results of the Copy Prior Scan algorithm from Main GUI, resulting in reflectivity images as discussed in section 2.5. It also has the ability to stitch the images together.

Create Spectrometer Calibration Processes the data as discussed in section A.2.1.

Create Beamsplitter Calibration Processes the data as discussed in section A.2.2.

A.5 Spec Fit Gui

Spec Fit GUI is used for quickly fitting one dimensional data with fitting functions. An example of Spec Fit GUI being used to fit some data from a PLE sweep is shown in Figure 78.

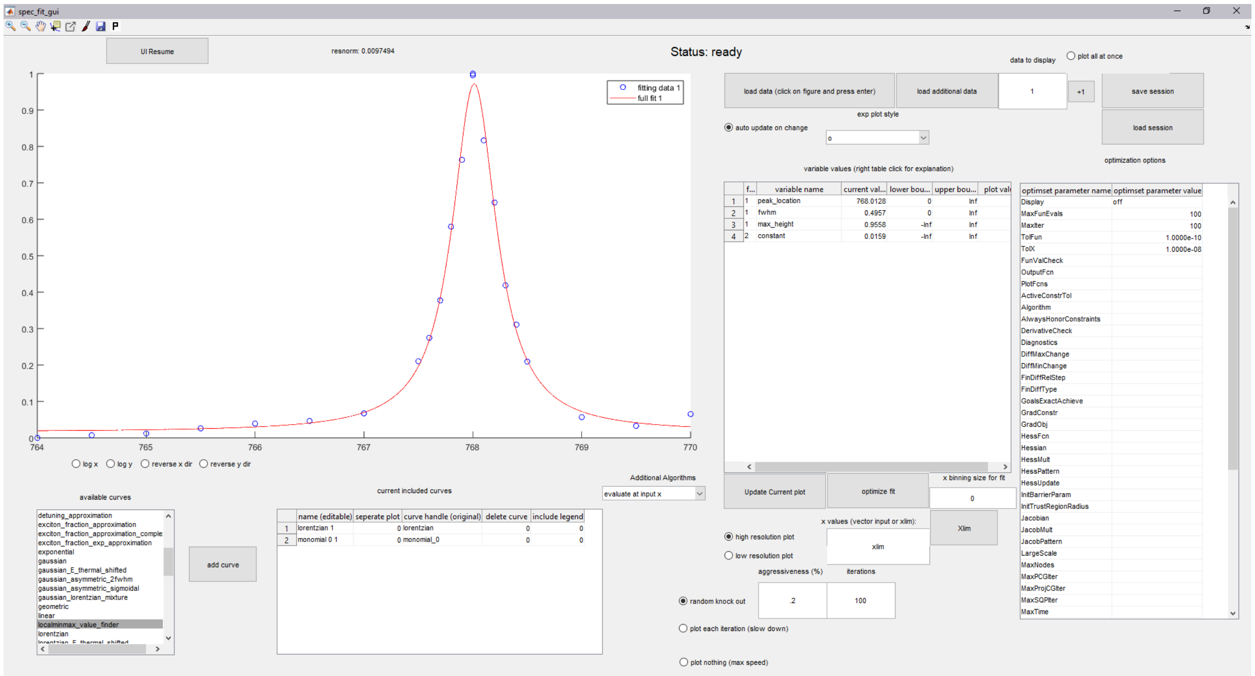

resnorm: 0.0097494

data to display: 1

● auto update on change

exp plot style: o

variable values (right table click for explanation)

	f...	variable name	current val...	lower bou...	upper bou...	plot val
1	1	peak_location	768.0128	0	Inf	
2	1	fwhm	0.4957	0	Inf	
3	1	max_height	0.9558	-Inf	Inf	
4	2	constant	0.0159	-Inf	Inf	

optimset parameter name	optimset parameter value
Display	off
MaxFunEvals	100
MaxIter	100
TolFun	1.0000e-10
TolX	1.0000e-08
FunValCheck	
OutputFcn	
PlotFcns	
ActiveConstrTol	
Algorithm	
AlwaysHonorConstraints	
DerivativeCheck	
Diagnostics	
DiffMaxChange	
DiffMinChange	
FinDiffRelStep	
FinDiffType	
GoalsExactAchieve	
GradConstr	
GradObj	
HessFcn	
Hessian	
HessMult	
HessPattern	
HessUpdate	
InitBarrierParam	
InitTrustRegionRadius	
Jacobian	
JacobMult	
JacobPattern	
LargeScale	
MaxNodes	
MaxPCGIter	
MaxProjCGIter	
MaxSQPIter	
MaxTime	

available curves:

detuning_approximation
exciton_fraction_approximation
exciton_fraction_approximation_comple
exciton_fraction_exp_approximation
exponential
gaussian
gaussian_E_thermal_shifted
gaussian_asymmetric_2fwhm
gaussian_asymmetric_sigmoidal
gaussian_lorentzian_mixture
geometric
linear
localminmax_value_finder
lorentzian
lorentzian_E_thermal_shifted

current included curves

	name (editable)	seperate plot	curve handle (original)	delete curve	include legend
1	lorentzian 1	0	lorentzian	0	0
2	monomial 0 1	0	monomial_0	0	0

Additional Algorithms: evaluate at input x

x binning size for fit: 0

x values (vector input or xlim): xlim

● high resolution plot
○ low resolution plot

	aggressiveness (%)	iterations
● random knock out	.2	100

○ plot each iteration (slow down)

○ plot nothing (max speed)

Figure 78: Spec Fit GUI being used to fit some PLE data.

Spec Fit GUI is designed to be easy to users to add new fitting functions to. Any function found in the folder:
`\working\spec_fit\fittable_functions` is automatically added to the Available Curves table in Spec Fit GUI. These functions must adhere to a strict format where if the function is called with no input the output is a structure array with variable names and typical ranges for the variable. If making a new function, it is a good idea to copy and paste the code from one of the simple functions such as the Lorentzian function.

By default, the random knock out function is turned on with an aggressiveness of 0.2% and 100 iterations. This means the curve will be fit using MATLAB's standard least squares minimization. Then the best fit parameters will be randomized to within 0.2% of their current values and the minimization will occur again. This process will be repeated 100 times, and the best fit found is the final fit displayed.

The value of the fitting parameters and the bounds for them are displayed in the variables table. By pressing the letter s while editing one of these values, the GUI enters Smart Input Mode. In this mode, you may click on the plot of the data as a quick way to input parameters. For example, if you are editing the value of a Lorentzian line center and press S, you may click once on the screen to set the line center as the location you clicked. If you are editing the linewidth you may click twice to set the width.

There are many useful functions inside of the Additional Algorithms dropdown menu. Specifically referenced in this work is the "Sweep Variable, Get Best Fits algorithm". This algorithm can be used to sweep one variable while letting the other variables be tuned to produce the best constrained fit. The result is a plot of the global error as a function of the sweep parameter. This may be used to test how tightly constrained a parameter is. This was used in section 4.4 to test how important the PLE measurement is in constraining a Rabi model, specifically in Figures 48 and 49.

Appendix B Software Manual For TMM GUIs

This appendix shall serve as a software manual for the GUIs related to transfer matrix method (TMM) simulation of our samples. Videos demonstrating the use of these GUIs are available online [11]. The central GUI is named Reflectivity GUI, which is responsible for running the simulations, and tuning the parameters described in section 4.5.1 and 4.5.2 to produce agreement between simulation and experimental data. Additionally, I will discuss the Sample Creator GUI which is used to make the structure files which will be loaded into Reflectivity GUI. The GUI named E Field GUI is used to display images of the sample with the simulated electric field plotted on top of it. The final GUI is named Reflectivity GUI Simulfit, and is used to run multiple instances of Reflectivity GUI in a linked manner (for example, forcing exciton energy to be the same in each instance).

B.1 Sample Creator GUI

The Sample Creator GUI is used to quickly build the sample structures which are loaded into the other GUIs. Sample structures are saved in a homemade file format; .jbsample, which is just a standard MATLAB .mat file with a different name. This alternative name was chosen just to make sorting and organization of files cleaner.

An example of Sample Creator GUI is shown in figure 79. Layers are added one at a time by selecting the material from the dropdown menu, and then entering the desired thickness. The layer may be added directly to the sample, or to the multi layer section below. In this case, we have added two layers to the multi layer section. Then, I stipulated 15 repeats, and added the multi layer section to the main sample. The 15 repetitions produces the DBRs. In order to construct the top DBR I used the Flip Multi Layer and then added it to the sample again. In between the two DBRs I manually placed a layer of GaAs quantum well, which has the tunable index discussed in section 4.5.2.

If a thick substrate has been included, it may be scaled down using the buttons in the

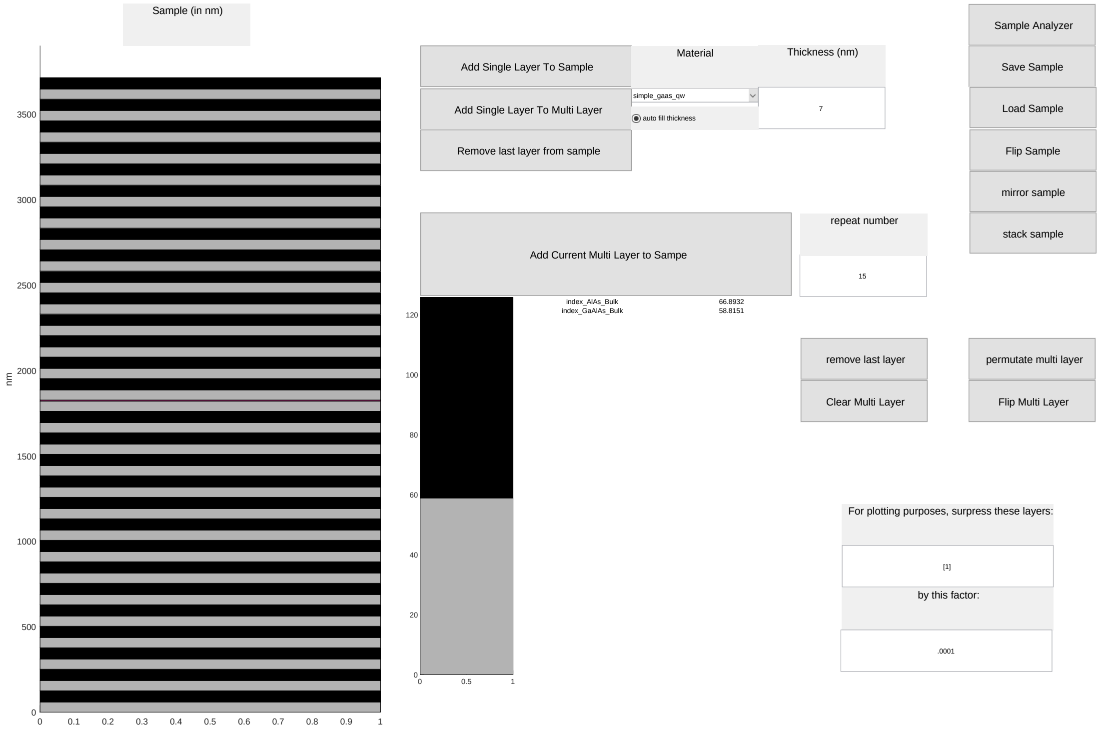

Material: simple_gaas_qw

● auto fill thickness

Thickness (nm): 7

repeat number: 15

index_AlAs_Bulk 66.8932

index_GaAlAs_Bulk 58.8151

For plotting purposes, surpress these layers: [1]

by this factor: .0001

Figure 79: The Sample Creator GUI in use. Here, we have made a simple microcavity sample using the auto fill feature.

bottom right. This scaling only affects what is plotted, but does not affect the actual saved sample file, or the resulting simulations.

The Auto Fill Thickness feature was used in this case. This feature is useful when designing new simple samples. Through proper configuration it may be used to create layers which are each one quarter of the design wavelength, the standard DBR design discussed in section 1.6. Although in real samples a small correction is necessary for the bottom DBR to accommodate the wetting layer.

In the top right are some buttons to help combine and stack samples together. This was useful in TMD sample design, which is often done in a multiple step process rather than one continuous growth. The Sample Analyzer button opens up a panel with a few helpful algorithms. In particular, it has algorithms to insert additional layers into an existing sample, as well as manually change specific layers or materials.

B.2 Reflectivity GUI

The Reflectivity GUI is the central GUI for the TMM simulation. It runs the underlying simulation, and is used by the user to define the various fitting parameters. It is also responsible for tuning the fitting parameters to bring the simulation into agreement with the experimental data.

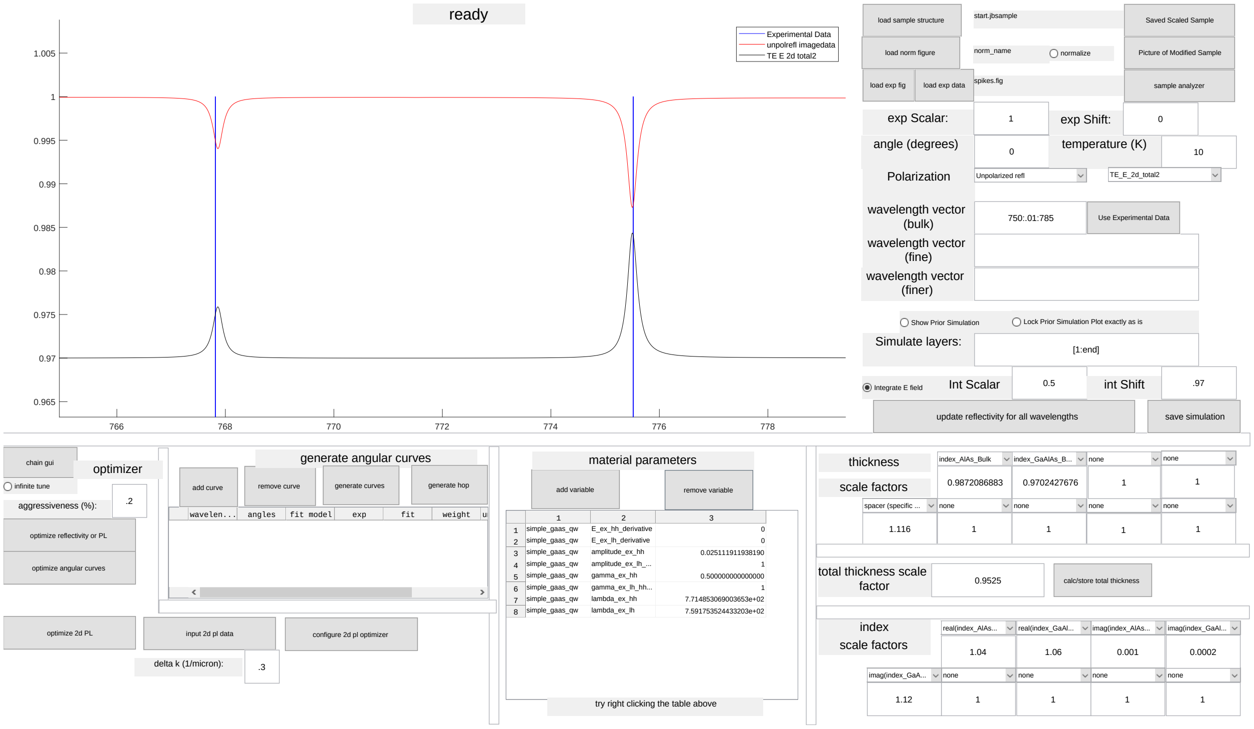

start.jbsample
norm_name ○ normalize
spikes.fig
exp Scalar: 1
exp Shift: 0
angle (degrees) 0
temperature (K) 10
Polarization Unpolarized refl TE_E_2d_total2
wavelength vector (bulk) 750:.01:785
wavelength vector (fine)
wavelength vector (finer)
○ Show Prior Simulation ○ Lock Prior Simulation Plot exactly as is
Simulate layers: [1:end]
● Integrate E field Int Scalar 0.5 int Shift .97
○ infinite tune
aggressiveness (%): .2
delta k (1/micron): .3

wavelen...	angles	fit model	exp	fit	weight	u

	1	2	3
1	simple_gaas_qw	E_ex_hh_derivative	0
2	simple_gaas_qw	E_ex_lh_derivative	0
3	simple_gaas_qw	amplitude_ex_hh	0.025111911938190
4	simple_gaas_qw	amplitude_ex_lh_...	1
5	simple_gaas_qw	gamma_ex_hh	0.500000000000000
6	simple_gaas_qw	gamma_ex_lh_hh...	1
7	simple_gaas_qw	lambda_ex_hh	7.714853069003653e+02
8	simple_gaas_qw	lambda_ex_lh	7.591753524433203e+02

thickness scale factors	index_AlAs_Bulk	index_GaAlAs_B...	none	none
	0.9872086883	0.9702427676	1	1
spacer (specific ...	none	none	none	none
1.116	1	1	1	1

total thickness scale factor 0.9525

index scale factors	real(index_AlAs...	real(index_GaAl...	imag(index_AlAs...	imag(index_GaAl...
	1.04	1.06	0.001	0.0002
imag(index_GaA...	none	none	none	none
1.12	1	1	1	1

Figure 80: Reflectivity GUI, after being configured by a user, ready to begin fitting the simulation to polariton data.

An example of Reflectivity GUI is shown in figure B.2. Here, I will discuss the various buttons and what they do.

Load Sample Structure Loads the sample file which was created by the Sample Creator GUI. The name of the currently loaded sample is shown to the right.

Load Norm Figure Loads a mirror correction file. This fixes the fact that the N images discussed in section 2.5 were not taken off a perfect reflector. It is the $R_M(\lambda, \theta)$ from equation 2.10. The curves are stored in the folder:

`\working\sample_simulation\normalizer_correction_figures`

Load Exp Fig Loads the experimental data from a figure. The name of the figure is displayed to the right. This is the data that the Optimize Reflectivity Or PL button fits the simulation to.

Load Exp Data Opens a Quick Movie GUI, which is linked to the current Reflectivity GUI. The crosshair profile in the linked Quick Movie GUI will be the experimental data inside Reflectivity GUI. I find this button not so useful, as the Chain GUI accomplishes a similar goal in a cleaner manner.

Save Scaled Sample Saves the sample structure after the various modifications made by Reflectivity GUI. In practice I find this button not so useful. Instead, I find the Save GUI button in the top left much more useful and more organized. The Save GUI button saves an instance of the reflectivity GUI. Double clicking the file inside of MATLAB will then open the GUI exactly as it was when saved.

Picture Of Modified Sample Opens the current sample in the E Field GUI. See section B.3.

Sample Analyzer Opens the same Sample Analyzer panel as discussed in section B.1.

Exp Scalar Scales the experimental data by a constant. One is no scaling. This feature is useful because often experimentally measured reflectivity has error resulting in similar scalings.

Exp Shift Shifts the experimental data up or down by a constant. Useful for offsetting noise.

Angle (degrees) The angle of incident light being simulated. If the input is a vector a two dimensional simulation will be displayed in a link instance of Quick Movie GUI.

Temperature (K) The temperature of the simulation. Not all index functions have temperature dependence built into them. The GaAs, AlAs, and AlGaAs index functions do.

Polarization The dropdown menu here allows the user to select what the simulation data is. Polarizations may be selected between TE, TM, unpolarized (average of TE and TM) and the ellipsometric ratio which is calculated as TM divided by TE. All four polarizations are available as power reflectivity, power transmission, and power absorption.

If the Integrate E Field button below is enabled, a second plot determined by the second dropdown menu is plotted. There are many options in this menu. The “2d” in the names is a book keeping trick, and not describing the fields in anyway, just ignore it. The user may select various spatial integrals inside of selected layers of the sample, such as the in-plane or out of plane components of the E or B field. As well as quantities such as dissipation and total energy density. In this case, the total E field’s magnitude squared is shown for a TE wave. Of course, the out of plane component for TE is zero, so this is also a plot of the in-plane E field’s magnitude squared.

Wavelength Vector These three inputs allow the user to stipulate what wavelengths the simulation occurs over. All three are combined, allowing the user to easily stipulate higher resolution in regions of interest. The User Experimental Data button fills in the bulk to be the same range as the experimental data.

Show Prior Simulation The previous simulation will be plotted as well. This makes it easy to see the affect of a change on the resulting reflectivity.

Lock Prior Simulation Plot Exactly As Is Locks the previous simulation, rather than updating it with each new simulation. This is helpful to see how a series of changes modifies the reflectivity curve.

Simulate Layers This input allows the user to remove or repeat layers of the sample. For example, if the user were fitting the reflectivity data of a sample in which the top DBR had been etched away, the following input might be warranted: $[30 : end]$. Utilize the Picture of Modified Sample button to verify the desired changes are being implemented with this input. Right click to bring up additional helpful options.

Integrate E Field This turns on E field integration (and a lot more). It will ask the user which layers they would like to integrate the E field inside of. The integration is a simple numerical spatial integral. This is accomplished by breaking the sample layers up into many smaller layers. The program will ask the user how many layers to break each layer in to.

When enabled, this option also results in the data necessary for the E Field GUI to display the electric field on top of the sample.

By default, the TMM simulation only calculates the E field at the boundaries of the sample. This can give the E field a very jagged appearance. To produce publication quality plots of the E field, turn this option on, and integrate the E field inside every layer of the sample. Breaking the layers into 20 pieces seems to work well.

The integration scalar and shift inputs move and scale the simulation as plotted inside Reflectivity GUI. They do not influence any of the physical results.

Update Reflectivity For All Wavelengths This is the main button to run the simulation and update the display of everything in the GUI.

Thickness Scale Factors These nine dropdown menus allow the user to select groups of layers and scale them. Typically the DBR materials are selected as a group. In the shown example, every AlAs layer is being scaled by a factor of 0.987. By Selecting the Specific Layers option, the user may define a custom group of layers and assign them a name. This is useful for scaling the spacer layers inside of the cavity.

Total Thickness Scale Factor This uniformly scales every layer of the sample. In the GaAs samples, moving across the sample can be captured fairly well by only adjusting this value.

Index Scale Factors These dropdown menus allows the user to select materials in the sample, and scale the real and imaginary portions of the index of those materials. In the shown example, the real part of the index of AlAs is being scaled up by 1.04, and the imaginary part is being scaled down by 0.001.

Material Parameters Certain materials have programmed index functions which support more complicated modification, as opposed to the simple scaling used in the Index Scale Factors portion of the GUI. For example, the GaAs quantum wells have the excitons which are controlled by several parameters discussed in section 4.5.2. These variables are first added to the table using the Add Variable button. Their value may then be set by the user modifying the table's values. The variables in this table may also be tuned by the various optimization buttons discussed below.

Chain GUI This opens up another instance of Reflectivity GUI. That instance may be programmed to operate at a different temperature, angle, and polarization and fitted to a different set of data. However, the chained GUIs will utilize the same sample. This allows the user to fit multiple sets of data from the same location as one unified fit. Three such GUIs are commonly used to fit the simulation to the broad reflectivity as discussed in section 4.5.1.

Infinite Tune When enabled, this modifies the behavior of the "Optimize Reflectivity Or PL" and "Optimize 2d PL' buttons. Normally, those two buttons perform a simple curve fitting error minimization routine to tune the fitting parameters as defined by the user. When Infinite Tune is enabled, those optimizers will then randomize the fitting parameters by the percentage given in the aggressiveness input, and optimize them again. This is repeated for as long as the Infinite Tune option is enabled. Typically, such fits are performed for an hour or more, until the user is satisfied with the fit. Then, the user may turn off the Infinite Tune, and wait a few minutes for the current fit to finish, and then view the final results.

Optimize Reflectivity Or PL This fits the current simulation to the loaded experimental data. See section B.5.

Optimize Angular Curves This is no longer a supported function. Don't use it.

Optimize 2d PL The 2d PL buttons are used to fit the polaritons extracted from the simulation to the experimental polariton data. See section B.6.

Generate Angular Curves This portion of the GUI is largely abandoned and no longer in use, as the 2d PL buttons have replaced it with more efficient algorithms. However, this is still the method of generating the exciton fraction, as discussed in section B.6.

B.3 E Field GUI

The E Field GUI is a convenient way of displaying images of the sample, along with various fields plotted on top of it. It may be opened with the corresponding button in Reflectivity GUI. The displayed sample is the sample after the modifications imposed by the parameters in Reflectivity GUI. An example of the E Field GUI is shown in Figure 81.

In order to produce electric fields as shown in the example image, the Integrate E Field

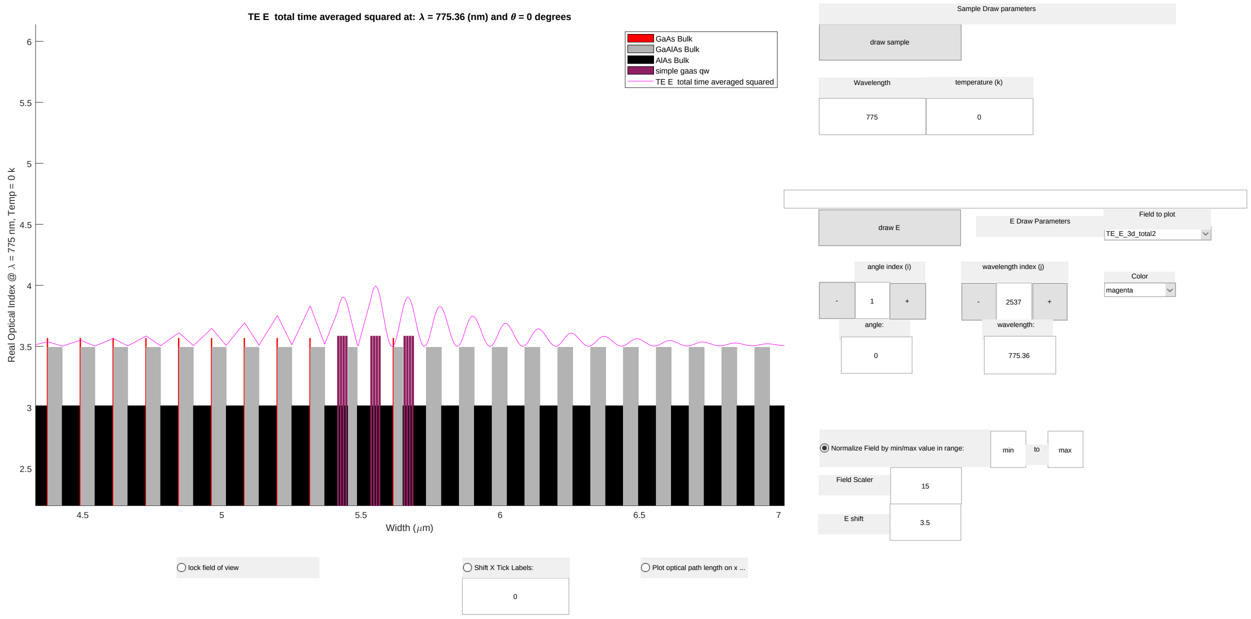

Wavelength: 775
temperature (k): 0
Field to plot: TE_E_3d_total2
angle index (i): 1
wavelength index (j): 2537
Color: magenta
angle: 0
wavelength: 775.36
● Normalize Field by min/max value in range: min to max
Field Scaler: 15
E shift: 3.5
○ Shift X Tick Labels: 0

Figure 81: The E Field GUI, with an electric field plotted on top of a sample.

option must be turned on inside of Reflectivity GUI. The electric fields are not explicitly calculated unless this option is enabled, so there are no fields to plot if it is turned off. When that feature is enabled, the user will be prompted for which layers of the material they wish to integrate the electric field in. This was originally implemented to measure how much energy was present inside of the quantum wells. The integral is performed numerically using a midpoint point rectangular Riemann sum. By default, the TMM scheme only calculates the fields at the boundaries of the sample. In order to improve the accuracy of the integration, the software will break the integration layers into a series of thinner layers (preserving total thickness). The number of thinner layers is given by the user's input. Typically twenty works well enough. This may be utilized to force the simulation to plot the electric field with a higher density of points.

In the example image the E field integration has been turned on in the top half of the sample, but not the bottom half of the sample. We see that the E field looks much smoother in the top half of the sample. If one is looking to produce a figure of the E field, turn on E field integration inside every layer of the sample and split the layers into twenty smaller layers.

The displayed E field may also be shifted and scaled by the two inputs in the bottom right corner of the GUI. The wavelength and angle of the simulation may be chosen from the index values. These values may only cover wavelengths and angles which were simulated by Reflectivity GUI.

Additionally, many other fields may be selected from the dropdown menu. The 3d part of the name is just a book keeping trick and does not hold any physical meaning.

The results of the field integration are displayed in Reflectivity GUI, as an additional line plot. The fields at each location in the sample are plotted here in the E Field GUI.

B.4 Simulfit GUI

The Simulfit GUI is a convenient way to link multiple instances of the Reflectivity GUI. This can be done for several reasons, the most prominent is that it allows all the fits and data to be saved into a single unified file, which when opened, gives the user access to all the simulations as they were when the characterization procedure was performed. Additionally, the Simulfit GUI may be used to impose constraints across multiple simulations, such as functional form for the thickness dependence of the exciton energy inside the quantum wells.

An example of the Simulfit GUI is shown in Figure 82. The nine linked Reflectivity GUIs are shown as the other tabs in the MATLAB figure window. In order to set up the Simulfit GUI, first all the desired instances of the Reflectivity GUI should be opened. Then, the Link All Existing Reflectivity GUIs button should be pressed. All the Reflectivity GUIs will be renamed to have a number added to their name. It is most convenient to dock all the figures into a single MATLAB window at this point. The save button in the top left corner of the Simulfit GUI will save an instance of all these GUIs together at any time.

Next, the Set Tuning Variables button should be pressed. This will prompt the user to define which variables should be tuned in all the Reflectivity GUIs. Variables may be stipulated as local or global. Global variables will be forced to have the same value in all linked Reflectivity GUIs. This creates the convenient Tuning Variables table allowing the user to compare the values of the parameters at the various locations on the sample.

The Minimize Global Error button may be pressed to begin the optimization process in all linked Reflectivity GUIs together. However, in my experimentation, I did not find this feature to be very helpful. Although I searched for one, I was not able to conclusively determine a trend among the various fitting parameters at the different locations on the sample. As the key feature which changes with location on the sample is thickness, I searched for functional forms related to the thickness of the various layers. For example, I hoped to find a relationship between exciton energy and thickness of the GaAs quantum wells. The thickness of the quantum wells was determine from the stopband fitting, and the assumption that all layers of the sample are scaling uniformly. However, the lack of a trend seems to indicate this is a bad assumption.

However, the Simulfit GUI is still a convenient way to store the simulations in one unified file, and the option is here to continue this work. Inside of the Sample Characterization folder in the bay 4 server computer where I've stored copies of the characterizations of all the samples, there are Simulfit GUIs with the associated Reflectivity GUIs.

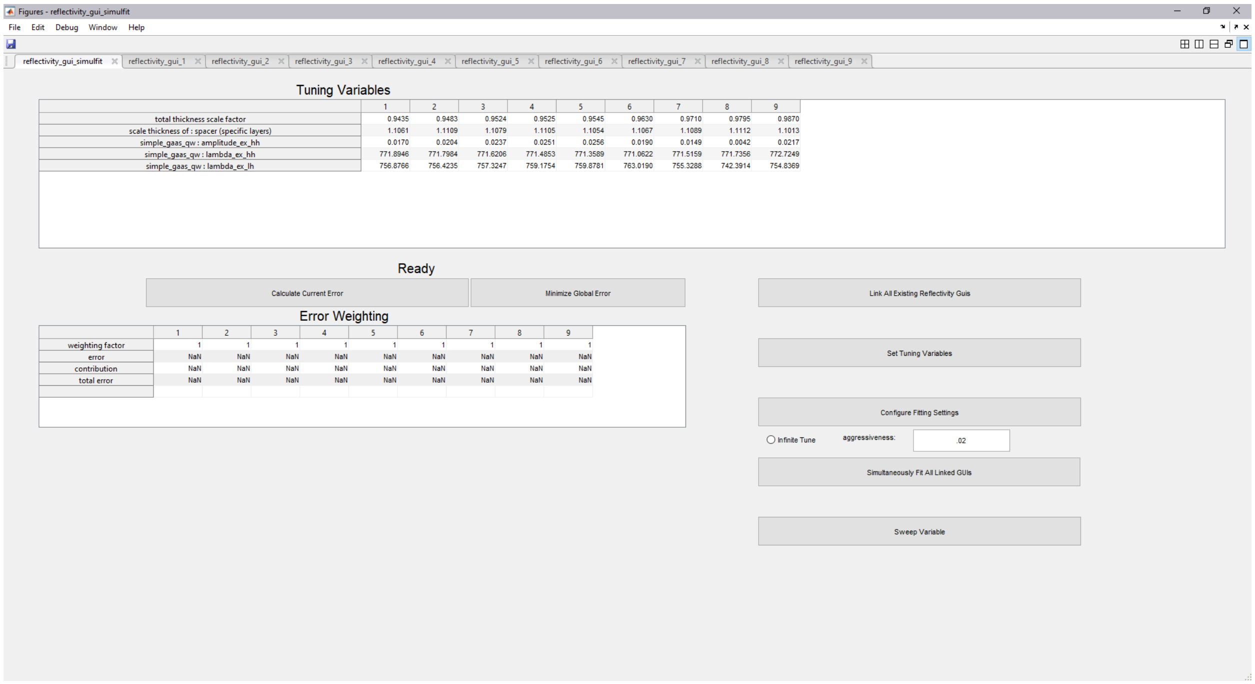

Tuning Variables

	1	2	3	4	5	6	7	8	9
total thickness scale factor	0.9435	0.9483	0.9524	0.9525	0.9545	0.9630	0.9710	0.9795	0.9870
scale thickness of : spacer (specific layers)	1.1061	1.1109	1.1079	1.1105	1.1054	1.1067	1.1089	1.1112	1.1013
simple_gaas_qw : amplitude_ex_hh	0.0170	0.0204	0.0237	0.0251	0.0256	0.0190	0.0149	0.0042	0.0217
simple_gaas_qw : lambda_ex_hh	771.8946	771.7984	771.6206	771.4853	771.3589	771.0622	771.5159	771.7356	772.7249
simple_gaas_qw : lambda_ex_lh	756.8766	756.4235	757.3247	759.1754	759.8781	763.0190	755.3288	742.3914	754.8369

Ready

Error Weighting

	1	2	3	4	5	6	7	8	9
weighting factor	1	1	1	1	1	1	1	1	1
error	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN
contribution	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN
total error	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN

○ Infinite Tune

aggressiveness: .02

Figure 82: An example of the Simulfit GUI.

B.5 Example: Fitting Stopband

The processing of fitting the broad reflectivity as discussed in section 4.5.1 is well covered in the posted videos online [11]. However, I will also include a brief outline of the steps taken here.

This process will make use of the Reflectivity GUI and experimentally measured reflectivity data as outlined in section 2.5, as well as a sample structure file created by the Sample Creator GUI based on the MBE growth sheet provided by the grower.

- Begin by opening Reflectivity GUI, and loading in the sample structure file created by the Sample Creator GUI using the Load Sample Structure button.
- Load in the experimental reflectivity data using the Load Exp Fig button.
- Create index scale factors for imaginary parts of the index of the AlAs and AlGaAs materials. Set these to 0.001 to begin.
- Run the simulation. At this point the simulation is probably producing a flat stopband, but at a higher or lower wavelengths than the data. Adjust the Total Thickness Scale Factor until the stopband is approximately centered on the data.
- The width of the stopband is determine by the contrast of the index of the DBR materials. Likely, the measured stopband is a bit wider or narrower than the simulation. Create index scale factors for the real parts of the index of AlAs and AlGaAs. Keep these two numbers multiplicative inverses of one another, and adjust until the width of the stopband matches the data. The GUI can handle inputs with division in them. For example, the scale factor for AlAs may be 1.01 and the scale factor for AlGaAs may be 1/1.01. As this ratio is adjusted, the total thickness scale factor may have to be adjusted too.
- Once the simulation is approximately lined up with the experimental data, additional data may be added to the simulation by pressing the Chain GUI button to open additional linked instances of Reflectivity GUI. Typically I use three total instances. The first is for fitting data at 0 degrees, the second for 15 degrees, and the third for 30 degrees.
- Create thickness scale factors for the AlAs and AlGaAs materials.

- The GUI is now prepared to perform a fit through error minimization. Press the Optimize Reflectivity Or PL button to begin. The GUI presents the option to fit either the reflectivity or the integration curve to the data. In this case we use the reflectivity option. There are two different error metrics to choose from. The regular and the first derivative. The first derivative can be useful sometimes as a first fit, to be followed up by the regular metric afterwards. In this case, we use the regular error metric right away. Finally, the GUI asks which parameters to use as fitting parameters. In this case, we select the index scale factors for the real parts of the index of AlAs and AlGaAs, as well as the thickness of those two materials. This is a total of four fitting parameters.
- While the error minimization is running (or before), the user may turn on the infinite tune option if desired. Once the user is satisfied with the quality of the fit, it may be turned off.
- The fitting is now complete. The simulation should be fitting the experimental data well everywhere except around the 805 nm location, which as discussed in the text, seems to be an effect due to the substrate GaAs excitons.
- The GUI (and any chained GUIs) may be saved using the floppy disk icon in the top left portion of the window. This will save a file which, when double clicked inside the MATLAB environment will open the simulation as is.

B.6 Example: Fitting Polaritons

The processing of fitting the finer reflectivity to the polariton data as discussed in section 4.5.2 is well covered in the posted videos online [11]. However, I will also include a brief outline of the steps taken here.

This process will make use of the the extracted dispersion curve of the lower polariton $\lambda(\theta)$ as discussed in section 4.2 and the PLE measurement of the upper polariton λ_{PLE}. Additionally, it begins with the fitted simulation from section B.5.

- Use the `make_spikes` function in MATLAB to create a spike plot showing the upper and lower polariton at normal incidence. This is a function I've made which is included with the software packages. Load this figure into Reflectivity GUI to replace the broad reflectivity data the simulation has already been fit to.

- Use the Add Variable button to add the exciton parameter to the material parameters table. This includes the heavy-hole and light-hole exciton's energies, coupling strengths, and linewidths. As discussed in section 4.5.2, I imposed that both excitons have the same coupling strength (amplitude) and linewidth. This may be done using the "ratio" variables inside the quantum well's index function.
- Create a thickness scale factor for the spacer layers inside the microcavity. This is done by selecting "specific layers" from the dropdown menu; the user will be prompted to name this group. Typically I name this group the spacers.
- The four fitting parameters are now accessible in the GUI. The user should fit by hand the simulation so that the dips approximately line up with the spikes from the spike plot. The amplitude of the excitons correlates with the splitting between the upper and lower polariton. The energy of the heavy-hole excitons will shift the more exciton branch and the thickness scale factor will shift the more photonic branch.
- Do not move on until the dips of the simulation are approximately matching the location of the spikes in the spike plot.
- Use the Input 2d PL Data button to load in the lower polariton dispersion curve and the upper polariton PLE data point. The GUI will guide the user through this process.
- Use the Configure 2d PL Optimizer button to configure the fitting procedure to extract the upper and lower polariton curves from the reflectivity simulation. The GUI will guide the user through this process.
- Right click on the Optimize 2d PL button to show the current fit. This will open up the two dimensional fitting window. Keep this window open for the optimization procedure.
- The GUI is now ready to perform the error minimization procedure to tune the fitting parameters. Left click on the Optimize 2d PL button. The user must select which parameters are tunable. Typically the fitting parameters are the spacer scale factor as well as the amplitude and energy of the two excitons. However, as mentioned the amplitudes are normally pinned to be equal, giving a total of four fitting parameters.

- While the error minimization is running (or before), the user may turn on the infinite tune option if desired. Once the user is satisfied with the quality of the fit, it may be turned off.
- Now that the simulation has been fit to the polariton data, the simulation should be saved using the floppy disk icon at the top of the window.
- The exciton fraction may now be extracted. This is done using the Generate Angular Curves portion of the GUI. Click on the Add Curve button twice. By default the two curves are set to be simulated from -30 degrees to 30 degrees. However, we only need this to be performed at normal incidence, so change the angle input to be just zero. Additionally, adjust the wavelengths range to cover the lower polariton range.
- Clicking on the zero inside the Fit Model column will bring up an instance of Spec Fit GUI. The user must define a fitting function, and then press the UI Resume button. Repeat this process for the second curve.
- Using the Add Variable button near the Material Parameters table, add the exciton energy derivative variables for the quantum wells. These are not variables to be fit, but rather variables that allow the GUI to easily perform the numerical derivative described in section 4.5.2.
- Scroll the table to the right and click on the zeros under the Calc Hop column. This will prompt the user for which parameter to take derivatives with respect to. Select the two derivative parameters. The GUI will prompt the user for units, select eV.
- It is now a good time to once again save this instance of the GUI using the floppy disk icon at the top of the window.
- Click the Generate Hop button. The GUI will go through a procedure to calculate the numerical derivatives. The final output will be figure windows from which the user may copy the heavy-hole exciton fraction and light-hole exciton fraction.
- This process may be repeated at every location on the sample to generate the data in figure 53A. This data may then be fit with equation 4.29 using Spec Fit GUI, which has it programmed in under the name Exciton Fraction Approximation.

Bibliography

[1] Sadao Adachi. GaAs, AlAs, and AlxGa1-xAs: Material parameters for use in research and device applications. *Journal of Applied Physics*, 58(3):R1–R29, 08 1985.

[2] Sadao Adachi. *Physical Properties of III-V Semiconductor Compounts*. John Wiley & Sons, Ltd, 1992.

[3] Hassan Alnatah, Qi Yao, Jonathan Beaumariage, Shouvik Mukherjee, Man Chun Tam, Zbigniew Wasilewski, Ken West, Kirk Baldwin, Loren N. Pfeiffer, and David W. Snoke. Coherence measurements of polaritons in thermal equilibrium reveal a power law for two-dimensional condensates. *Science Advances*, 10(18):eadk6960, 2024.

[4] Hassan Alnatah, Qi Yao, Qiaochu Wan, Jonathan Beaumariage, Ken West, Kirk Baldwin, Loren N. Pfeiffer, and David W. Snoke. Bose-einstein condensation of polaritons at room temperature in a gaas/algaas structure, 2024.

[5] A. Askitopoulos, L. Mouchliadis, I. Iorsh, G. Christmann, J. J. Baumberg, M. A. Kaliteevski, Z. Hatzopoulos, and P. G. Savvidis. Bragg polaritons: Strong coupling and amplification in an unfolded microcavity. *Phys. Rev. Lett.*, 106:076401, Feb 2011.

[6] D. E. Aspnes, S. M. Kelso, R. A. Logan, and R. Bhat. Optical properties of AlxGa1-x As. *Journal of Applied Physics*, 60(2):754–767, 07 1986.

[7] R. Balili, V. Hartwell, D. Snoke, L. Pfeiffer, and K. West. Bose-einstein condensation of microcavity polaritons in a trap. *Science*, 316(5827):1007–1010, 2007.

[8] R. B. Balili, D. W. Snoke, L. Pfeiffer, and K. West. Actively tuned and spatially trapped polaritons. *Applied Physics Letters*, 88(3):031110, 01 2006.

[9] María Barra-Burillo, Unai Muniain, Sara Catalano, Marta Autore, Fèlix Casanova, Luis E. Hueso, Javier Aizpurua, Ruben Esteban, and Rainer Hillenbrand. Microcavity phonon polaritons from the weak to the ultrastrong phonon–photon coupling regime. *Nature Communications*, 12(1):6206, Oct 2021.

[10] Gerald Bastard. *Wave Mechanics Applied to Semiconductor Heterostructures*. Wiley-Interscience, 1991.

[11] Jonathan Beaumariage. Jonathan beaumariage youtube channel. `https://www.youtube.com/@jonathanbeaumariage6480`, 2022. [Online; accessed May 02, 2022].

[12] Jonathan Beaumariage. Tmm reflectivity simulator for dbr microcavity polaritons. `https://www.mathworks.com/matlabcentral/fileexchange/111005`, 2022. [Online; accessed May 02, 2022].

[13] Jonathan Beaumariage, Zheng Sun, Hassan Alnatah, Qi Yao, David M. Myers, Mark Steger, Ken West, Kirk Baldwin, Loren N. Pfeiffer, Man Chun Alan Tam, Zbig R. Wailewski, and David W. Snoke. Measurement of exciton fraction of microcavity exciton-polaritons using transfer-matrix modeling, 2024.

[14] P A Belov and E S Khramtsov. The binding energy of excitons in narrow quantum wells. *Journal of Physics: Conference Series*, 816(1):012018, mar 2017.

[15] Michael E. Berkowitz, Brian S. Y. Kim, Guangxin Ni, Alexander S. McLeod, Chiu Fan Bowen Lo, Zhiyuan Sun, Genda Gu, Kenji Watanabe, Takashi Taniguchi, Andrew J. Millis, James C. Hone, Michael M. Fogler, Richard D. Averitt, and D. N. Basov. Hyperbolic cooper-pair polaritons in planar graphene/cuprate plasmonic cavities. *Nano Letters*, 21(1):308–316, Jan 2021.

[16] James R. Chelikowsky and Marvin L. Cohen. Nonlocal pseudopotential calculations for the electronic structure of eleven diamond and zinc-blende semiconductors. *Phys. Rev. B*, 14:556–582, Jul 1976.

[17] Xingzhou Chen, Hassan Alnatah, Danqun Mao, Mengyao Xu, Yuening Fan, Qiaochu Wan, Jonathan Beaumariage, Wei Xie, Hongxing Xu, Zhe-Yu Shi, David Snoke, Zheng Sun, and Jian Wu. Bose condensation of upper-branch exciton-polaritons in a transferable microcavity. *Nano Letters*, 23(20):9538–9546, 2023. PMID: 37818838.

[18] Xingzhou Chen, Hassan Alnatah, Danqun Mao, Mengyao Xu, Yuening Fan, Qiaochu Wan, Jonathan Beaumariage, Wei Xie, Hongxing Xu, Zhe-Yu Shi, David Snoke, Zheng Sun, and Jian Wu. Bose condensation of upper-branch exciton-polaritons in a transferable microcavity. *Nano Letters*, 23(20):9538–9546, 2023. PMID: 37818838.

[19] Shun Lien Chuang. *Physics of Photonic Devices*. Wiley Series in Pure and Applied Optics. Wiley, 2009.

[20] Luis V. Rodríguez de Marcos, Juan I. Larruquert, José A. Méndez, and José A. Aznárez. Self-consistent optical constants of sio2 and ta2o5 films. *Opt. Mater. Express*, 6(11):3622–3637, Nov 2016.

[21] Yago del Valle-Inclan Redondo, Helgi Sigurdsson, Hamid Ohadi, Ivan A. Shelykh, Yuri G. Rubo, Zacharias Hatzopoulos, Pavlos G. Savvidis, and Jeremy J. Baumberg. Observation of inversion, hysteresis, and collapse of spin in optically trapped polariton condensates. *Phys. Rev. B*, 99:165311, Apr 2019.

[22] Hui Deng, Hartmut Haug, and Yoshihisa Yamamoto. Exciton-polariton bose-einstein condensation. *Rev. Mod. Phys.*, 82:1489–1537, May 2010.

[23] Hui Deng, David Press, Stephan Götzinger, Glenn S. Solomon, Rudolf Hey, Klaus H. Ploog, and Yoshihisa Yamamoto. Quantum degenerate exciton-polaritons in thermal equilibrium. *Phys. Rev. Lett.*, 97:146402, Oct 2006.

[24] Hui Deng, Gregor Weihs, David Snoke, Jacqueline Bloch, and Yoshihisa Yamamoto. Polariton lasing vs. photon lasing in a semiconductor microcavity. *Proceedings of the National Academy of Sciences*, 100(26):15318–15323, 2003.

[25] Dmitriy Dovzhenko, Denis Aristov, Lucinda Pickup, Helgi Sigurðsson, and Pavlos Lagoudakis. Next-nearest-neighbor coupling with spinor polariton condensates. *Phys. Rev. B*, 108:L161301, Oct 2023.

[26] M. A. Dupertuis, M. Proctor, and B. Acklin. Generalization of complex snell–descartes and fresnel laws. *J. Opt. Soc. Am. A*, 11(3):1159–1166, Mar 1994.

[27] M. V. Gurudev Dutt. *Coherent Optical Manipulation of Electron Spin In Charged Semiconductor Quantum Dots*. PhD thesis, The University of Michigan, 2004.

[28] F. Dybała, M. P. Polak, J. Kopaczek, P. Scharoch, K. Wu, S. Tongay, and R. Kudrawiec. Pressure coefficients for direct optical transitions in mos2, mose2, ws2, and wse2 crystals and semiconductor to metal transitions. *Scientific Reports*, 6(1):26663, May 2016.

[29] Arthur Edelstein, Nenad Amodaj, Karl Hoover, Ron Vale, and Nico Stuurman. Computer control of microscopes using µmanager. *Current Protocols in Molecular Biology*, 92(1):14.20.1–14.20.17, 2010.

[30] Yuening Fan, Qiaochu Wan, Qi Yao, Xingzhou Chen, Yuanjun Guan, Hassan Alnatah, Daniel Vaz, Jonathan Beaumariage, Kenji Watanabe, Takashi Taniguchi, Jian Wu, Zheng Sun, and David Snoke. High efficiency of exciton-polariton lasing in a 2d multilayer structure. *ACS Photonics*, 11(7):2722–2728, Jul 2024.

[31] R.F.C. Farrow. Recent developments in molecular beam epitaxy. *Journal of Crystal Growth*, 104(2):556–577, 1990.

[32] R. E. Fern and A. Onton. Refractive Index of AlAs. *Journal of Applied Physics*, 42(9):3499–3500, 12 2003.

[33] Karen G. Fleming. Fluorescence theory. In John C. Lindon, editor, *Encyclopedia of Spectroscopy and Spectrometry (Second Edition)*, pages 628–634. Academic Press, Oxford, second edition edition, 2010.

[34] Alexander Franzen. Componentlibrary. `https://www.gwoptics.org/ComponentLibrary/`, 2006. This work is licensed under the Creative Commons Attribution-NonCommercial 3.0 Unported License. To view a copy of this license, visit `https://creativecommons.org/licenses/by-nc/3.0/`.

[35] S. Gehrsitz, F. K. Reinhart, C. Gourgon, N. Herres, A. Vonlanthen, and H. Sigg. The refractive index of AlxGa1-xAs below the band gap: Accurate determination and empirical modeling. *Journal of Applied Physics*, 87(11):7825–7837, 06 2000.

[36] Rudresh Ghosh, Joon-Seok Kim, Anupam Roy, Harry Chou, Mary Vu, Sanjay K. Banerjee, and Deji Akinwande. Large area chemical vapor deposition growth of monolayer mose2 and its controlled sulfurization to mos2. *Journal of Materials Research*, 31(7):917–922, Apr 2016.

[37] Juan A. Gonzalez-Cuevas, Tamer F. Refaat, M. Nurul Abedin, and Hani E. Elsayed-Ali. Calculations of the temperature and alloy composition effects on the optical properties of AlxGa1-xAsySb1-y and GaxIn1-xAsySb1-y in the spectral range 0.5–6 eV. *Journal of Applied Physics*, 102(1), 07 2007. 014504.

[38] Manuel Guizar. Efficient subpixel image registration by cross-correlation. `https://www.mathworks.com/matlabcentral/fileexchange/18401-efficient-subpixel-image-registration-by-cross-correlation`, 2016. [Online; accessed May 02, 2022].

[39] Manuel Guizar-Sicairos, Samuel T. Thurman, and James R. Fienup. Efficient subpixel image registration algorithms. *Opt. Lett.*, 33(2):156–158, Jan 2008.

[40] Bertrand I. Halperin. On the hohenberg–mermin–wagner theorem and its limitations. *Journal of Statistical Physics*, 175(3):521–529, May 2019.

[41] Eugene Hecht. *Optics*. Pearson, fifth edition, 2015.

[42] John David Jackson. *Classical Electrodynamics*. Wiley, 1998.

[43] Peng Ji, R. Balili, J. Beaumariage, S. Mukherjee, D. Snoke, and M. V. Gurudev Dutt. Multiple-photon excitation of nitrogen vacancy centers in diamond. *Phys. Rev. B*, 97:134112, Apr 2018.

[44] Saju Joseph, Jainy Mohan, Seetha Lakshmy, Simil Thomas, Brahmananda Chakraborty, Sabu Thomas, and Nandakumar Kalarikkal. A review of the synthesis, properties, and applications of 2d transition metal dichalcogenides and their heterostructures. *Materials Chemistry and Physics*, 297:127332, 2023.

[45] B A Joyce. Molecular beam epitaxy. *Reports on Progress in Physics*, 48(12):1637, dec 1985.

[46] Alexey V. Kavokin, Jeremy J. Baumberg, Guillaume Malpuech, and Fabrice P. Laussy. *Microcavities*. Oxford Science Publications, 2007.

[47] Alexey V. Kavokin, Jeremy J. Baumberg, Guillaume Malpuech, and Fabrice P. Laussy. *Microcavities*, chapter 1.1.1 Q-factor and finesse. Oxford Science Publications, 2007.

[48] Alexey V. Kavokin, Jeremy J. Baumberg, Guillaume Malpuech, and Fabrice P. Laussy. *Microcavities*, chapter 4.4.5 Motional Narrowing. Oxford Science Publications, 2007.

[49] Charles Kittel. *Introduction to Solid State Physics*. Wiley, 2004.

[50] Kenneth S. Krane. *Modern Physics*. John Wiley & Sons Inc, 2012.

[51] R. B. Laughlin and David Pines. The theory of everything. *Proceedings of the National Academy of Sciences*, 97(1):28–31, 2000.

[52] Kevin Luke, Yoshitomo Okawachi, Michael R. E. Lamont, Alexander L. Gaeta, and Michal Lipson. Broadband mid-infrared frequency comb generation in a si3n4 microresonator. *Opt. Lett.*, 40(21):4823–4826, Nov 2015.

[53] Xuekai Ma, Bernd Berger, Marc Aßmann, Rodislav Driben, Torsten Meier, Christian Schneider, Sven Höfling, and Stefan Schumacher. Realization of all-optical vortex switching in exciton-polariton condensates. *Nature Communications*, 11(1):897, February 2020.

[54] F. Manni, K. G. Lagoudakis, T. C. H. Liew, R. André, and B. Deveaud-Plédran. Spontaneous pattern formation in a polariton condensate. *Phys. Rev. Lett.*, 107:106401, Sep 2011.

[55] M. V. Marquezini, J. Tignon, T. Hasche, and D. S. Chemla. Refractive index and absorption of GaAs quantum wells across excitonic resonances. *Applied Physics Letters*, 73(16):2313–2315, 10 1998.

[56] S. Mukherjee, V. K. Kozin, A. V. Nalitov, I. A. Shelykh, Z. Sun, D. M. Myers, B. Ozden, J. Beaumariage, M. Steger, L. N. Pfeiffer, K. West, and D. W. Snoke. Dynamics of spin polarization in tilted polariton rings. *Phys. Rev. B*, 103:165306, Apr 2021.

[57] S. Mukherjee, D. M. Myers, R. G. Lena, B. Ozden, J. Beaumariage, Z. Sun, M. Steger, L. N. Pfeiffer, K. West, A. J. Daley, and D. W. Snoke. Observation of nonequilibrium motion and equilibration in polariton rings. *Phys. Rev. B*, 100:245304, Dec 2019.

[58] D. M. Myers, S. Mukherjee, J. Beaumariage, D. W. Snoke, M. Steger, L. N. Pfeiffer, and K. West. Polariton-enhanced exciton transport. *Phys. Rev. B*, 98:235302, Dec 2018.

[59] D. M. Myers, J. K. Wuenschell, B. Ozden, J. Beaumariage, D. W. Snoke, L. Pfeiffer, and K. West. Edge trapping of exciton-polariton condensates in etched pillars. *Applied Physics Letters*, 110(21):211104, 2017.

[60] V. Negoita, D. W. Snoke, and K. Eberl. Stretching quantum wells: A method for trapping free carriers in GaAs heterostructures. *Applied Physics Letters*, 75(14):2059–2061, 10 1999.

[61] M. Ochiai, G. E. Giudice, H. Temkin, J. W. Scott, and T. M. Cockerill. Kinetics of thermal oxidation of AlAs in water vapor. *Applied Physics Letters*, 68(14):1898–1900, 04 1996.

[62] Shunji Ozaki and Sadao Adachi. Spectroscopic ellipsometry and thermoreflectance of GaAs. *Journal of Applied Physics*, 78(5):3380–3386, 09 1995.

[63] Konstantinos Papatryfonos, Todora Angelova, Antoine Brimont, Barry Reid, Stefan Guldin, Peter Raymond Smith, Mingchu Tang, Keshuang Li, Alwyn J. Seeds, Huiyun Liu, and David R. Selviah. Refractive indices of MBE-grown AlxGa(1-x)As ternary alloys in the transparent wavelength region. *AIP Advances*, 11(2), 02 2021. 025327.

[64] R. K. Pathria and Paul D. Beale. *Statistical Mechanics*. Academic Press, 2011.

[65] Maciej Pieczarka, Eliezer Estrecho, Sanjib Ghosh, Matthias Wurdack, Mark Steger, David W. Snoke, Kenneth West, Loren N. Pfeiffer, Timothy C. H. Liew, Andrew G. Truscott, and Elena A. Ostrovskaya. Topological phase transition in an all-optical exciton-polariton lattice. *Optica*, 8(8):1084–1091, Aug 2021.

[66] Mikhail N. Polyanskiy. Refractiveindex.info database of optical constants. *Scientific Data*, 11(1):94, Jan 2024.

[67] Aleksandar D. Rakić and Marian L. Majewski. Modeling the optical dielectric function of GaAs and AlAs: Extension of Adachi's model. *Journal of Applied Physics*, 80(10):5909–5914, 11 1996.

[68] M. Saba, C. Ciuti, J. Bloch, V. Thierry-Mieg, R. André, Le Si Dang, S. Kundermann, A. Mura, G. Bongiovanni, J. L. Staehli, and B. Deveaud. High-temperature ultrafast polariton parametric amplification in semiconductor microcavities. *Nature*, 414(6865):731–735, Dec 2001.

[69] D Sanvitto, F M Marchetti, M H Szymańska, G Tosi, M Baudisch, F P Laussy, D N Krizhanovskii, M S Skolnick, L Marrucci, A Lemaître, J Bloch, C Tejedor, and L Viña. Persistent currents and quantized vortices in a polariton superfluid. *Nature Physics*, 6(7):527–533, July 2010.

[70] P. G. Savvidis, J. J. Baumberg, R. M. Stevenson, M. S. Skolnick, D. M. Whittaker, and J. S. Roberts. Angle-resonant stimulated polariton amplifier. *Phys. Rev. Lett.*, 84:1547–1550, Feb 2000.

[71] Daniel V. Schroeder. *An Introduction to Thermal Physics*. Addison Wesley Longman, 2021.

[72] D. W. Snoke, V. Hartwell, J. Beaumariage, S. Mukherjee, Y. Yoon, D. M. Myers, M. Steger, Z. Sun, K. A. Nelson, and L. N. Pfeiffer. Reanalysis of experimental determinations of polariton-polariton interactions in microcavities. *Phys. Rev. B*, 107:165302, Apr 2023.

[73] David W. Snoke. *Solid State Physics*. Cambridge University Press, 2020.

[74] David W. Snoke. *Solid State Physics*, chapter 11 SpontaneousCoherence in Matter. Cambridge University Press, 2020.

[75] David W. Snoke. *Solid State Physics*, chapter 7 The Complex Susceptibility. Cambridge University Press, 2020.

[76] Mark Steger, Chitra Gautham, David W. Snoke, Loren Pfeiffer, and Ken West. Slow reflection and two-photon generation of microcavity exciton–polaritons. *Optica*, 2(1):1–5, Jan 2015.

[77] Yongbao Sun, Patrick Wen, Yoseob Yoon, Gangqiang Liu, Mark Steger, Loren N. Pfeiffer, Ken West, David W. Snoke, and Keith A. Nelson. Bose-einstein condensation of long-lifetime polaritons in thermal equilibrium. *Phys. Rev. Lett.*, 118:016602, Jan 2017.

[78] Yongbao Sun, Yoseob Yoon, Mark Steger, Gangqiang Liu, Loren N. Pfeiffer, Ken West, David W. Snoke, and Keith A. Nelson. Direct measurement of polariton–polariton interaction strength. *Nature Physics*, 13(9):870–875, 2017.

[79] Zheng Sun, Jonathan Beaumariage, Qingrui Cao, Kenji Watanabe, Takashi Taniguchi, Benjamin Matthew Hunt, and David Snoke. Observation of the interlayer exciton gases in wse2 -p:wse2 heterostructures. *ACS Photonics*, 7(7):1622–1627, 2020.

[80] Zheng Sun, Jonathan Beaumariage, Hema C.P. Movva, Sayema Chowdhury, Anupam Roy, Sanjay K. Banerjee, and David W. Snoke. Stress-induced bandgap renormalization in atomic crystals. *Solid State Communications*, 288:18–21, 2019.

[81] Zheng Sun, Jonathan Beaumariage, Qiaochu Wan, Hassan Alnatah, Nicholas Hougland, Jessica Chisholm, Qingrui Cao, Kenji Watanabe, Takashi Taniguchi, Benjamin Matthew Hunt, Igor V. Bondarev, and David Snoke. Charged bosons made of fermions in bilayer structures with strong metallic screening. *Nano Letters*, 21(18):7669–7675, 2021. PMID: 34516139.

[82] Zheng Sun, Jonathan Beaumariage, Ke Xu, Jierui Liang, Shaocong Hou, Stephen R. Forrest, Susan K. Fullerton-Shirey, and David W. Snoke. Electric-field-induced optical hysteresis in single-layer wse2. *Applied Physics Letters*, 115(16):161103, 2019.

[83] Zheng Sun, Ke Xu, Chang Liu, Jonathan Beaumariage, Jierui Liang, Susan K. Fullerton-Shirey, Zhe-Yu Shi, Jian Wu, and David Snoke. Photoluminescence switching effect in a two-dimensional atomic crystal. *ACS Nano*, 0(0):null, 0. PMID: 34878266.

[84] A. I. Tartakovskii, M. Emam-Ismail, R. M. Stevenson, M. S. Skolnick, V. N. Astratov, D. M. Whittaker, J. J. Baumberg, and J. S. Roberts. Relaxation bottleneck and its suppression in semiconductor microcavities. *Phys. Rev. B*, 62:R2283–R2286, Jul 2000.

[85] Gwo-Ching Wang and Toh-Ming Lu. *RHEED Transmission Mode and Pole Figures*. Springer, 2014.

[86] Zefang Wang, Daniel A. Rhodes, Kenji Watanabe, Takashi Taniguchi, James C. Hone, Jie Shan, and Kin Fai Mak. Evidence of high-temperature exciton condensation in two-dimensional atomic double layers. *Nature*, 574(7776):76–80, Oct 2019.

[87] A M White, E W Williams, P Porteous, and C Hilsum. Applications of photoluminescence excitation spectroscopy to the study of indium gallium phosphide alloys. *Journal of Physics D: Applied Physics*, 3(9):1322, sep 1970.

[88] Q. Yao, E. Sedov, S. Mukherjee, J. Beaumariage, B. Ozden, K. West, L. Pfeiffer, A. Kavokin, and D. W. Snoke. Ballistic transport of a polariton ring condensate with spin precession. *Phys. Rev. B*, 106:245309, Dec 2022.

[89] Hui Zhang, Yaoguang Ma, Yi Wan, Xin Rong, Ziang Xie, Wei Wang, and Lun Dai. Measuring the refractive index of highly crystalline monolayer mos2 with high confidence. *Scientific Reports*, 5(1):8440, Feb 2015.

[90] Junxi Zhang, Lide Zhang, and Wei Xu. Surface plasmon polaritons: physics and applications. *Journal of Physics D: Applied Physics*, 45(11):113001, feb 2012.