

**ELECTRICAL CURRENT INJECTION INTO  
EXCITON-POLARITON CONDENSATES AND  
POLARITON-ENHANCED EXCITON TRANSPORT**

by

**David M. Myers**

B.A. in Physics, Covenant College, 2012

M.S. in Physics, University of Pittsburgh, 2014

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

2019

UNIVERSITY OF PITTSBURGH  
KENNETH P. DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

David M. Myers

It was defended on

March 22nd 2019

and approved by

David W. Snoke, Professor, Department of Physics and Astronomy

Robert P. Devaty, Associate Professor, Department of Physics and Astronomy

W. Vincent Liu, Professor, Department of Physics and Astronomy

Ayres Freitas, Associate Professor, Department of Physics and Astronomy

Elias Towe, Professor, Department of Electrical Engineering, Carnegie Mellon University

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

Copyright © by David M. Myers  
2019

# **ELECTRICAL CURRENT INJECTION INTO EXCITON-POLARITON CONDENSATES AND POLARITON-ENHANCED EXCITON TRANSPORT**

David M. Myers, PhD

University of Pittsburgh, 2019

Exciton-polaritons are composite quasiparticles that are the result of coupling excitons with photons. They bring together many of the advantages of both underlying parts: The excitons lend their strong interactions, and the photons give an extremely light mass. Polaritons exist in solid-state systems, which gives them a high potential for device applications, especially as their most interesting quantum effects, Bose-Einstein condensation in particular, become accessible at higher temperatures.

In this thesis, I focus primarily on the effects of injecting electrical current into a polariton condensate. I was able to take advantage of the well-established fabrication methods for GaAs-based systems, as well as the high-quality samples developed by my predecessors, which allow the formation of nearly equilibrium condensates over large areas. My colleagues and I then made devices that allow direct injection of electrical charge carriers through polariton condensates in an in-plane direction of the 2D system. I report here three major observations. The first is a superlinear increase in photocurrent, which we attribute to stimulation of an Auger-like scattering of excitons into both the condensate and free carrier states. The second is an increase in the condensate density with current injection, which we explain as stimulated pairing of injected free electrons with intrinsic free holes. The third effect is a shift in the in-plane momentum of polariton condensates with injected current, which is evidence of drag on the polaritons by the free electrons.

In addition to the work with electrical current injection, I also report on exciton transport distances in the polariton system. Specifically, very highly excitonic particles are found far

from the creation point, with transport distances comparable to those of the much more photonic polaritons. We explain this as an effect of the unusual shape of the polariton energy dispersion.

## TABLE OF CONTENTS

<b>PREFACE</b> . . . . .	xix
<b>1.0 INTRODUCTION</b> . . . . .	1
1.1 The exciton-polariton system . . . . .	1
1.1.1 Polaritons at equilibrium . . . . .	3
1.1.2 Advantages of polaritons . . . . .	4
1.1.3 Dissertation outline . . . . .	5
1.2 Cavity photons . . . . .	6
1.2.1 The cavity photon energy dispersion . . . . .	6
1.2.2 Free spectral range and microcavities . . . . .	11
1.2.3 $Q$ -factor: the cavity as an harmonic oscillator . . . . .	12
1.2.4 Cavity finesse . . . . .	15
1.2.5 Distributed Bragg reflectors . . . . .	16
1.2.6 DBR microcavities . . . . .	22
1.3 Excitons . . . . .	27
1.3.1 Excitons as hydrogen-like particles . . . . .	27
1.3.2 The electron band structure . . . . .	29
1.3.3 The Luttinger-Kohn model . . . . .	33
1.3.4 Excitons in quantum wells . . . . .	35
1.4 Polaritons . . . . .	40
1.4.1 Exciton-photon coupling . . . . .	40
1.4.2 The polariton eigenstates . . . . .	47
1.4.3 Detuning . . . . .	50

1.4.4	Polariton effective mass . . . . .	53
1.4.5	Rate equations and polariton lifetime . . . . .	56
1.4.6	Strong coupling . . . . .	60
1.4.7	Interactions . . . . .	61
1.4.8	Strain . . . . .	62
1.4.9	Bose-Einstein condensation . . . . .	63
<b>2.0</b>	<b>FABRICATION AND EXPERIMENTAL METHODS . . . . .</b>	<b>66</b>
2.1	Microcavity sample design . . . . .	66
2.1.1	The cavity wedge . . . . .	67
2.2	Fabrication methods . . . . .	70
2.2.1	Review of confinement and fabrication methods . . . . .	70
2.2.1.1	Advantages of etched pillars . . . . .	72
2.2.2	Etching . . . . .	73
2.2.3	Electrical contacts . . . . .	77
2.3	Experimental Methods . . . . .	78
2.3.1	Optical pumping . . . . .	78
2.3.1.1	Resonant pumping . . . . .	78
2.3.1.2	Nonresonant pumping . . . . .	80
2.3.2	Optics and imaging . . . . .	83
2.3.3	Spectroscopy and detection . . . . .	86
2.3.4	Electrical measurements . . . . .	88
<b>3.0</b>	<b>TRAPPING AND CONDENSATION IN ETCHED PILLARS . . . . .</b>	<b>91</b>
3.1	Motivation . . . . .	91
3.2	The experiment . . . . .	91
3.3	Results . . . . .	93
3.4	Effect on the polariton lifetime . . . . .	104
3.5	Strain as an explanation of the trapping . . . . .	104
3.6	Three-threshold behavior . . . . .	105
3.7	Conclusions . . . . .	108
<b>4.0</b>	<b>CURRENT INJECTION INTO A POLARITON CONDENSATE . . . . .</b>	<b>109</b>

4.1	Motivation . . . . .	109
4.2	The experiment . . . . .	110
4.3	Basic electrical characteristics: the device as a phototransistor . . . . .	112
4.4	Stimulated injection into a condensate . . . . .	116
4.5	Superlinear increase of photocurrent . . . . .	123
4.5.1	Variations in the experimental geometry . . . . .	125
4.5.2	Theoretical model . . . . .	129
4.6	Conclusions . . . . .	132
<b>5.0</b>	<b>MOMENTUM DRAG ON A POLARITON CONDENSATE . . . . .</b>	<b>135</b>
5.1	Motivation . . . . .	135
5.2	Experiment . . . . .	137
5.3	Results . . . . .	141
5.3.1	Drag on the polariton momentum . . . . .	141
5.3.2	Pumping on the high-energy end . . . . .	145
5.3.3	Current-voltage characteristics . . . . .	148
5.3.4	Maximum velocity . . . . .	148
5.4	Conclusions . . . . .	149
<b>6.0</b>	<b>LONG RANGE EXCITON TRANSPORT . . . . .</b>	<b>151</b>
6.1	Motivation . . . . .	151
6.2	Some opening considerations . . . . .	152
6.3	Getting particle counts from measurements . . . . .	154
6.4	Experiment . . . . .	156
6.5	Long-range transport . . . . .	159
6.5.1	Transport distance estimate . . . . .	163
6.6	Backscattering at high in-plane momentum . . . . .	165
6.7	Application to ring pump geometry . . . . .	167
6.8	Conclusions . . . . .	172
<b>7.0</b>	<b>FINAL REMARKS . . . . .</b>	<b>174</b>
	<b>APPENDIX A. CAVITY REFRACTIVE INDEX ANGLE DEPENDENCE</b>	<b>177</b>
	<b>APPENDIX B. THE TRANSFER MATRIX METHOD . . . . .</b>	<b>180</b>

<b>APPENDIX C. FABRICATION DETAILS</b>	185
C.1 Lithography for etching	185
C.2 Reactive ion etching	187
C.3 Lithography for contact deposition	188
C.4 Metal contact deposition	190
<b>APPENDIX D. OPTICAL SETUPS</b>	192
<b>APPENDIX E. ADDITIONAL INFORMATION: LONG-RANGE EXCI-</b>	
<b>TON TRANSPORT</b>	197
E.1 Determining the lower polariton detuning	197
E.2 Calculating polariton counts and error bounds	198
E.2.1 Determining the pump spot size	201
E.2.2 Additional data sets	202
<b>BIBLIOGRAPHY</b>	205

## LIST OF TABLES

E.1 Parameters for various data sets used in the exciton transport study . . . . .	204
--	-----

## LIST OF FIGURES

1.1	Optical cavity interference diagram . . . . .	7
1.2	Cavity photon energy dispersion . . . . .	9
1.3	Bragg interference diagram . . . . .	17
1.4	DBR reflectance as a function of energy . . . . .	19
1.5	Phase shift of a DBR as a function of energy . . . . .	21
1.6	DBR cavity reflectance as a function of energy . . . . .	25
1.7	Free spectral range for a DBR microcavity . . . . .	26
1.8	Outline sketch for understanding the GaAs band structure . . . . .	30
1.9	The $\mathbf{k} \cdot \mathbf{p}$ theory for two bands . . . . .	32
1.10	Energy dispersions of electron bands in GaAs quantum wells . . . . .	37
1.11	The effect of quantum well confinement on the exciton states . . . . .	38
1.12	Polariton energy, Hopfield coefficients, and lifetimes as functions of in-plane wavenumber . . . . .	52
1.13	The polariton mode anticrossing and the polariton masses . . . . .	54
2.1	The microcavity sample design . . . . .	68
2.2	The sample gradient . . . . .	69
2.3	Overview of the fabrication process . . . . .	75
2.4	Comparison of resonant and nonresonant pumping . . . . .	81
2.5	Basic optical setups . . . . .	84
2.6	Optical setups for both real-space and $k$ -space imaging and filtering . . . . .	85
2.7	Synchronization setup for optical and electrical measurements . . . . .	89
2.8	Diagram of the electrical measurement timing . . . . .	90

3.1	SEM, diagram, and measured energy profile of an etched pillar . . . . .	94
3.2	Energy-resolved polariton luminescence from an etched pillar . . . . .	96
3.3	Real space polariton luminescence from an etched pillar . . . . .	97
3.4	Energy- and angle-resolved polariton luminescence from an etched pillar . . .	99
3.5	Intensity as a function of pump power in an etched pillar . . . . .	100
3.6	Map of the threshold powers for various points near the edge of an etched pillar	102
3.7	Energy of the polariton luminescence as a function of pump power . . . . .	103
3.8	Three-threshold behaviour of square pillar devices . . . . .	106
4.1	Diagram, SEM, energy profile, and condensate luminescence of an etched pillar with electrical contacts . . . . .	111
4.2	Calculated band device band diagram . . . . .	113
4.3	Measured and modeled current-voltage characteristics . . . . .	115
4.4	Electrical current vs. pump power from Ebers-Moll Model . . . . .	115
4.5	Polariton luminescence intensity as a function of pump power with and without applied voltage . . . . .	116
4.6	Polariton luminescence intensity as a function of current and voltage . . . . .	118
4.7	Diagram and result for free electron injection model . . . . .	119
4.8	Polariton luminescence intensity at the pump spot as a function of pump power and current . . . . .	122
4.9	Polariton luminescence intensity and electrical current as a function of pump power . . . . .	124
4.10	Pump spot size dependence of the current increase . . . . .	126
4.11	Illustration of a wire device . . . . .	127
4.12	Polariton luminescence intensity and electrical current as a function of pump power in a wire device . . . . .	128
4.13	Example of numerical results of the model for current increase . . . . .	131
4.14	Modeling results for occupancy of the condensate and free carrier reservoir as functions of pump power . . . . .	133
5.1	Feynman diagram for electron interaction with a photon . . . . .	137
5.2	Diagram and SEM of a wire device and the experimental arrangement . . . . .	139

5.3	Photoluminescence intensity vs. energy and position in a wire device . . . . .	140
5.4	Photoluminescence intensity vs. energy and in-plane wavenumber for a wire device . . . . .	142
5.5	The shift in polariton $k_{\parallel}$ with injected electrical current . . . . .	143
5.6	Overview of the results for pumping on the high-energy end of a wire . . . . .	146
5.7	Current vs. voltage in a wire with different pump spot locations . . . . .	147
5.8	Average in-plane wavevector vs. polariton counts . . . . .	150
6.1	Diagram of the experimental setup . . . . .	157
6.2	Energy- and angle-resolved polariton number . . . . .	158
6.3	Polariton population as a function of in-plane wavenumber . . . . .	160
6.4	Population of various in-plane wavenumber ranges as a function of distance from the pump spot . . . . .	162
6.5	Calculated group velocity and transport distance as a function of in-plane wavenumber . . . . .	164
6.6	Backscattered population at very large angle . . . . .	166
6.7	Calculated DBR transmission as a function of energy and in-plane wavenumber	167
6.8	Spatial distributions of the excitons and polaritons applied to a ring-pump geometry . . . . .	170
6.9	Exciton to polariton number ratio in a ring-pump geometry as a function of ring diameter . . . . .	171
A.1	Emission angle dependence of the refractive index of a multi-layer cavity . . .	179
C.1	Diagram of lift-of resist process . . . . .	188
D.1	High magnification and long working distance optical setup . . . . .	194
D.2	High magnification and long working distance optical setup with large numerical aperture . . . . .	195
D.3	Optical setup with large numerical aperture microscope objective . . . . .	196
E.1	Measurement and fit of the pump spot spatial profile . . . . .	202
E.2	Additional data for polariton-enhanced exciton transport . . . . .	203

## LIST OF EQUATIONS

1.1	6
1.2	8
1.3	8
1.4	8
1.5	8
1.6	10
1.7	10
1.8	11
1.9	11
1.10	12
1.11	12
1.12	13
1.13	13
1.14	13
1.15	13
1.16	14
1.17	14
1.18	14
1.19	14
1.20	15
1.21	15
1.22	15

1.23 .....	15
1.24 .....	16
1.25 .....	16
1.26 .....	18
1.27 .....	27
1.28 .....	27
1.29 .....	27
1.30 .....	28
1.31 .....	29
1.32 .....	29
1.33 .....	31
1.34 .....	31
1.35 .....	31
1.36 .....	31
1.37 .....	31
1.38 .....	34
1.39 .....	34
1.40 .....	35
1.41 .....	35
1.42 .....	36
1.43 .....	39
1.44 .....	40
1.45 .....	41
1.46 .....	41
1.47 .....	41
1.48 .....	41
1.49 .....	41
1.50 .....	42
1.51 .....	42
1.52 .....	42

1.53	43
1.54	43
1.55	44
1.56	44
1.57	45
1.58	45
1.59	45
1.60	46
1.61	46
1.62	47
1.63	48
1.64	48
1.65	49
1.66	49
1.67	49
1.68	49
1.69	49
1.70	50
1.71	50
1.72	51
1.73	55
1.74	55
1.75	55
1.76	55
1.77	56
1.78	57
1.79	57
1.80	57
1.81	58
1.82	58

1.83 .....	59
1.84 .....	59
1.85 .....	59
1.86 .....	60
1.87 .....	62
1.88 .....	63
1.89 .....	64
1.90 .....	64
1.91 .....	64
1.92 .....	65
1.93 .....	65
2.1 .....	87
4.1 .....	114
4.2 .....	114
4.3 .....	120
4.4 .....	129
4.5 .....	129
4.6 .....	129
4.7 .....	131
4.8 .....	131
4.9 .....	131
6.1 .....	151
6.2 .....	151
6.3 .....	155
6.4 .....	155
6.5 .....	163
6.6 .....	164
6.7 .....	168
6.8 .....	168
6.9 .....	168

6.10	168
6.11	168
6.12	169
6.13	169
A.1	177
A.2	177
A.3	178
B.1a	180
B.1b	180
B.2	181
B.3	181
B.4	181
B.5	182
B.6	182
B.7	182
B.8	182
B.9	183
B.10	183
B.11	183
B.12	183
B.13	183
B.14	184
B.15	184
B.16	184
B.17	184
B.18	184
E.1	199
E.2	199
E.3	200

## PREFACE

First of all, I am thankful for my daughter, Nelleke. She has been a deep source of joy these last six months, and has helped me to find the motivation to push this work to completion. I dedicate this work to her.

I am incredibly thankful to my wife, Jo, for her constant support over the past seven years. They were difficult and demanding years, yet also enjoyable, and I will always look back on them with fondness. I would not have wanted to do it without her.

I would also especially like to thank my advisor, Prof. David Snoke. He is a great teacher and guide, who constantly pushed me to become a better scientist. He is a dependable advocate for his students, and always has their best interests in mind. He is incredibly skilled at the many aspects of scientific endeavors, from designing experiments and solving problems in the lab, to organizing and stimulating enlightening discussions and lively debates. I have learned so much from him over the years about how science, as a whole, is done.

I am greatly appreciative of the guidance I received from the senior members of the lab, Jeff Wuenschell, Mark Steger, Chitra Gautham, and Gangquiang Liu. They helped me find my footing in those early years when a project worthy of a dissertation seems so far off. In particular, I am thankful to Mark for his patient introduction to the methods of optics and spectroscopy, and to Jeff for the same in fabrication techniques. In addition, the excitement, optimism, and drive displayed by Ryan Balili during his several visiting periods to the lab were, and continue to be, a source of inspiration to me.

The work I discuss in this dissertation was possible due to the efforts of several people. My first successful experimental project was moved forward during the early stages in large part due to Jeff Wuenschell's help. My optics and programming skills were honed over countless hours of bouncing ideas back-and-forth with Jonathan Beaumariage. I am also indebted to

him for his work in the initial characterization of each of our samples, which is no simple task. His programming contributions to the lab were also invaluable, from our joint effort to automate and control everything with MATLAB, to his easy-to-use runs-everything-in-the-lab program. The majority of the devices I studied were fabricated by Burcu Ozden, and I would never have made the progress that I did if it were not for her dependable production of quality samples. I also learned much from her about how much better I should be at keeping up with the latest published research. I am thankful to Shouvik Mukherjee for a countless number of discussions, arguments, and debates about many scientific topics, including experimental methods and general physics theory. They helped me become a better and more knowledgeable scientist, especially about how much I don't know. And as for discussions, arguments, and debates about things not scientific, I am thankful to Jonny, Shouvik, and all the others in the lab, because, while they were perhaps detrimental to our personal scientific progress, they were certainly enjoyable and helpful in maintaining sanity.

I would like to thank Qi Yao for help with my latest experiments, and generally taking over my projects. It is good to know that at least some of the loose ends I am leaving will be picked up. I would also like to thank Zheng Sun for his encouragement over these last months, especially as I have been looking toward the next step.

I would like to thank the members of my committee, both past and present, including those listed above and Prof. David Pekker, for guiding me toward a completed PhD over the years. In particular, I am thankful to them for reading this dissertation and giving input, which is no quick task.

I am also thankful to many other members of the faculty and staff who helped me along this journey. To the teachers of my courses, for helping to build up my physics foundation. To the machinists, who are so much faster and better at making things than I am. To the staff in the office, for taking care of all of the many administrative things. And especially to Leyla Hirschfeld, for making sure I did all the things I needed to do when I needed to do them, and always being available and willing the answer all the little questions that came up over the years.

I am also indebted to my professors at Covenant College, especially Profs. Phill Broussard and Don Petcher. They are truly devoted to teaching their students the beauty of physics

and the physical world, and they gave me the foundation I needed to dive into the study of physics.

My family, both the part from birth and the part from marriage, has been an incredible source of support during this time. They remind me that there is a life outside of science, academia, the city, or any other bubble in which I find myself. They remind me that that life is not only out there: It is also wonderful! I am especially appreciative of my parents, for always providing a place for Jo and me to get away, if only for an evening.

To Matt and Bethany Mierski, and Autumn Matta, I am thankful for the friendship that was shared with both Jo and me. Likewise, I am thankful to our many friends at First Reformed, for welcoming us in and coming alongside us as a family.

And finally, I am thankful to God, who is eternally faithful, and ever-present in times of need. *Soli deo gloria.*

## 1.0 INTRODUCTION

### 1.1 THE EXCITON-POLARITON SYSTEM

Of the two classes of particles in nature, bosons are the more mysterious. Many of us have our first introduction to the other class, fermions, in high school chemistry courses, where we are told that only one electron can occupy each state in an atom. This Pauli exclusion principle, while at that point typically representing nothing more than a rule to be memorized, becomes normal and expected. When the rule is later applied to any general quantum state, not simply to the atomic orbitals, we nod our heads in recognition of a familiar axiom. We learn even later on that this is because particle wavefunctions must be either symmetric or antisymmetric, with the latter giving the exclusion principle. Bosons, with their fully symmetric wavefunctions, have no requirement to follow this comfortable rule. Instead, they can all pile into the same state, leading to very different behavior. This includes stimulated emission (or, more generally, stimulated scattering), Bose-Einstein condensation (BEC), and superfluidity.

The unique statistics of bosons were first described by Bose in 1924 [1], and expanded by Einstein to predict BEC in a gas of massive particles [2]. BEC was then used by London [3] to explain the first major realization of macroscopic bosonic phenomena: the superfluid phase transition in liquid  $^4\text{He}$  [4, 5]. However, requirements of low temperatures, high densities, and weak interactions caused a long delay before the first observation of BEC of a gas [6], which was seen in cold atoms in 1995 [7, 8].

Photons are probably the most familiar elementary bosons, and they exhibit the most generally familiar bosonic phenomena: stimulated emission. Lasers have brought this effect into our daily lives, with huge populations of photons moving with the same phase, energy,

and momentum. While stimulation into a particular quantum state (which I briefly discuss in Section 1.4.5), is a general feature of bosons, which arises whenever many bosons are already occupying a particular state, BEC and superfluidity require a bit more work. For each of these, the particles need mass and there needs to be some process by which the particles reach (or at least approach) equilibrium [9]. This presents a problem for photons, since they are massless and directly interact with each other only very weakly (such as through the creation of particle-antiparticle pairs in vacuum), and therefore cannot effectively exchange momentum and energy. Adding mass is easy, and, as I discuss in Section 1.2.1, can be accomplished by simply trapping photons in an optical cavity. Giving photons a means of equilibration is harder. One way to overcome this is by using a dye that emits at a lower frequency than it absorbs [10], allowing incoherent loss of energy by photons to the “heat bath” provided by the dye. Another way is by coupling photons to polarizable material. This forms new *polariton* states which are part photon and part matter [11]. Since matter typically has interactions, this effectively gives the photons interactions. One arrangement is coupling photons to excitons, forming *exciton-polaritons*. Excitons themselves are composite bosons, being formed by bound electron-hole pairs (technically only approximately bosons [12]), and thus the composite photon-exciton system is also bosonic. Interest in exciton-polaritons has greatly increased since condensation was observed in both CdTe [13] and GaAs [14] systems. GaAs-based exciton-polaritons are the system I used for the work described in this thesis.

One of the primary advantages of exciton-polaritons over cold atoms is the extremely low mass,  $\sim 10^{-4}$  times the vacuum electron mass. This leads to condensation at much higher temperatures (see Section 1.4.9), with the early observations done at  $\sim 10$  K [13–15]. This is much easier to achieve than the nano- or micro-kelvin temperatures required in the cold atom systems mentioned above. In fact, the temperature limitations of exciton-polaritons are caused by weak exciton binding energies and photon-exciton coupling strengths, and there is no reason that condensation would not occur even at room temperature. Polariton lasing, which is inversion-less coherent emission, though not from a local ground state (see the discussion of “quasiconsensates” in Section 3.6), has been observed at room temperature in both GaN [16] and organic [17] systems, both of which have very large exciton binding energies and exciton-photon coupling strengths. While condensation into a well-defined

ground state remains elusive at room temperature, it should be possible in principle, with the only hurdles being caused by the difficulty of working with the most promising materials [9].

In addition to showing stimulated “polariton lasing” and condensation, exciton-polariton systems have shown a number of other effects. This includes signs of superfluidity, especially frictionless flow as evidenced by a suppression of defect scattering. This was first seen at low temperature in a GaAs-based system [18], but then later demonstrated at room temperature in an organic system [19]. The quantized vortices in phase associated with superfluidity have also been observed in exciton-polaritons [20, 21], as well as half-quantum vortices due to a combination of phase and polarization rotation [22]. Quantized rotation of polarization has also been observed in a polariton condensate in a ring geometry [23]. Another interesting effect revealed in exciton-polariton experiments are Josephson oscillations from the coupling of two separate macroscopically occupied bosonic states through a tunneling barrier [24]. In GaAs-based systems, essentially identical to those studied in this thesis, very long polariton lifetimes (as compared to the thermalization time) and long-range ballistic flow have been observed, due to extremely low defect density and highly reflective mirrors used to make the cavity [25]. These long-lifetime systems have also been shown to contain condensates very near to true thermal equilibrium, via direct measurement of the expected Bose-Einstein distribution [26]. Finally, work has been done to produce polaritons by electrical injection of free carriers, [27–29], moving polaritons closer to possible electronic applications.

### 1.1.1 Polaritons at equilibrium

Until recently, cold atoms have had the distinct advantage over polaritons of being able to reach thermal equilibrium. A useful comparison in determining how well a particular population can equilibrate is the ratio of its thermalization time to its lifetime, with equilibrium only attainable by those populations with thermalization times much shorter than their lifetimes. In cold atom systems, the lifetimes are typically  $\sim 1 - 10$  seconds, while the particle-particle scattering times are  $\sim 1 - 100$  ms [7, 8, 30]. Many individual scattering events are required to achieve thermal equilibrium, leading to thermalization times consistently  $\sim 10$  times longer than the scattering times [31]. The thermalization times are then

actually  $\sim 0.01 - 1$  seconds. This gives lifetimes about  $10 - 1000$  times greater than the thermalization times in the cold atom systems mentioned, leading to nicely equilibrated populations. For polaritons, the scattering time is estimated to be  $\sim 1$  ps [32], and this value was successfully used to fit experimental energy distributions [33]. This leads to thermalization times of  $\sim 10$  ps, which also agrees with experimental work that indicated thermalization in about 30-40 ps [34]. This was obviously problematic for reaching equilibrium when the longest polariton lifetimes were about 10 ps [35]. However, the latest generation of polaritons have lifetimes around 200 ps [25], so a thermalized population is certainly possible and has been confirmed experimentally [26]. With lifetimes now significantly longer than their thermalization times, polaritons are approaching levels of equilibrium comparable to cold atomic gases.

### 1.1.2 Advantages of polaritons

The fact that condensation takes place at much higher temperatures in polariton systems is obviously an important advantage over other systems, especially cold atoms. Being able to perform experiments even at liquid helium temperatures ( $\sim 4$  K) is much easier than working in the milli- and nano-Kelvin temperature range, and requires relatively simple cryogenics. The conservation of energy and momentum between polaritons and the photons into which they decay also makes probing the polariton system relatively simple (see Section 2.3.2). The ease of performing polariton experiments, compared to more complicated systems like cold atoms, is itself a case for studying the physics of bosons using polaritons.

From a practical point of view, another important advantage of polaritons over cold atoms is that they exist in a solid-state system. This brings the physics of boson condensation into the realm of device application. Stimulated scattering is a powerful effect, and a polariton condensate provides the essential ingredient: macroscopic occupation of a single quantum state. One of the most obvious applications is in low-threshold lasing. Just like in a conventional laser, emission from a polariton condensate is coherent, but, unlike conventional lasing, polariton lasing does not require population inversion. As mentioned above, polaritons in a “quasicondensate” also give coherent emission without population inversion.

They instead rely only on the fact that the polaritons occupy the same quantum state, and coherent emission can therefore be achieved with much lower input power [36]. In fact, just such a laser, complete with electrical pumping, has already been produced, though it only functioned at low temperature [27].

The work I present in this thesis suggests further applications in electronics beyond simply making lasers. In Chapter 4, I discuss how electrical current injected through a polariton device exhibits nonlinear behaviour linked to the formation of a polariton condensate. This could be used to build new types of electronic devices, akin to transistors, but now ruled by the effects of stimulated scattering. In Chapter 5, I show how electrical current is able to shift the momentum state of a polariton condensate. The result is essentially a laser whose emission angle is electronically tunable. This could lead to low-threshold polariton lasers with output beams that can be steered by simply changing an applied voltage. As work in new material systems progresses and polaritons move to room temperature, the effects observed here and elsewhere in polariton systems will become even more viable for use in actual devices [9].

### 1.1.3 Dissertation outline

The organization of this thesis is fairly straightforward. I begin with a theoretical introduction to the system, discussing cavity photons, excitons, and polaritons. Chapter 2 is focused on the samples and methods, with a brief introduction to the actual samples, a more detailed discussion of the fabrication I used after the initial growth, and then a discussion of the general optical and electrical measurements. Chapter 3 discusses the basic results of the fabrication, focusing on the resultant ground-state energy profile and appearance of polariton condensates. Chapter 4 provides the first major focus of the electrical current injection work: emphasizing studies of the photocurrent and increases to the polariton density. Chapter 5 presents a study of the drag effect of injected electrons on the polaritons. The work reported in Chapter 6 is a bit different, with a focus on below-threshold transport distances of both polaritons and excitons. I give some final remarks in Chapter 7, and provide several appendices with further pertinent calculations, derivations, and experimental results

and details.

## 1.2 CAVITY PHOTONS

### 1.2.1 The cavity photon energy dispersion

A cavity photon can be thought of most simply as a photon temporarily trapped between two opposing mirrors. Like any confined wave, only those wavelengths that constructively interfere will survive while the others will be suppressed. The result is a resonator with particular modes, separated from each other by the free spectral range of the cavity. The most typical arrangement uses two planar mirrors, and is called a Fabry-Pérot cavity. The condition for constructive interference is simply that a wave must undergo an integer number of  $2\pi$  changes in phase over one full round trip within the cavity. This allows a wave that has already made  $N$  round trips to constructively interfere with another wave that has made  $N'$  round trips (or, similarly, with a different part of itself). For waves propagating at an angle not exactly normal to the planes of the mirrors, the condition for constructive interference brings about one of the most interesting features of a cavity photon: an effective mass for the normally massless photon.

Starting with a standard plane wave, given by

$$E(x, t) = E_0 e^{i(kz - \omega t)}, \quad (1.1)$$

we can see that, for a particular point in time, the total change in phase  $\phi$  across a distance  $d$  is given by  $kd$ . For propagation normal to the plane of the mirrors, and a cavity of length  $\ell$ , constructive interference is obtained when  $2k\ell = (2\pi)m$ , where  $m$  is an integer.

For the general case of non-normal propagation, consider Figure 1.1. For two in-phase and parallel waves, which overlap in space after one has made a single round trip of the cavity, we see that the extra distance traveled is not simply twice the length of the cavity. Both the wave inside and the wave outside the cavity gain some phase due to propagation in their respective media. The phase change in each medium (distinguished by subscript

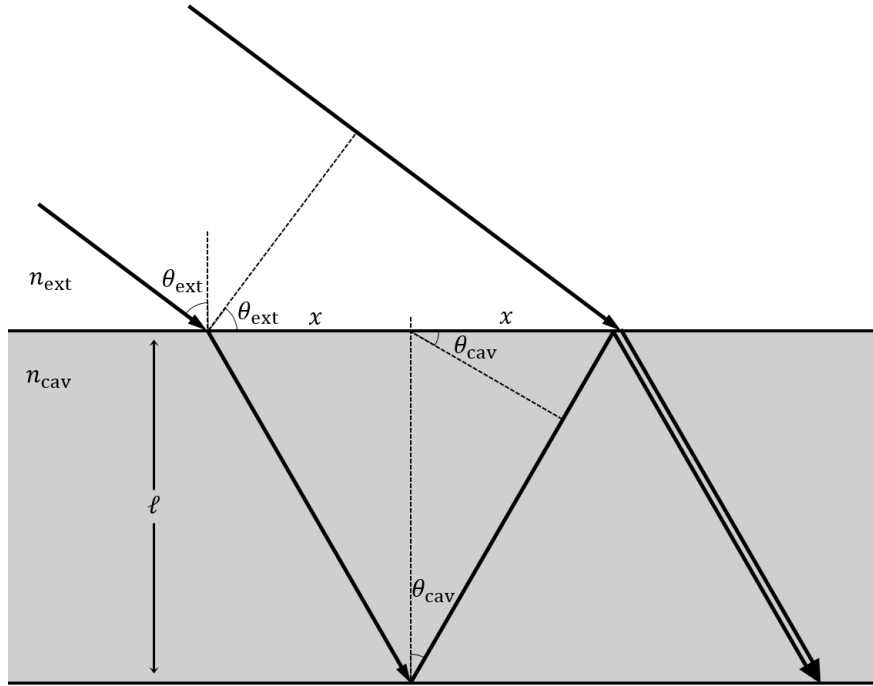


Figure 1.1: Diagram for the cavity interference condition. Constructive interference will occur when the difference in the extra path length of the left-side ray within the cavity and the extra path length of the right-side ray before reaching the cavity are equal to  $2\pi$  times an integer.

$i$ ) is given by the extra distance traveled in that medium  $d_i$  multiplied by the wavenumber  $k_i = k_{\text{vac}} n_i$ , where  $k_{\text{vac}}$  is the wavenumber in vacuum. The total difference in phase of the two waves is then be given by

$$\begin{aligned}
\Delta\phi &= k_{\text{cav}} d_{\text{cav}} - k_{\text{ext}} d_{\text{ext}} \\
&= 2k_{\text{vac}} n_{\text{cav}} (\ell \cos \theta_{\text{cav}} + x \sin \theta_{\text{cav}}) - k_{\text{vac}} n_{\text{ext}} (2x) \sin \theta_{\text{ext}} \\
&= 2\ell k_{\text{vac}} n_{\text{cav}} \cos \theta_{\text{cav}} + 2x k_{\text{vac}} (n_{\text{cav}} \sin \theta_{\text{cav}} - n_{\text{ext}} \sin \theta_{\text{ext}}) \\
&= 2\ell k_{\text{vac}} n_{\text{cav}} \cos \theta_{\text{cav}},
\end{aligned} \tag{1.2}$$

where  $2x$  is the in-plane distance traveled by the wave during the round trip,  $\theta_{\text{cav}}$  is the angle of propagation within the cavity,  $n_{\text{cav}}$  is the refractive index of the cavity,  $\theta_{\text{ext}}$  is the angle of propagation outside of the cavity, and  $n_{\text{ext}}$  is the refractive index outside of the cavity. Snell's law allows the elimination of two of the contributions, leaving only one term dependent on  $\cos \theta_{\text{cav}}$ . The condition for constructive interference is now

$$\Delta\phi = 2\ell k_{\text{cav}} \cos \theta_{\text{cav}} = (2\pi)m, \quad m = 1, 2, 3, \dots \tag{1.3}$$

In general, there may or may not be a phase shift of  $\pi$  upon reflecting at each mirror. If we assume that the cavity is symmetric, with the same material on the interface of each mirror, then there will be a total phase shift of either 0 or  $2\pi$  after the two reflections during a round trip. Either way, there is no overall effect on the condition for constructive interference.

It is helpful to recognize that  $k_{\text{cav}} = |\mathbf{k}_{\text{cav}}| = \sqrt{k_{\perp}^2 + k_{\parallel}^2}$ , where  $k_{\perp}$  and  $k_{\parallel}$  are the perpendicular and in-plane components of the wavevector  $\mathbf{k}_{\text{cav}}$ , respectively. This allows us to make the substitution  $k_{\perp} = k_{\text{cav}} \cos \theta_{\text{cav}}$ , giving us

$$k_{\perp} = \frac{m\pi}{\ell}. \tag{1.4}$$

Thus, using only the condition of constructive interference, we can see that the cavity simply constrains the perpendicular component of the electromagnetic wave at any arbitrary angle of propagation. This is the intuitive result, and allows us to write the energy dispersion as

$$E_{\text{cav}} = \frac{\hbar c |\mathbf{k}_{\text{cav}}|}{n_{\text{cav}}} = \frac{\hbar c}{n_{\text{cav}}} \sqrt{k_{\perp}^2 + k_{\parallel}^2} = \frac{m\pi \hbar c}{\ell n_{\text{cav}}} \sqrt{1 + \frac{\ell^2 k_{\parallel}^2}{(m\pi)^2}}. \tag{1.5}$$

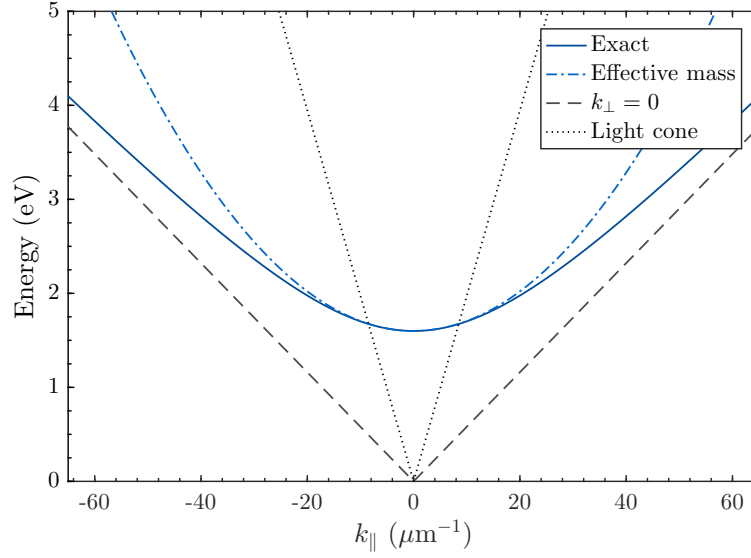


Figure 1.2: The cavity photon dispersion for a cavity with photon energy of 1.6 eV at  $k_{\parallel} = 0$  and a refractive index of 3.4. The exact result is from Equation 1.5, while the plot using the effective mass is from Equation 1.6. The result using  $k_{\perp} = 0$  is simply  $E = \hbar c k_{\parallel} / n_{\text{cav}}$ . The light cone shows the range of  $k_{\parallel}$  for each photon energy that will escape the cavity into vacuum. Anything outside of the light cone is totally-internally reflected.

This gives the cavity photon energy dispersion as a function of  $k_{\parallel}$ . For sufficiently small  $k_{\parallel}$ , we can make the approximation

$$E_{\text{cav}}(k_{\parallel}) \approx \frac{m\pi\hbar c}{\ell n_{\text{cav}}} \left( 1 + \frac{\ell^2 k_{\parallel}^2}{2(m\pi)^2} \right) = E_{\text{cav}}(0) + \frac{\hbar^2 k_{\parallel}^2}{2m_{\text{cav}}}, \quad (1.6)$$

which introduces the effective cavity mass  $m_{\text{cav}} = (n_{\text{cav}}^2/c^2)E_{\text{cav}}(0)$ . This shows that, at small in-plane momentum, the confinement of the cavity gives the cavity photon a dispersion very similar to a particle with mass. Figure 1.2 shows a summary of these results for a cavity refractive index of 3.4 and photon energy of 1.6 eV at  $k_{\parallel} = 0$ , which is close to the values for the cavities in the actual samples I will discuss in later chapters. For comparison, the light cone of the cavity has been plotted on the same axes. This simply gives the range of energies at each  $k_{\parallel}$  that escape the cavity into vacuum: all photons outside of the light cone are totally-internally reflected at the cavity mirrors. The effective-mass approximation is clearly very good at small values of  $k_{\parallel}$ , and only begins to slightly diverge from the exact result at the maximum values within the light cone. At large  $k_{\parallel}$ , the cavity dispersion approaches the standard photon dispersion for light traveling within and parallel to the plane of the cavity,  $E = \hbar c k_{\parallel} / n_{\text{cav}}$ .

While this is a fairly obvious result, one advantage in introducing the interference condition is that it reveals a peculiarity of cavities that are made of more than one material layer. In such cavities, the phase change in Equation 1.2 becomes a sum of the phase changes across each individual layer, with different refractive indices for each layer. Using the index  $i$  to designate each individual layer of the cavity, the phase difference becomes

$$\Delta\phi = 2k_{\text{vac}} \sum_i \ell_i n_i \cos \theta_i = 2k_{\text{vac}} \ell n_{\text{cav}} \cos \theta_{\text{cav}}, \quad (1.7)$$

where  $\ell$ ,  $n_{\text{cav}}$  and  $\theta_{\text{cav}}$  are now overall *effective* values for the cavity length, cavity refractive index, and propagation angle within the cavity, respectively. A reasonable choice for the effective length is simply the total length, given by  $\ell = \sum_i \ell_i$ . Snell's law can be used to relate all of the propagation angles to the external propagation angle  $\theta_{\text{ext}}$  and the refractive indices, using  $n_{\text{ext}} \sin \theta_{\text{ext}} = n_{\text{cav}} \sin \theta_{\text{cav}} = n_i \sin \theta_i$ . It is then clear from this expression that the effective cavity refractive index  $n_{\text{cav}}$  is somewhat dependent on the the propagation angle of the light within the cavity, and not simply the average of all the layers. Intuitively, this can

be understood by imagining the zig-zag path of a light ray propagating through the multiple layers of the cavity, determined by Snell's law at each interface. For normal incidence, the path is straight and  $n_{\text{cav}}$  is given by the weighted average of all of the refractive indices within the cavity:  $n_{\text{cav}} = \ell^{-1} \sum_i \ell_i n_i$ . But as the angle becomes larger,  $n_{\text{cav}}$  is slightly changed from the simple average. This change is obviously greater for greater contrast between the indices of refraction of the layers (resulting in a more zig-zaggy path). For the most part, this effect is small, but, for the sake of completeness, it is important to recognize that the  $n_{\text{cav}}$  in Equation 1.5 is itself slightly dependent on  $k_{\parallel}$  when multiple material layers are used in the cavity. For a more complete discussion and example of this effect, see Appendix A.

### 1.2.2 Free spectral range and microcavities

One important property of an optical cavity is the free spectral range (FSR). This is simply the frequency spacing between adjacent modes of the cavity. From Equation 1.4, the mode separation  $\Delta\omega_{\text{cav}}$  is given by

$$\begin{aligned} \omega_m &= \frac{\pi c m}{\ell n_{\text{cav}}} \\ \Rightarrow \Delta\omega_{\text{cav}} &= \omega_{m+1} - \omega_m = \frac{\pi c}{\ell n_{\text{cav}}}. \end{aligned} \tag{1.8}$$

A typical laser cavity is primarily designed to function at high orders (large  $m$  values) of the cavity. Thus, the cavity length  $\ell$  is usually much larger than the wavelength of the light emitted. A microcavity, on the other hand, is distinguished by a cavity length on the order of the wavelength of light within the cavity material, which is typically  $\lesssim 1 \mu\text{m}$ . This means that the mode orders considered in designing a microcavity are also much smaller (often  $m = 1, 2, 3, \dots$ ). Microcavities are often described with the length given by  $\ell = m\lambda/2$ , where  $\lambda$  is the design wavelength within the cavity ( $\lambda = \lambda_{\text{vac}}/n_{\text{cav}}$ , with  $\lambda_{\text{vac}}$  representing the wavelength of the light in vacuum).

The result is that, for a given design wavelength, the free spectral range of a microcavity is typically much larger than that of a conventional cavity. Rewriting the FSR in terms of a design wavelength,

$$\Delta\omega_{\text{cav}} = \frac{2\pi c}{m\lambda_{\text{vac}}}, \tag{1.9}$$

we see that the mode separation will be much larger for cavities designed to work at lower mode orders.

We have thus far explained how a cavity can be formed in general, and the properties of the cavity that affect the resultant energy dispersion. A real physical system has a few other important properties, which are generally affected by the mirrors themselves. I will discuss these next.

### 1.2.3 $Q$ -factor: the cavity as an harmonic oscillator

For any physical system, the mirrors will not be perfect, leading to loss of photons as they leak through the mirrors. One way to describe the “leakiness” of the cavity is the *quality factor*, typically called the  $Q$ -factor. This is usually defined as the ratio

$$Q \equiv \frac{\omega_{\text{cav}}}{\delta\omega} \quad (1.10)$$

where  $\omega_{\text{cav}}$  is the cavity mode frequency and  $\delta\omega$  is the cavity mode full-width at half-maximum (FWHM) linewidth, both of which are empirical values. In this section, I will show that  $Q$  is related to both the cavity lifetime (the characteristic time a photon spends within the cavity), and the frequency linewidth.

The  $Q$ -factor is a standard way of describing resonators in electronics (see e.g. pages 41-42 of [37]). However, it can be applied to any resonator, including optical cavities. In the case of an optical cavity, the oscillating function is the electromagnetic field of the light. If we imagine a cavity already populated with a quantity of photons, we can see that the amplitude of the electric field will decrease over time as the photons escape from the cavity. This is equivalent to a damped harmonic oscillator with some initial amplitude, which is described by the equation

$$\ddot{E} + \gamma\dot{E} + \omega_0^2 E = 0, \quad (1.11)$$

where  $E(t)$  is the time dependent electric field within the cavity. I will guess  $E(t) = E_0 e^{i\omega t} e^{-\gamma t/2}$  as a solution, and put it into Equation 1.11, getting

$$\begin{aligned}
& (i\omega - \frac{\gamma}{2})^2 + \gamma(i\omega - \frac{\gamma}{2}) + \omega_0^2 = 0 \\
\Rightarrow \omega^2 &= \omega_0^2 - \frac{\gamma^2}{4} \\
\Rightarrow \omega &= \omega_0 \sqrt{1 - [\gamma/(2\omega_0)]^2} \\
\Rightarrow E(t) &= E_0 e^{i\omega_0 \sqrt{1 - [\gamma/(2\omega_0)]^2} t} e^{-\gamma t/2}.
\end{aligned} \tag{1.12}$$

From this we can see that the electric field amplitude is exponentially damped, with a decay lifetime of  $2/\gamma$ , and that the frequency is also shifted by the damping. The frequency shift is small when  $\gamma \ll \omega_0$ . We can now relate this field strength to the number of photons  $N$ , recognizing that

$$N(t) \propto |E(t)|^2 = |E_0|^2 e^{-\gamma t}. \tag{1.13}$$

Therefore, the damping term  $\gamma$  is simply the inverse cavity photon lifetime ( $\tau_{\text{cav}} = 1/\gamma$ ).

Next, I will introduce a driving field to the oscillator, giving an inhomogeneous version of the homogeneous Equation 1.11:

$$\ddot{E} + \gamma \dot{E} + \omega_0^2 E = A e^{i\omega t}. \tag{1.14}$$

Since the driving field has a constant amplitude in time, we should not expect an exponentially damped solution as in the undriven case. Instead, after some time, the system should reach a steady state, with a similarly constant field amplitude in time. Following Snoke (chapter 7 of Ref. [38]), I guess the solution  $E(t) = E_0 e^{i(\omega t + \delta)}$ , where  $\delta$  is introduced to account for a phase difference between the driving field and the field within the cavity. Plugging this into Equation 1.14 gives

$$[(i\omega)^2 + \gamma(i\omega) + \omega_0^2] E_0 e^{i\delta} = A. \tag{1.15}$$

It is worth pointing that  $A$  could also be complex without affecting the result. In such a case, we could write  $A = A_0 e^{i\phi}$ , and then Equation 1.15 would look exactly the same but with  $A \rightarrow A_0$  and  $\delta \rightarrow \delta - \phi$ . Thus, in Equation 1.15, we can assume both  $A$  and  $E_0$  are real and the important complex phase information can be completely described by  $\delta$  without

loss of generality. Since this equation must be fulfilled individually for both the real and imaginary parts, it actually provides two equations:

$$\begin{aligned}
\text{Imaginary} &\implies \gamma\omega E_0 \cos \delta + (\omega_0^2 - \omega^2) E_0 \sin \delta = 0 \\
&\implies -(\omega_0^2 - \omega^2) \sin \delta = \gamma\omega \cos \delta \\
&\implies \tan \delta = \frac{-\gamma\omega}{\omega_0^2 - \omega^2} \\
\text{Real} &\implies (\omega_0^2 - \omega^2) E_0 \cos \delta - \gamma\omega E_0 \sin \delta = A \\
&\implies E_0 = \frac{A}{(\omega_0^2 - \omega^2) \cos \delta - \gamma\omega \sin \delta}.
\end{aligned} \tag{1.16}$$

The equation for the imaginary part gives the phase, while the real part gives the field amplitude. To finish solving for the amplitude, we can take the square of the denominator and use the second line of Equation 1.16:

$$\begin{aligned}
E_0 &= \frac{A}{\sqrt{(\omega_0^2 - \omega^2)^2 \cos^2 \delta - 2\gamma\omega(\omega_0^2 - \omega^2) \cos \delta \sin \delta + \gamma^2 \omega^2 \sin^2 \delta}} \\
&= \frac{A}{\sqrt{(\omega_0^2 - \omega^2)^2 \cos^2 \delta + [(\omega_0^2 - \omega^2)^2 \sin^2 \delta + \gamma^2 \omega^2 \cos^2 \delta] + \gamma^2 \omega^2 \sin^2 \delta}} \\
&= \frac{A}{\sqrt{(\omega_0^2 - \omega^2)^2 + \gamma^2 \omega^2}}.
\end{aligned} \tag{1.17}$$

This gives the particular solution for this driving field, but the undriven solution can also be added for the general case. The full solution is then

$$E(t) = \alpha e^{(i\omega_0 \sqrt{1 - [\gamma/(2\omega_0)]^2} + i\phi - \gamma/2)t} + \frac{A e^{i(\omega t + \delta)}}{\sqrt{(\omega_0^2 - \omega^2)^2 + \gamma^2 \omega^2}}, \tag{1.18}$$

where  $\alpha$  and  $\phi$  are set by the initial conditions. If we consider a situation long after the decaying part has disappeared, we can find the frequency response of the cavity to the driving field. As I mentioned above, the photon number (and field intensity) are related to the field amplitude by

$$N(t) \propto I(t) \propto |E(t)|^2 \xrightarrow{t \rightarrow \infty} \frac{A^2}{(\omega_0^2 - \omega^2)^2 + \gamma^2 \omega^2}, \tag{1.19}$$

which is clearly a peaked function of the driving frequency  $\omega$ , with the peak near the resonant frequency  $\omega_0$ . For  $\gamma \ll \omega_0$ , this will be very strongly peaked at the resonance. Thus,  $\omega$  is

approximately just  $\omega_0$  for all of the frequency region that matters. This means I can replace  $\omega$  with  $\omega_0$  everywhere except where one is subtracted from the other, giving

$$|E(t)|^2 = \frac{A^2}{((\omega_0 - \omega)(\omega_0 + \omega))^2 + \gamma^2 \omega^2} \xrightarrow{\gamma \ll \omega_0} \frac{A^2/(2\omega_0)^2}{(\omega_0 - \omega)^2 + (\gamma/2)^2}. \quad (1.20)$$

This is now simply a Lorentzian, with a FWHM in frequency of  $\gamma$ . This is exactly the value we need for Equation 1.10 to find  $Q$ :

$$Q = \frac{\omega_{\text{cav}}}{\delta\omega} = \frac{\omega_0}{\gamma} = \omega_0 \tau_{\text{cav}}. \quad (1.21)$$

The quality factor is therefore closely related to both the lifetime of a photon within the cavity, and the frequency linewidth of the cavity field intensity.

#### 1.2.4 Cavity finesse

The finesse of a cavity is simply the ratio of the mode separation to the frequency linewidth. It is therefore closely related to both  $Q$  and the free spectral range, since the former gives the linewidth and the latter gives the mode separation. Designating the finesse as  $\mathcal{F}$ , it is given by

$$\mathcal{F} = \frac{\Delta\omega_{\text{cav}}}{\delta\omega} = \tau_{\text{cav}} \Delta\omega_{\text{cav}} = \frac{Q \Delta\omega_{\text{cav}}}{\omega_{\text{cav}}}. \quad (1.22)$$

The finesse can be thought of as a measure of the distinctness of each individual mode of the cavity. If the linewidth  $\delta\omega$  is similar to the FSR (corresponding to  $\mathcal{F} \sim 1$ ), then the individual modes begin to blend together. Since the free spectral range  $\Delta\omega_{\text{cav}}$  is typically very large for a microcavity, the finesse can also be very high. From Equation 1.9, if we use the relationship  $\lambda_{\text{vac}} = 2\pi c/\omega_{\text{cav}}$ , we can write the finesse as

$$\mathcal{F} = \frac{Q}{\omega_{\text{cav}}} \frac{\omega_{\text{cav}}}{m} = \frac{Q}{m}, \quad (1.23)$$

which shows that the finesse of a microcavity is essentially given by the quality factor and the mode order. For a design using the first order mode, corresponding to the half-wavelength cavity length,  $Q$  and  $\mathcal{F}$  are equivalent, and higher order designs will greatly reduce the finesse.

### 1.2.5 Distributed Bragg reflectors

As I showed in Section 1.2.3, the cavity mode linewidth and lifetime are both dependent on the “lossiness” of the mirrors. The choice of mirrors is then very important when designing a cavity, and will greatly affect its characteristics. I will limit this discussion to distributed Bragg reflectors (DBRs), since those are what I used in my work. I will also show that, while at first glance the general work of Section 1.2.1 seems to be independent of the mirrors, the cavity length is not so simple in the case of DBRs.

A DBR is essentially a stack of alternating dielectric material layers that utilize the standard thin film interference condition. Similar to the condition for a cavity, in this case the total phase change on a round trip through each film must allow constructive interfere with the portion of the wave that was reflected at the initial film interface. This calls for an analysis similar that in Section 1.2, with an overall phase difference between the wave reflected at the first interface and the wave reflected at the second interface. Referring to Figure 1.3, which shows a single layer, this phase is given by

$$\begin{aligned}\Delta\phi &= 2k_{\text{vac}}n_q(d_q \cos \theta_q + x \sin \theta_q) - k_{\text{vac}}n_p(2x) \sin \theta_p \\ &= 2d_qk_{\text{vac}}n_q \cos \theta_q.\end{aligned}\tag{1.24}$$

In this case,  $d_q$  is the thickness of the  $q^{\text{th}}$  layer,  $n_p$  and  $n_q$  give the refractive indices of the two materials,  $k_{\text{vac}}$  is the vacuum wavenumber, and  $\theta_p$  and  $\theta_q$  give the propagation angles in each material.  $x$  is again the in-plane propagation distance traveled during the round trip through the layer. While the phase difference after traveling one round trip is the same as for the cavity, the interference condition is slightly different due to phase changes upon reflection at certain interfaces. In the case of a symmetric cavity, these always either cancel out (since cavities involve the interference of the wave with itself upon multiple reflections) or do not take place due to the refractive index of the cavity being larger than that of the mirrors. For the DBRs, we can simply add a phase change term for each interface, changing Equation 1.3 into

$$\Delta\phi - \phi_{p \rightarrow q} + \phi_{q \rightarrow p} = 2d_qk_q \cos \theta_q - \phi_{p \rightarrow q} + \phi_{q \rightarrow p} = (2\pi)m, \quad m = 0, 1, 2, 3, \dots, \tag{1.25}$$

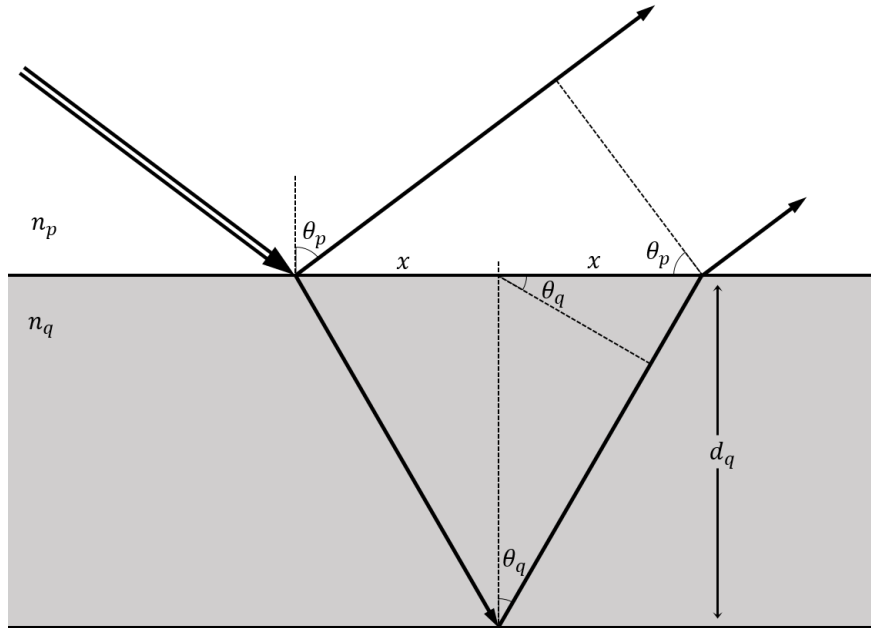


Figure 1.3: Diagram for Bragg interference of a thin film. Constructive interference will occur when the waves reflected at the first and second surface have a phase difference of  $2\pi m$ , where  $m$  is an integer. That phase difference will be caused by both differences in effective path length as well as  $\pi$  phase changes when reflected at certain interfaces.

where  $\phi_{p \rightarrow q}$  is the phase change of the wave while in material  $p$  and reflected from material  $q$ , and I also made the substitution  $k_q = k_{\text{vac}} n_q$ . For a DBR, the next material below any layer is the same as the one above (with the exception of the initial and final layers), so one phase change will always be zero while the other will be  $\pi$ . The interference condition also allows for the  $m = 0$  case since the phase shift given by travel in a layer  $q$  can be exactly compensated by the phase shift upon reflecting at that layer ( $\phi_{p \rightarrow q} = \pi$ ). The  $m = 0$  case with  $\phi_{q \rightarrow p} = \pi$  is clearly unphysical since it requires a negative phase shift from passage through the material, so we can ignore it. Equation 1.25 can now be simplified to

$$\Delta\phi = 2d_q k_q \cos \theta_q = (2m + 1)\pi, \quad m = 0, 1, 2, 3, \dots \quad (1.26)$$

which for the lowest order case and normal incidence gives  $d_q = \pi/2k_q = \lambda_{\text{ext}}/4n_q$ . For end layers (initial and final), the condition can be maintained as long as the external material continues the pattern. If the external material has a lower refractive index than the end material, then the end material must have the higher refractive index of the two materials in the DBR. The opposite is true if the external material has higher refractive index.

The essential concept of a DBR is to meet the interference condition in Equation 1.26 for each layer of the stack, causing the constructive interference to build with the number of layers. While this will clearly lead to high reflectivity near the frequency that meets the condition, the full picture is much more complicated. First, in order to get a sense of the total reflectivity, it is not enough to simply know the reflectivity of each material. This is because there will actually be an infinite series of reflections *within each layer*, all of which will add a contribution to the interference. Furthermore, we also need to know the details of the frequency and incident angle dependence of the DBR reflectivity. To fully characterize a DBR, or, for that matter, the optical characteristics of any multi-layer stack, the *transfer matrix method* is typically employed. This method essentially fulfills the boundary conditions of Maxwell's equations at each layer interface, and propagates the waves through homogeneous materials. For a discussion of this method, see Appendix B.

Figure 1.4 shows the reflectance of a DBR as a function of photon energy with a variety of design parameters. The design wavelength was 780 nm, corresponding to an energy of about 1.59 eV, and the structure was considered to be in vacuum. In every case the layer thicknesses

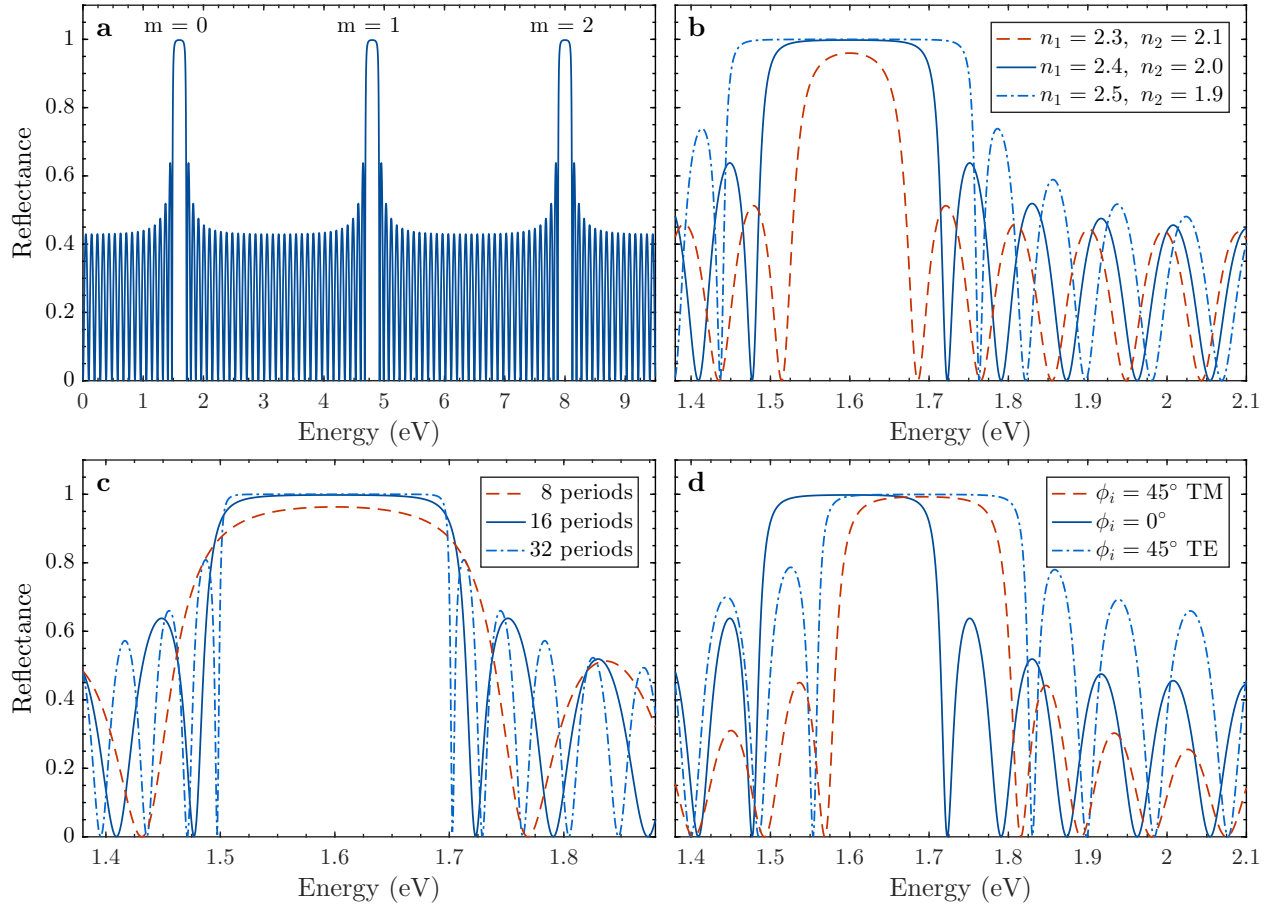


Figure 1.4: The reflectance of a DBR in a vacuum with a design wavelength of 780 nm (about 1.59 eV) as a function of photon energy, calculated using the transfer matrix method. In every case the layer thicknesses were set to exactly a quarter of the design wavelength ( $\lambda/4n$ ). (a) A wide range of energies, showing nearly unity reflectance at the first three “stop bands” corresponding to  $m = 0, 1, 2$  in Equation 1.26. (b) The reflectance of the first order stop band, using different refractive indices in the structure. 16.5 layer periods were used for all three cases. (c) The reflectance of the first order stop band with different numbers of alternating layer periods (assuming an extra half-period in each case). In each case, the refractive indices of the materials were 2.4 and 2.0. (d) The reflectance at large incident angle and both TE and TM polarization. All three cases used 16.5 layer periods and 2.4 and 2.0 as refractive indices.

were set to exactly a quarter of the design wavelength ( $\lambda/4n$ ). Figure 1.4(a) shows a wide range of energies, revealing equally spaced ranges of nearly unity reflectance. These regions in energy correspond to  $m = 1, 2, 3$  in Equation 1.26 and are called the DBR “stop bands.” The characteristics of these stop bands are determined mostly by a few basic parameters, as shown in the other panels. Figure 1.4(b) shows the effect of changing the refractive index contrast. While keeping the average refractive index the same, there are clearly two changes. The first is that the maximum reflectance at the design energy is higher when the index contrast is higher. The difference is nearly indistinguishable in the two higher-contrast cases due to the reflectance asymptotically approaching unity as the contrast gets higher, but the effect is still significant when considering the change in “loss” via transmittance: the transmittance for the intermediate case (blue solid line) is 0.0020, which is actually about 20 times higher than highest-contrast case (the dashed blue line). The second effect is that the width of the stop band is also increased as the index contrast is increased, giving a wider range of energies with nearly unity reflectance. Figure 1.4(c) show the effects of increasing the number of alternating layer periods in the structure. As the period number is increased the maximum reflectance is increased. The sharpness of the stop band profile is also increased with the number of periods, which, while not significantly changing the width of the stop band, does increase the usable range of energies with very high reflectance. Figure 1.4(d) shows the effect of high incident angle at both TE and TM polarizations. As expected from Equation 1.26, an increase in incident angle must be matched by an increase in wavenumber, which corresponds to an increase in photon energy. This is clearly seen in the shift of the center of the stop band to higher energy for both polarizations. It is also interesting to note that TE polarization at high angle actually results in a sharper and wider stop band profile, effectively increasing the quality of the DBR. TM polarization does the opposite. The reason for this is that, for TE polarization, increasing the incident angle has the same effect as increasing the index contrast (as seen in Equation B.3). The opposite is true for TM polarization, for which increasing the incident angle has the same effect as reducing the index contrast (as seen in Equation B.17).

Another important feature of a DBR is the phase change that occurs at the external interfaces. These shifts are plotted in Figure 1.5 for both the reflected and transmitted

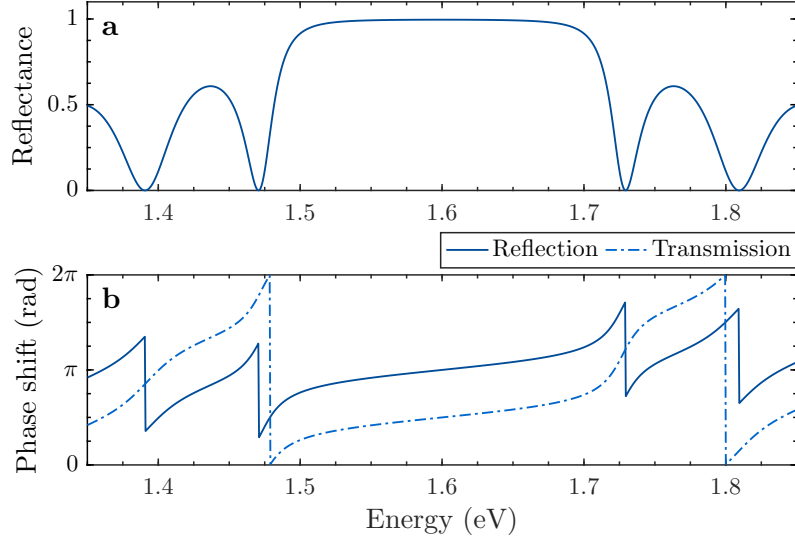


Figure 1.5: (a) The reflectance at normal incidence as a function of photon energy for a DBR with 16.5 periods and refractive indices of 2.4 and 2.0, and a central design energy of 1.6 eV, plotted here for comparison with (b) and produced using the transfer matrix method. The layer thicknesses were set to exactly a quarter of the design wavelength ( $\lambda/4n$ ). (b) The phase shift with respect to the incident wave of the complex reflection and transmission coefficients.

waves. At the exact center of the stop band, there is a  $\pi$  phase shift for the reflected part, which is the same result as for a wave reflecting from an interface between two infinite media while in the medium with lower refractive index. The transmitted part has an overall  $\pi/2$  phase shift at the central energy. This is to be expected since, for lowest order in Equation 1.26, each layer will induce a  $\pi$  phase shift on a full round trip. This corresponds to a  $\pi/2$  phase shift for the transmitted wave, which only makes a single pass through each layer. In this case, there are 16 full periods, each contributing a shift of  $\pi$ , and a final single layer of the higher refractive index material was added to meet the interference condition when passing into vacuum. This gives a total phase shift of  $16.5 \times \pi$ . The most interesting result is the slight phase shift for the reflected wave within the stop band but at energies not exactly in the center. This reveals a shifting of the node of the reflected wave.

One way to think about this is in terms of interference conditions, which usually require a particular shift in phase. If a shift of  $\pi$  is expected at the surface of the mirror, then when below the central DBR stop band energy, the DBR has less of an effect than expected, so a longer path-length will be required to meet the same interference condition. The opposite is true for energies above the central stop band energy. An equivalent picture is shifting the nodes of the wave away from the interface. For the case of energies below the central energy, since the shift from traveling in the material is  $< \pi$ , the node, which corresponds to exactly a  $\pi$  phase shift, will be out slightly along the path of the reflected wave in the original medium. On the other hand, the higher energy waves effectively “skip ahead,” and therefore interfere with waves that have nodes slightly inside of the DBR. These results are important when considering cavities made out of DBRs, since they result in shifts of the nodes of the cavity mode into or out of the mirrors. This effectively shifts the cavity mode to either higher or lower energies.

### 1.2.6 DBR microcavities

For my work, a single DBR by itself is not the full picture. Instead, two DBRs are placed at close proximity, with a narrow cavity in between. The two DBRs can be different in principle, but must have similar stop band energies. To obtain high  $Q$ , the cavity spacer

must be designed to support modes that overlap with the stop bands of the DBRs. Using the lowest order, the result is two stacks of  $\lambda/4n$  layers with a  $\lambda/2n$  spacer in between. Just like the regular DBR, this stack of materials can also be simulated with the transfer matrix method (discussed in Appendix B). The result of one such design is shown in Figure 1.6(a). Obviously, a  $\lambda/2n$  layer does not meet the interference condition for a DBR given in Equation 1.26, so we should not expect the same flat stop band. What appears is a narrow dip in reflectance, corresponding to a peak in transmission through the structure at the cavity mode. Since the cavity spacer gives a cavity mode energy the same as the central energy of the DBR stop bands, the cavity mode appears exactly in the center.

Some interesting features emerge when the cavity thickness is idealized for a mode energy not quite aligned with the central stop band energy (when the detuning  $\Delta = E_{\text{cav}} - E_{\text{DBR}} \neq 0$ ). Figure 1.6(b) shows the result for the same DBR structure as panel (a) but with a slightly wider cavity spacer, corresponding to an energy of 1.5 eV instead of 1.6 eV. Even at normal incidence, the cavity mode appears as a dip in the reflectance at considerably higher energy than the energy expected purely from the width of the cavity. To explain this, it is necessary to recall that the phase shift for the reflected light from a DBR is not exactly either  $\pi$  or zero. As shown in Figure 1.5(b), for photon energies less than the central stop band energy, the phase shift is less than expected. This corresponds to an effective shift of the electric field node to a point outside of the DBR. While the details of the interference for a DBR microcavity obviously require a more careful study of the structure as a whole and not simply bouncing between DBR surfaces, the concept is the same. In this case, the effect is to shift the nodes to a point inside the cavity spacer, effectively reducing the cavity length and increasing the cavity mode energy. The opposite is true of modes with energies slightly above the central stop band energy [an example of both cases can be seen in Figure 1.7(b)].

Another interesting feature that appears at non-zero incident angle and detuning between the stop band and cavity mode is a TE-TM splitting of the cavity mode. Figure 1.6(b) shows an example at fairly large angle and large detuning, resulting in a TM mode about 10 meV above the TE mode. The splitting is reversed when the cavity mode is above the central DBR stop band energy. This effectively means that angle dependence of cavity mode energy is different for TE and TM polarizations, resulting in slightly different effective masses for

the two states. For a more full discussion of this splitting, see Ref. [39].

One final consideration is worth mentioning when it comes to DBR microcavities. In a system where the desired result is a single (excepting the TE-TM splitting, of course) cavity photon state, the free spectral range of the cavity must be more than twice as large as the DBR stop band width. The stop band width, as shown in Figure 1.4, can be controlled mostly by the choice of the refractive indices, corresponding to the choice of materials. This is often dictated by other concerns, such as absorption in certain energy ranges or material compatibility. The free spectral range is primarily controlled by the choice of the cavity mode order, given by the different integer values in Equation 1.3. The case of a  $\lambda/2$  cavity is shown in Figure 1.7(a), with the “nominal” cavity modes (those expected simply from the cavity thickness without considering DBR phase shifts) marked with vertical gray lines. Since the DBR stop bands appear on odd integer values only, there is an extra cavity mode between each stop band. Clearly, in this case, the free spectral range  $\hbar\omega_{\text{cav}}$  is much wider than the stop band widths, resulting in a single high- $Q$  cavity mode in a wide energy range. A contrasting case is shown in Figure 1.7(b). In this case, the structure is designed with a  $15\lambda/2n$  cavity between the same DBRs as panel (a). In this case, the cavity modes just above and just below the central mode also slightly overlap with the stop band. The result is three high- $Q$  cavity modes in a narrow range of energies. It is also interesting to note that the actual modes away from the center are both shifted away from the expected energies towards the center of the stop band. This is due to the phase shift when not at the central stop band energy, which shifts the position of the cavity mode nodes to points away from the actual interfaces between the cavity spacer and the DBRs, as discussed above.

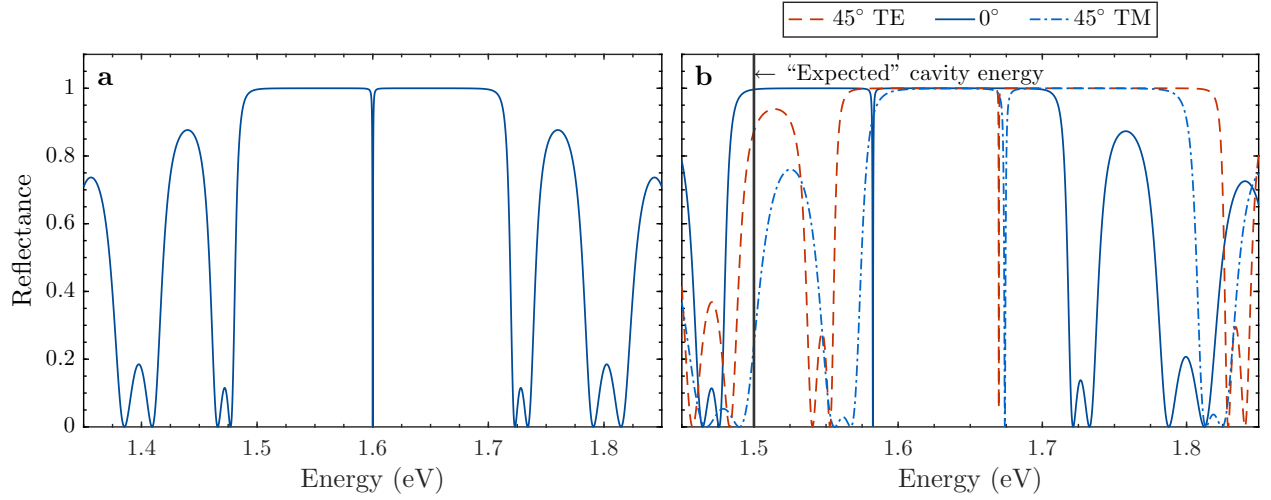


Figure 1.6: (a) Reflectance as a function of photon energy for a DBR cavity structure. The refractive indices of the DBR were 2.4 and 2.0, and the refractive index of the cavity material was 2.5. The DBR layers were each exactly  $\lambda/4n$  thick while the cavity was exactly  $\lambda/2n$  thick for a design energy of 1.6 eV. Exactly 16 periods were used in both DBRs, resulting in two identical structures arranged opposite each other with a cavity space between. (b) The same structure as in (a), but with the cavity spacer at a thickness of  $\lambda/2n$  for an energy of 1.5 eV, and also showing the case of light incident at 45° for both TE and TM polarization. The DBR design energy was still 1.6 eV.

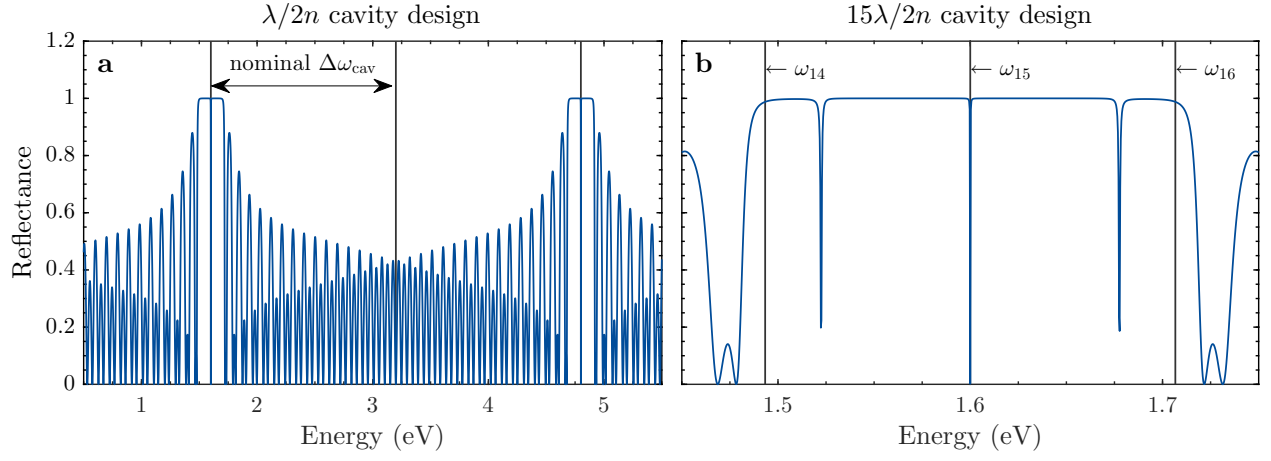


Figure 1.7: (a) The reflectance of a DBR microcavity, with 16 periods for both the top and bottom DBRs and a  $\lambda/2n$  cavity spacer, and the same material parameters as in Figure 1.6. The cavity and DBR layer thicknesses are all designed to work ideally at 1.6 eV. The vertical gray lines show the expected positions of the cavity modes based on the thickness of the cavity spacer. (b) Similar to (a) but with a  $15\lambda/2n$  cavity spacer, showing the overlap of the neighboring modes with the same DBR stop band.

## 1.3 EXCITONS

### 1.3.1 Excitons as hydrogen-like particles

In any semiconductor, the Fermi level lies within a relatively narrow energy gap between a filled electron energy band and an empty (at zero temperature) energy band. With an energy input large enough to cross the gap, an electron can jump from the filled valence band up into the empty conduction band, leaving a positively charge “hole” in the valence band. An exciton is formed by the Coulomb attraction between that conduction-band electron and the valence-band hole, resulting in a bound state analagous to hydrogen or positronium. Conveniently, the Schrödinger equation for bound states between an electron and a positive charge are very familiar.

Following Ref. [38], the energy levels of a hydrogen atom are given by

$$E_{\text{hyd}} = -\frac{\text{Ry}}{n^2} + \frac{\hbar^2 k^2}{2(m_e + m_p)}, \quad (1.27)$$

where  $m_e$  is the free electron mass and  $m_p$  is the proton mass. This contains the binding energy term based on the Rydberg (Ry), which is simply the ground-state binding energy for a hydrogen atom alone in space, and a center-of-mass kinetic energy term. The Rydberg is defined as

$$\text{Ry} = \frac{m_e e^4}{(4\pi)^2 2\hbar^2 \epsilon_0^2} \quad (1.28)$$

where  $e$  is the electron charge, and  $\epsilon_0$  is the permittivity of vacuum. An exciton is obviously not alone, but actually exists within a material. Thus, the same results can be used to describe the exciton by simply replacing  $\epsilon_0$  with the specific permittivity of the medium  $\epsilon$  (effectively reducing the Coulomb interaction), replacing the free electron and proton masses in the kinetic term with the conduction band electron and valence band hole masses  $m_c$  and  $m_v$ , and replacing the free electron mass in the binding energy term with the exciton reduced mass  $\mu_{\text{exc}} = m_c m_v / (m_c + m_v)$ . This gives the exciton version of Equation 1.27:

$$E_{\text{exc}} = E_{\text{gap}} - \frac{\mu_{\text{exc}} e^4}{(4\pi)^2 2\hbar^2 \epsilon^2 n^2} + \frac{\hbar^2 k^2}{2(m_c + m_v)}. \quad (1.29)$$

This equation assumes an isotropic dielectric with energy bands that can be perfectly described by effective masses, and that the exciton is alone in the material. These are obviously all approximations, but it is helpful in revealing some of the general characteristics of excitons. There are three important features to notice when comparing excitons to hydrogen atoms. The first is that there is an extra term  $E_{\text{gap}}$  which gives the bandgap of the material. Since excitons are fundamentally an excitation of valence band electrons, it makes sense to measure their energy above the unexcited state. The second is that the binding energy of the exciton is much lower, since the conduction band mass is typically significantly lower than the free electron mass (more on that in Section 1.3.2) and the permittivity in a solid material is significantly higher than in vacuum. For example, in GaAs, the ground state binding energy is 4.1 meV, compared to 13.6 eV for a hydrogen atom in vacuum.

The third feature of interest is the size of the excitons. For the ground state, the average size can be expressed as the Bohr radius:

$$a_{\text{B}} = \frac{4\pi\hbar^2\epsilon}{\mu_{\text{exc}}e^2}. \quad (1.30)$$

For the same reasons that the exciton binding energy is much lower than that of hydrogen, the exciton Bohr radius is much larger than for hydrogen (which also makes sense since the binding energy is lower). This leads to an important distinction between two types of excitons, Frenkel and Wannier-Mott varieties. Frenkel excitons have a Bohr radius less than the lattice spacing of the material, so that the electron and hole both sit at the same lattice site. Wannier-Mott excitons are just the opposite, with a large enough Bohr radius that they are spread out across multiple lattice sites.

One final important point in introducing excitons is that they have a finite lifetime. Since they are created by an excitation, they can just as easily be annihilated by simply reversing the process. This can take place whenever the electron and hole encounter each other in space (see page 74 of Ref. [38]). In the case of a *direct bandgap* semiconductor, the excitation and annihilation are given by simply absorbing and reemitting a photon. Indirect bandgaps require a second process to shift the momentum of one of the particles, such as a phonon collision. These differences, along with differences in size, lead to a wide variety of exciton lifetimes.

While excitons can be generally understood using the analogy with hydrogen, a better understanding of the exciton states requires consideration of the electronic band structures of the material in which they exist.

### 1.3.2 The electron band structure: $\mathbf{k} \cdot \mathbf{p}$ perturbation theory.

For solid state systems, the typical starting point is the Bloch theorem. In this section, I will outline the basics of that theorem, and introduce the  $\mathbf{k} \cdot \mathbf{p}$  method for analyzing the band energies near critical points, drawing from Refs. [38, 40]. The overall process followed by this discussion is outlined in Figure 1.8.

The Bloch theorem is basically that the electron wavefunctions within a periodic crystal lattice can be expressed as functions (called Bloch functions), which have the same periodicity as the lattice. Thus, for the general Schrödinger equation with a potential  $V(\mathbf{r})$  which is invariant upon translation by a lattice vector  $\mathbf{R}$  (i.e.  $\mathbf{r} \rightarrow \mathbf{r} + \mathbf{R}$ ), the solutions are given by the Bloch functions

$$\psi_{n\mathbf{k}}(\mathbf{r}) = e^{i\mathbf{k} \cdot \mathbf{r}} u_{n\mathbf{k}}(\mathbf{r}), \quad (1.31)$$

where the  $u_{n\mathbf{k}}$  functions are also invariant upon translation by a lattice vector, such that  $u_{n\mathbf{k}}(\mathbf{r}) = u_{n\mathbf{k}}(\mathbf{r} + \mathbf{R})$ . The subscript  $n$  gives the particular electron band of the function. The Schrödinger equation then for an electron in band  $n$  is given by

$$\begin{aligned} \left( \frac{p^2}{2m_e} + V(\mathbf{r}) \right) \psi_{n\mathbf{k}} &= E_n(\mathbf{k}) \psi_{n\mathbf{k}} \\ \Rightarrow \left( \frac{p^2}{2m_e} + V(\mathbf{r}) \right) e^{i\mathbf{k} \cdot \mathbf{r}} u_{n\mathbf{k}}(\mathbf{r}) &= E_n(\mathbf{k}) e^{i\mathbf{k} \cdot \mathbf{r}} u_{n\mathbf{k}}(\mathbf{r}) \\ \Rightarrow \left( \frac{p^2}{2m_e} + \frac{\hbar}{m_e} \mathbf{k} \cdot \mathbf{p} + \frac{\hbar^2 k^2}{2m_e} + V(\mathbf{r}) \right) u_{n\mathbf{k}}(\mathbf{r}) &= E_n(\mathbf{k}) u_{n\mathbf{k}}(\mathbf{r}). \end{aligned} \quad (1.32)$$

It is important to remember that  $\mathbf{k}$  is simply a number giving the wavevector of the electron, while  $\mathbf{p}$  is the quantum mechanical momentum operator. Since the first and last terms on the left-hand-side are just the original Hamiltonian of the system, we can identify the various

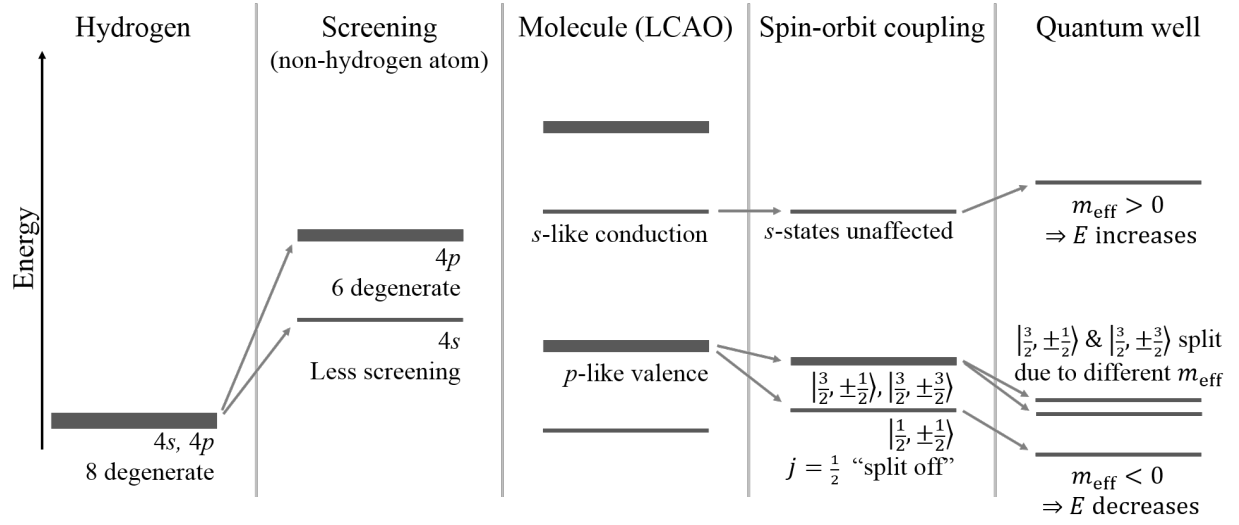


Figure 1.8: A sketch outlining the process of arriving at the heavy and light hole valence bands with an  $s$ -like conduction band in GaAs. The initial hydrogen  $s$ -type and  $p$ -type states are split in any atom with occupied lower bands due to screening. Then, in the case of a covalently bonded molecule like GaAs, the  $4s$  and  $4p$  states, which are at the valence level of both gallium and arsenic, are split nicely into the same total number of states (sixteen total), with the valence band now made up of six completely filled and degenerate  $p$ -like states. Spin-orbit coupling then has a stronger effect on the  $j = \frac{1}{2}$  state, causing a “split-off” band at lower energy. Quantum well confinement then splits the degenerate heavy hole and light hole states due to different effective mass (though the state that is higher depends on the chosen momentum direction within the lattice).

terms, following Ref. [38], as

$$\begin{aligned} H_0 &= \frac{p^2}{2m_e} + V(\mathbf{r}) \\ H_1 &= \frac{\hbar}{m_e} \mathbf{k} \cdot \mathbf{p} \\ H_2 &= \frac{\hbar^2 k^2}{2m_e}, \end{aligned} \tag{1.33}$$

and then apply the standard methods of time-independent perturbation theory where  $H_1$  and  $H_2$  are first- and second-order perturbations. For a discussion of the general method, see Refs. [40–42]. For bands with maxima and minima at the Brillouin zone center  $k = 0$ , the overall result near zone center for the energies is given by

$$E_n(\mathbf{k}) \approx E_n(0) + \frac{\hbar^2 k^2}{2m_e} + \frac{\hbar^2}{m_e^2} \sum_{n' \neq n} \frac{|\mathbf{k} \cdot \langle u_{n'0} | \mathbf{p} | u_{n0} \rangle|^2}{E_n(0) - E_{n'}(0)} \tag{1.34}$$

while the states are given by

$$|u_{n\mathbf{k}}\rangle \approx |u_{n0}\rangle + \frac{\hbar}{m_e} \sum_{n' \neq n} \frac{\mathbf{k} \cdot \langle u_{n'0} | \mathbf{p} | u_{n0} \rangle}{E_n(0) - E_{n'}(0)} |u_{n'0}\rangle. \tag{1.35}$$

Equation 1.34 is clearly still in the general form of a ground state energy with a kinetic term. The *effective mass approximation* takes advantage of this for isotropic bands, with the energy rewritten as [43]

$$E_n(\mathbf{k}) \approx E_n(0) + \frac{\hbar^2 k^2}{2m_{\text{eff}}} \tag{1.36}$$

where

$$\frac{1}{m_{\text{eff}}} = \frac{1}{m_e} + \frac{2}{m_e^2 k^2} \sum_{n' \neq n} \frac{|\mathbf{k} \cdot \langle u_{n'0} | \mathbf{p} | u_{n0} \rangle|^2}{E_n(0) - E_{n'}(0)}. \tag{1.37}$$

This result reveals a very interesting effect within the energy bands called *level repulsion*. From the above relation, it is clear that the effective mass of an electron in a given band is affected by the nearness of other bands. Other bands that are closer to the same energy give a smaller denominator in the final term and larger overall effect on the energy in Equation 1.34. If we consider just two bands  $n$  and  $n'$ , if  $n'$  is above  $n$  then the term will be negative, reducing energy of band  $n$  at  $k \neq 0$ . The opposite is true if band  $n'$  is below band  $n$ . Thus, the bands repel each other. For an electron conduction band located just above a negative mass valence band (typical for the  $k = 0$  point in a direct bandgap semiconductor

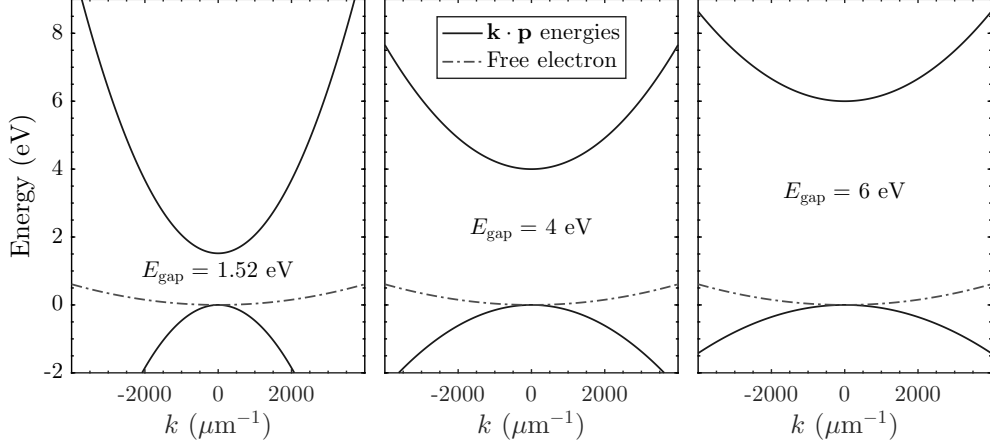


Figure 1.9: The  $\mathbf{k} \cdot \mathbf{p}$  theory results for two bands with the energies  $E_1(0) = 0$  and  $E_2(0) = E_{\text{gap}}$ , for three different gap energy values. The free electron dispersion is also plotted for comparison. As expected, the curvature of the upper band is increased, giving a smaller effective mass, while the lower band obtains a negative curvature and negative effective mass. Both masses approach the free electron mass as the gap is increased.

such as GaAs), the effect is to shift both bands away from each other. The upper band gets a larger curvature (i.e. smaller effective mass), while the lower band is bent in the negative direction, gaining a negative curvature (and negative effective mass) if the bands are close enough. This effect is shown in Figure 1.9 for a very simple system with two bands separated by a gap energy  $E_{\text{gap}}$ . The matrix elements used for the calculation were given by  $2|\langle u_{n'0} | p | u_{n0} \rangle|^2 / m_e = 20$  eV, which actually gives effective masses for the conduction band that are pretty close to the experimentally determined values in many of the common direct bandgap semiconductors (see Ref. [43], page 71). The energy gap value of 1.52 eV is close to the value for GaAs near 0 K [39, 43], though the situation in actual GaAs is very different from a simple two-band picture. The specific case of GaAs and similar materials requires a more careful treatment, which I will discuss in the next section.

### 1.3.3 The Luttinger-Kohn model

The material whose band structure is most important to my work is GaAs. Some of the basic properties of the band structure can be easily intuited by considering the states of a hydrogen atom (see e.g. Ref. [41]). For gallium, the outermost (most weakly bonded) electrons will have a principle quantum number  $n = 4$ , with two electrons in a  $4s$  state and one in a  $4p$  state. Arsenic also has two electrons in a  $4s$  state and then three in a  $4p$  state. In both cases, the  $4s$  states have significantly lower energy than the  $4p$  states due to a lower amount of screening: Since electrons in the  $s$  states have some probability of being found near the nucleus, they can partially bypass the electrons in inner states, feeling more of the full potential from the nucleus. Meanwhile the  $p$  state probability goes to zero at the nucleus, so they are more effectively screened by the electrons in the inner states.

When the two elements combine in a molecule, the overlap in space of their electronic wavefunctions gives rise to new states, which can be considered as a linear combination of the original states, with equal numbers of lower energy *bonding* states and higher energy *antibonding* states (See chapter 1 of Ref. [38] and chapter 2 Ref. [43] for nice discussions of this, based on the linear combination of atomic orbitals approximation.). In the case of GaAs, the eight total outer electrons fit nicely into the eight bonding states, while the antibonding states are empty. A tight-binding approach, which assumes the electron states are mostly unchanged from the unbound atomic states, works well for the bound valence states of a semiconductor. For the *zincblende* arrangement of GaAs, which is the lattice structure of the samples used throughout this work, the result is eight new bonding states including two  $s$ -like and six  $p$ -like states, and a lowest antibonding  $s$ -like state. The six  $p$ -like states are all degenerate at  $k = 0$  (when ignoring spin-orbit interactions, spin-spin interactions, etc.), which is simply due to the cubic symmetry of a *zincblende* lattice. In contrast to the *zincblende* structure of GaAs, in a material with a *wurtzite* structure there is no cubic symmetry so none of the three  $p$  states is degenerate with the other, even before considering spin-orbit interactions.

For the case of multiple degenerate bands in *zincblende*, there is a very useful approach using  $\mathbf{k} \cdot \mathbf{p}$  theory called the Luttinger-Kohn model [44]. This model is explained nicely in

Ref. [40], but I will again outline the major points. The general idea is to consider the six total  $p$ -bands as the primary set of states, while the conduction band (the lowest antibonding  $s$ -like state) and all other bands are considered only as a perturbation. Spin-orbit coupling is also explicitly included since it has a major effect on the character of these bands.

The six states are divided into three sets, each doubly degenerate because of the two possible electrons spins. In terms of total-angular momentum states  $|j, m_j\rangle$ , the heavy-hole is the  $|\frac{3}{2}, \pm\frac{3}{2}\rangle$  state, the light-hole is the  $|\frac{3}{2}, \pm\frac{1}{2}\rangle$  state, and the split-off band is the  $|\frac{1}{2}, \pm\frac{1}{2}\rangle$  state. The names describe their basic features: The split-off band is split to a lower energy due to spin-orbit interactions (see e.g. Ref. [41] chapter 11 or Ref. [42] chapter 20), while the heavy-hole has a generally higher effective mass than the light-hole.

An approximate value for the effective masses of the heavy-hole and light-hole states can be found using the Luttinger-Kohn Hamiltonian but ignoring the coupling to the split-off band. As worked out in Ref. [45] following the process and notation of Ref. [40], this results in a  $4 \times 4$  matrix Hamiltonian that couples the four heavy-hole and light-hole states. The resulting energies of the states are given by

$$E(\mathbf{k}) = -P \pm \sqrt{|Q|^2 + |R|^2 + |S|^2} \quad (1.38)$$

where

$$\begin{aligned} P &= \frac{\hbar^2 \gamma_1}{2m_e} (k_x^2 + k_y^2 + k_z^2) \\ Q &= \frac{\hbar^2 \gamma_2}{2m_e} (k_x^2 + k_y^2 - 2k_z^2) \\ R &= \frac{\hbar^2}{2m_e} (-\sqrt{3}\gamma_2(k_x^2 - k_y^2) + i2\sqrt{3}\gamma_3 k_x k_y) \\ S &= \frac{\hbar^2 \gamma_3}{m_e} \sqrt{3}(k_x - ik_y)k_z. \end{aligned} \quad (1.39)$$

The  $\gamma$ 's are experimentally-determined parameters, and are some of the more interesting results of the Luttinger-Kohn model since they account for the coupling to the other bands, including the conduction band, via perturbation. By using this phenomenologically informed approach (where the details are reduced to a few experimentally-determined parameters), the effect of the many other bands on the valence bands can be included in a simple way that also achieves good agreement with measurements. Their approximate values are  $\gamma_1 \approx 7$ ,  $\gamma_2 \approx 2$ ,

and  $\gamma_3 \approx 3$ , with a pretty wide variation across experiments (see Table II in Ref. [46]). As expected, since this solution is in the basis of the degenerate  $|\frac{3}{2}, \pm\frac{3}{2}\rangle$  and  $|\frac{3}{2}, \pm\frac{1}{2}\rangle$ , there are two eigenvalues. In the samples used for my work, the crystal growth direction (perpendicular to the plane of the sample wafer) was [001], corresponding to the  $z$ -direction. Using  $\mathbf{k} = k_z \hat{k}$ , this gives

$$E_{\pm} = -\frac{\hbar^2 \gamma_1}{2m_e} k_z^2 \pm \frac{\hbar^2 \gamma_2}{m_e} k_z^2 = -\frac{\hbar^2}{2m_e} k_z^2 (\gamma_1 \mp 2\gamma_2), \quad (1.40)$$

with the positive sign corresponding to the light hole and the negative sign to the heavy hole. This gives two effective masses, for the light hole (lh) and heavy hole (hh):

$$\begin{aligned} m_+ = m_{\text{hh}} &= -\frac{m_e}{\gamma_1 - 2\gamma_2} \\ m_- = m_{\text{lh}} &= -\frac{m_e}{\gamma_1 + 2\gamma_2}. \end{aligned} \quad (1.41)$$

Since the parameters  $\gamma_1$  and  $\gamma_2$  are both positive, electrons in the heavy-hole band are indeed heavier than those in the light-hole band along this direction. Their values are  $m_{\text{hh}} \approx -0.34m_e$  and  $m_{\text{lh}} \approx -0.09m_e$  [46]. This will be important when we consider confinement within a quantum well in Section 1.3.4. Finally, note that the mass of an electron in these bands is negative. This means that the positive hole states are actually positive mass, positive charge states that exist within this band. Thus the basic picture of an exciton that I gave in Section 1.3.1 still holds for these valence and conduction bands. The overall result for the electron dispersions is shown in Figure 1.10(a).

### 1.3.4 Excitons in quantum wells

A quantum well (QW) is simply a three-layer material stack: with a thin layer that supports lower energy exciton states between two “barrier” layers of material with higher energy states. Since the exciton energy (Equation 1.29) is strongly dependent on the bandgap of the material, quantum wells are typically realized by placing a lower bandgap layer between two higher bandgap layers. In principle, the offsets of the bands can lead to three different situations: confinement in both bands, confinement of the conduction band electron with a barrier for the valence band hole, or confinement of the valence band hole with a barrier for the conduction band electron. Both electrons and holes were confined in the samples used for

my work, which was realized by using GaAs ( $E_{\text{gap}} \simeq 1.5$  eV) for the middle layer and AlAs ( $E_{\text{gap}} \simeq 2.1$  eV) for the barrier layers. For these undoped semiconductors, the Fermi level falls somewhere within the energy bandgap, which allows for a simple approximate depiction of the energy band diagram in space, which is shown for the case of two nearby quantum wells in Figure 1.11.

In order to reasonably qualify as a quantum well, the lower energy layer must be narrow enough to produce a non-negligible quantum confinement. A full treatment requires considering the individual confinement potentials on both the electron and hole states, solving the Schrödinger equation (see Ref. [39] chapter 4)

$$\left( -\frac{\hbar^2}{2m_e} \nabla_e^2 - \frac{\hbar^2}{2m_h} \nabla_h^2 + V_e(z_e) + V_h(z_h) - \frac{e^2}{4\pi\epsilon|\mathbf{r}_e - \mathbf{r}_h|} \right) \Psi = E\Psi, \quad (1.42)$$

where the confinement is in the  $z$ -direction. A simplification can be made in the case of quantum wells that are thinner than the exciton Bohr radius  $a_B$ , which confines the electron and hole states individually. This allows factorization of the wave function, and results in a solution analogous to that of a 2D hydrogen atom. For quantum wells thicker than the Bohr radius, it is more reasonable to consider the exciton trapped as a whole, rather than the individual electron and hole components. Since the Bohr radius in GaAs is about 15 nm, and the quantum well layers used for my work were about 7 nm thick, the case of thinner quantum wells is more relevant.

An important result of this treatment [39] is that the exciton binding  $E_B$  energy in the ideal 2D case is increased by a factor of four ( $E_B^{2D} = 4E_B$ ). For non-ideal quantum wells, this situation is not fully realized, but the effect is still present. Instead, there is a continuous increase in the binding energy as the quantum wells are made thinner until the binding energy again decreases at a quantum well thickness of about  $0.5a_B$  due to tunnelling of the electrons and holes into the barrier layers.

As mentioned above, a careful treatment requires considering the confinement on both the electron and hole components in the presence of a finite square-well potential. Technically, the well is also not exactly square, since there must be a length over which the bend-bending occurs (see, e.g. chapter 29 of Ref. [47]). However, an approximate picture of a particle in an

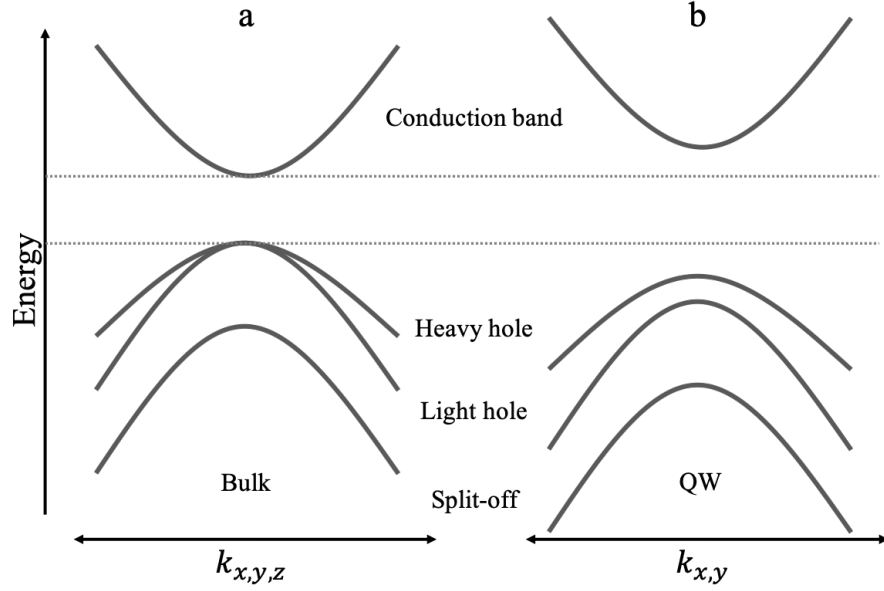


Figure 1.10: A sketch of the energy dispersions of the electron bands in GaAs (a) in bulk material near  $k = 0$  for  $k$  along one of the lattice directions and (b) in a quantum well (QW) near  $k_{\parallel} = 0$ , where  $k_{\parallel}$  is in a direction parallel to the plane of the quantum well (perpendicular to the  $z$ -direction), along one of the lattice directions within the plane of the quantum well. This is obviously not to scale, as the conduction-valence bandgap is about 1.5 eV while the splitting of the split-off band is about 0.34 eV (see table 2.24 in Ref. [43]). As discussed in the text, the quantum well confinement shifts the energies of all of the bands, introducing a splitting between the heavy and light holes.

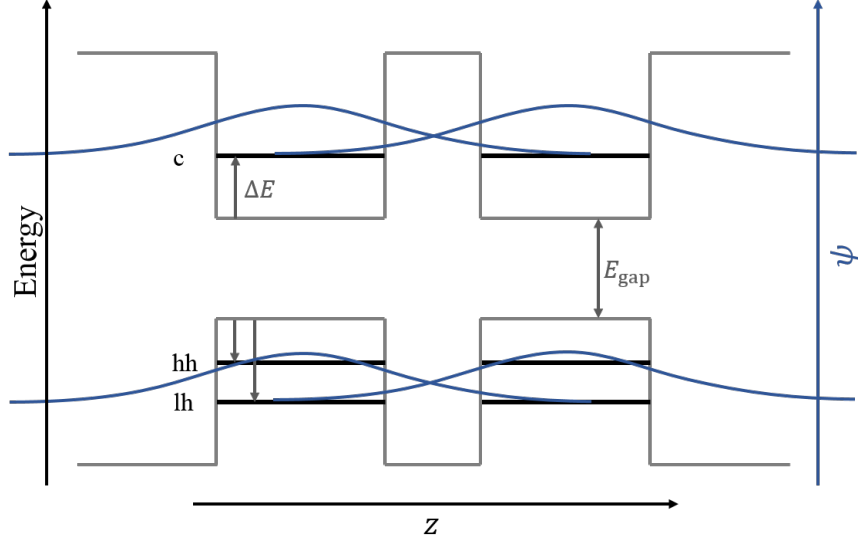


Figure 1.11: A sketch of the effect of quantum well confinement on the exciton states. The horizontal axis gives the position along the growth direction (perpendicular to the individual layers), while the vertical axis shows the approximate electron energies at  $k_{\parallel} = 0$ . Two undoped semiconductors with different band gap energies ( $E_{\text{gap}}$ ) are used to make two quantum wells separated by a narrow barrier. The original energy of the conduction band (c) is shifted higher by energy  $\Delta E$  due to the confinement and the positive effective mass, while the valence bands are shifted to lower energy due to negative effective masses. The light hole (lh) and heavy hole (hh) are split because of their different effective masses. The overall trend of the particle ground state wavefunctions is shown by the blue curves, with some penetration of the bound states into the neighboring quantum wells. A careful treatment would lead to two new states split in energy and evenly distributed between the two wells, though the overall effects of the confinement would be the same.

infinite square potential reveals the most important effects. For the infinite square potential, the original energies are shifted by the confinement energy

$$E_{\text{QW}} = \frac{\pi^2 \hbar^2 n^2}{2mL^2} \quad (1.43)$$

where  $m$  is the mass of the particle,  $L$  is the width of the potential well, and  $n$  is an integer. Since the confinement is inversely proportional to the particle mass, lighter particles will be shifted in energy by a larger amount. As discussed in Section 1.3.3, the mass of the light hole along the  $z$ -direction (which is the confinement direction in the case of my samples) is significantly lower than the heavy hole, resulting in a larger shift. Since the hole states both have negative mass, they are shifted to lower energy, while the conduction band is shifted to higher energy [see Figures 1.11 and 1.10(b)]. The result is that the light holes and heavy holes are no longer degenerate, with the lowest gap between the heavy hole and conduction band electron. Thus, QW excitons are preferentially made up of  $|\frac{1}{2}, \pm\frac{1}{2}\rangle$  conduction band electrons and  $|\frac{3}{2}, \pm\frac{3}{2}\rangle$  heavy holes. The overall situation is sketched out in Figure 1.11, but in the case of two nearby quantum wells. While the energy level trends of the infinite square well mostly apply, it is important to note that, in the finite case, there is some penetration into the barrier regions. The ground state wavefunction will be a symmetric cosine function within the well but an exponentially damped function outside (see page 97ff of Ref. [42]). As shown by the blue curves in Figure 1.11, this results overlap between the wavefunctions in neighboring wells. If the overlap is significant, the situation is referred to as “coupled quantum wells.” Any particle in such a system is shared between the wells due to tunneling. This has the advantage of producing bound states with a particular energy characteristic of the well width while also allowing a greater spatial extent of the bound states along the growth direction ( $z$ -direction in this case).

In summary, determining the exciton energy dispersion is actually a pretty complicated process. The steps followed in this section are outlined in Figure 1.8, showing the overall process of arriving at the heavy hole  $|\frac{3}{2}, \pm\frac{3}{2}\rangle$  valence band with the  $s$ -like  $|\frac{1}{2}, \pm\frac{1}{2}\rangle$  conduction band for a GaAs quantum well. As I will show in Section 1.4, the overall result is that we can actually mostly ignore details of the band, with the exception of the spin states, since

the exciton mass is so large that its energy dispersion can mostly be treated as flat (having negligible curvature) on the scales of the polaritons.

## 1.4 POLARITONS

Now that excitons and cavity photons have been explained independently, I will show how the coupling between them actually leads to the new *exciton-polariton* states.

### 1.4.1 Exciton-photon coupling

There are two essentially equivalent intuitive pictures of the exciton-photon interaction. The first is that, since an exciton is made up of a negative charge electron and a positive charge hole with some spatial separation, it will have some dipole moment. Photons are simply (classically) an oscillating electromagnetic field, so there will be coupling between the electric field of the photon and the dipole moment of the exciton. The second way of thinking about it is that the exciton requires an excitation of a valence band electron into an excited state (the exciton state). One way to do this is to simply absorb a photon with enough energy to kick an electron directly from the valence band into the conduction band. At a basic level, the two pictures are the same, with the first more classical and the second more quantum mechanical.

Fundamentally, the problem is a specific application of light-matter coupling, which has been worked out in many places (e.g. Refs. [12, 38, 42, 48]). I will again produce the result, with the specific application to an exciton system. The Hamiltonian for a single particle with charge  $q$  in an electromagnetic (EM) field can be written as

$$\mathcal{H} = \frac{[\mathbf{p} - q\mathbf{A}(\mathbf{x}, t)]^2}{2m} + q\phi(\mathbf{x}, t) + U(\mathbf{x}, t), \quad (1.44)$$

where  $\mathbf{A}$  is the electromagnetic vector potential field,  $\phi$  is the scalar potential field,  $U$  contains all other potentials,  $\mathbf{p}$  is the quantum mechanical momentum operator, and  $m$  is the particle mass (see chapter 13 of Baym [42] for a general discussion, especially regarding gauges).

Following Ref. [38], it is very convenient to work in the Coulomb gauge, in which  $\nabla \cdot \mathbf{A} = 0$ . This gauge allows  $\mathbf{A}$  and  $\mathbf{p}$  to commute, since, using a general function  $\psi$ ,

$$\mathbf{p} \cdot \mathbf{A}\psi = -i\hbar(\nabla \cdot \mathbf{A})\psi - i\hbar\mathbf{A} \cdot \nabla\psi = \mathbf{A} \cdot \mathbf{p}\psi. \quad (1.45)$$

The first term in Equation 1.44 can then be expanded to become  $(\mathbf{p}^2 - q\mathbf{p} \cdot \mathbf{A} - q\mathbf{A} \cdot \mathbf{p} + q^2\mathbf{A}^2)/2m \approx (\mathbf{p}^2 - 2q\mathbf{A} \cdot \mathbf{p})/2m$ , where the approximation is made to ignore the term that goes as  $\mathbf{A}^2$ , which holds for low amplitude fields. The Hamiltonian for a single particle in a electromagnetic field is then

$$\mathcal{H} = \left( \frac{\mathbf{p}^2}{2m} + U(\mathbf{x}, t) \right) - \frac{q\mathbf{A}(\mathbf{x}) \cdot \mathbf{p}}{m} = \mathcal{H}_0 + \mathcal{H}_{\text{int}}, \quad (1.46)$$

where the term in parentheses is the normal Hamiltonian for a free particle ( $\mathcal{H}_0$ ). The scalar field  $\phi(\mathbf{x}, t)$  is zero everywhere since I am dealing with free EM waves and no local charge density. The EM field is also taken to be constant in time, allowing me to drop the explicit time dependence. The vector potential field operator is given by

$$\mathbf{A}(\mathbf{x}) = \sum_{\mathbf{q}\lambda} \sqrt{\frac{\hbar}{2\epsilon_0 V \omega_{\mathbf{q}}}} \boldsymbol{\epsilon}_{\mathbf{q}\lambda} (a_{\mathbf{q}} e^{i\mathbf{q} \cdot \mathbf{x}} + a_{\mathbf{q}}^\dagger e^{-i\mathbf{q} \cdot \mathbf{x}}) = \sum_{\mathbf{q}\lambda} \frac{A_{\mathbf{q}}}{\sqrt{V}} \boldsymbol{\epsilon}_{\mathbf{q}\lambda} (a_{\mathbf{q}} e^{i\mathbf{q} \cdot \mathbf{x}} + a_{\mathbf{q}}^\dagger e^{-i\mathbf{q} \cdot \mathbf{x}}), \quad (1.47)$$

where the sum is over all wavevectors  $\mathbf{q}$  and polarization states  $\lambda$ ,  $\boldsymbol{\epsilon}_{\mathbf{q}\lambda}$  is the polarization unit vector, and  $a_{\mathbf{q}}$  and  $a_{\mathbf{q}}^\dagger$  are the annihilation and creation operators for photons, giving the quantization of the field. The volume  $V$  is the total volume of the region of interest. I have also simplified the expression by defining the term  $A_{\mathbf{q}} \equiv \sqrt{\hbar/(2\epsilon_0 \omega_{\mathbf{q}})}$ . The Hamiltonian in Equation 1.46 is only for a single particle. In the case of excitons, we are interested in processes where the interaction leads to a transition between bands. To describe a material, this Hamiltonian needs to be integrated over all space, looking for the probability of a transition from band  $n$  to band  $n'$ :

$$H_{n,n'} = \int d^3\mathbf{x} \Psi_{n'}^\dagger(\mathbf{x}) \mathcal{H}_{\text{int}} \Psi_n(\mathbf{x}). \quad (1.48)$$

For electrons in a solid the field operators are best described using the Bloch functions  $\psi_{n\mathbf{k}}(\mathbf{x}) = e^{i\mathbf{k} \cdot \mathbf{x}} u_{n\mathbf{k}}(\mathbf{x})$ , which were introduced in Section 1.3.2. These are then given by

$$\Psi_n(\mathbf{x}) = \sum_{\mathbf{k}} \frac{e^{i\mathbf{k} \cdot \mathbf{x}}}{\sqrt{V}} u_{n\mathbf{k}}(\mathbf{x}) b_{n\mathbf{k}}. \quad (1.49)$$

The volume  $V$  is again the total volume of the region of interest, included for normalization and usually much larger than a unit cell for the periodic lattice. The sum is over all wavevectors  $\mathbf{k}$ ,  $b_{n\mathbf{k}}$  is the annihilation operator of an electron in band  $n$ , and its hermitian conjugate is the creation operator. The interaction part of the Hamiltonian (now specifically applied to an electron) is now

$$H_{n,n'} = \sum_{\mathbf{k}, \mathbf{k}', \mathbf{q}, \lambda} \int d^3\mathbf{x} \left( \frac{e^{-i\mathbf{k}' \cdot \mathbf{x}}}{\sqrt{V}} u_{n'\mathbf{k}'}^\dagger(\mathbf{x}) b_{n'\mathbf{k}'}^\dagger \right) \times -\frac{e}{m} \frac{A_{\mathbf{q}}}{\sqrt{V}} (a_{\mathbf{q}} e^{i\mathbf{q} \cdot \mathbf{x}} + a_{\mathbf{q}}^\dagger e^{-i\mathbf{q} \cdot \mathbf{x}}) \boldsymbol{\epsilon}_{\mathbf{q}\lambda} \cdot \mathbf{p} \left( \frac{e^{i\mathbf{k} \cdot \mathbf{x}}}{\sqrt{V}} u_{n\mathbf{k}}(\mathbf{x}) b_{n\mathbf{k}} \right). \quad (1.50)$$

The momentum operator  $\mathbf{p}$  does not commute with  $\mathbf{x}$ . However, if we again consider a general function  $\psi$ , we get  $\mathbf{p}e^{i\mathbf{k} \cdot \mathbf{x}}\psi = -i\hbar(\mathbf{k})e^{i\mathbf{k} \cdot \mathbf{x}}\psi - i\hbar e^{i\mathbf{k} \cdot \mathbf{x}}\nabla\psi = -i\hbar(\mathbf{k})e^{i\mathbf{k} \cdot \mathbf{x}}\psi + e^{i\mathbf{k} \cdot \mathbf{x}}\mathbf{p}\psi$ . We can see the commutator  $[\mathbf{p}, e^{i\mathbf{k} \cdot \mathbf{x}}]$  gives an extra term with no momentum operator. Ignoring that term for now, it is helpful to consolidate the exponentials in the other terms, giving

$$H_{n,n'} = -\frac{e}{m} \sum_{\mathbf{k}, \mathbf{k}', \mathbf{q}, \lambda} \frac{A_{\mathbf{q}}}{\sqrt{V}} \int d^3\mathbf{x} \left( \frac{e^{i(-\mathbf{k}' + \mathbf{q} + \mathbf{k}) \cdot \mathbf{x}}}{V} b_{n'\mathbf{k}'}^\dagger a_{\mathbf{q}\lambda} b_{n\mathbf{k}} + \frac{e^{i(-\mathbf{k}' - \mathbf{q} + \mathbf{k}) \cdot \mathbf{x}}}{V} b_{n'\mathbf{k}'}^\dagger a_{\mathbf{q}\lambda}^\dagger b_{n\mathbf{k}} \right) \times u_{n'\mathbf{k}'}^\dagger(\mathbf{x}) (\boldsymbol{\epsilon}_{\mathbf{q}\lambda} \cdot \mathbf{p}) u_{n\mathbf{k}}(\mathbf{x}). \quad (1.51)$$

This can be simplified by making the approximation that  $k$ ,  $k'$ , and  $q$  are all very small compared to the boundaries of the first Brillouin zone. This means that the wavelengths of the fields are much longer than the unit cell size, and they therefore change very little across a particular unit cell. The contribution for each unit cell can then be taken as a single value, given by setting  $\mathbf{x}$  to the lattice vector  $\mathbf{R}$  of the cell. This allows the exponentials to be pulled out of the integral. Meanwhile, the  $u_{n\mathbf{k}}$  functions are, by definition, invariant under a change from one unit cell to another, so the remaining integral is identical at each unit cell. The Hamiltonian can then be written as

$$H_{n,n'} = -\frac{e}{m} \sum_{\mathbf{k}, \mathbf{k}', \mathbf{q}, \lambda} \frac{A_{\mathbf{q}}}{\sqrt{V}} \frac{1}{N} \sum_{\mathbf{R}} \left( e^{i(-\mathbf{k}' + \mathbf{q} + \mathbf{k}) \cdot \mathbf{R}} b_{n'\mathbf{k}'}^\dagger a_{\mathbf{q}\lambda} b_{n\mathbf{k}} + e^{i(-\mathbf{k}' - \mathbf{q} + \mathbf{k}) \cdot \mathbf{R}} b_{n'\mathbf{k}'}^\dagger a_{\mathbf{q}\lambda}^\dagger b_{n\mathbf{k}} \right) \times \frac{1}{V_{\text{cell}}} \int_{\text{cell}} d^3\mathbf{x} u_{n'\mathbf{k}'}^\dagger(\mathbf{x}) (\boldsymbol{\epsilon}_{\mathbf{q}\lambda} \cdot \mathbf{p}) u_{n\mathbf{k}}(\mathbf{x}), \quad (1.52)$$

where  $N$  is the total number of lattice sites in the original volume  $V$ , and  $V_{\text{cell}}$  is the volume of a single unit cell of the lattice. The sum over  $\mathbf{R}$  turns the exponential terms into Kronecker deltas via the identity  $(1/N) \sum_{\mathbf{R}} e^{i(\mathbf{k}-\mathbf{k}')\cdot\mathbf{R}} = \delta_{\mathbf{k},\mathbf{k}'}$ . It is now clear why the terms without a  $\mathbf{p}$  dependence, mentioned above, can be ignored: the  $\epsilon_{\mathbf{q}\lambda}$  can actually be pulled out of the integral since it is not dependent on position, and without the momentum operator to give transitions between  $n$  and  $n'$ , the orthogonality of the  $u_{n\mathbf{k}}$  functions will cause those terms to be zero for transitions between bands when  $\mathbf{k} = \mathbf{k}'$ . The Kronecker deltas nearly fulfill this condition by setting  $\mathbf{k}' = \mathbf{k} \pm \mathbf{q}$ . It is then necessary to use the approximation that the photon momentum  $\mathbf{q}$  is very small compared to the first Brillouin zone of the electron bands, such that  $\mathbf{k} \pm \mathbf{q} \approx \mathbf{k}$ , which are vertical transitions between the bands. This is a very good approximation due to the very steep photon dispersion compared to those of the electrons, whether within a cavity or not (compare the scales of Figure 1.2 and Figure 1.9). Thus, the terms without the momentum operator  $\mathbf{p}$  are negligible. Using the shorthand notation  $\langle \mathbf{p} \rangle_{\mathbf{k}n'n} = (1/V_{\text{cell}}) \int_{\text{cell}} d^3\mathbf{x} u_{n'\mathbf{k}}^\dagger(\mathbf{x}) \mathbf{p} u_{n\mathbf{k}}(\mathbf{x})$ , and not making the approximation  $\mathbf{k} \pm \mathbf{q} \approx \mathbf{k}$  for the electron creation operators for now, the Hamiltonian for a transition between two bands becomes

$$H_{n,n'} = -\frac{e}{m} \sum_{\mathbf{k},\mathbf{q},\lambda} \frac{A_{\mathbf{q}}}{\sqrt{V}} \left( b_{n',\mathbf{k}+\mathbf{q}}^\dagger a_{\mathbf{q}\lambda} b_{n\mathbf{k}} + b_{n',\mathbf{k}-\mathbf{q}}^\dagger a_{\mathbf{q}\lambda}^\dagger b_{n\mathbf{k}} \right) \epsilon_{\mathbf{q}\lambda} \cdot \langle \mathbf{p} \rangle_{\mathbf{k}n'n}. \quad (1.53)$$

The total Hamiltonian will include a sum of Equation 1.53 over all transitions. For considering an exciton system, we are really only interested in direct transitions between the valence and conduction bands. Thus, the total interaction Hamiltonian is

$$\begin{aligned} H_{\text{int}} &= H_{c,v} + H_{v,c} \\ &= -\frac{e}{m} \sum_{\mathbf{k},\mathbf{q},\lambda} \frac{A_{\mathbf{q}}}{\sqrt{V}} \left[ \left( b_{c,\mathbf{k}+\mathbf{q}}^\dagger a_{\mathbf{q}\lambda} b_{v\mathbf{k}} + b_{c,\mathbf{k}-\mathbf{q}}^\dagger a_{\mathbf{q}\lambda}^\dagger b_{v\mathbf{k}} \right) \epsilon_{\mathbf{q}\lambda} \cdot \langle \mathbf{p} \rangle_{\mathbf{k}cv} \right. \\ &\quad \left. + \left( b_{v,\mathbf{k}+\mathbf{q}}^\dagger a_{\mathbf{q}\lambda} b_{c\mathbf{k}} + b_{v,\mathbf{k}-\mathbf{q}}^\dagger a_{\mathbf{q}\lambda}^\dagger b_{c\mathbf{k}} \right) \epsilon_{\mathbf{q}\lambda} \cdot \langle \mathbf{p} \rangle_{\mathbf{k}vc} \right]. \end{aligned} \quad (1.54)$$

Two of the four terms in this equation can be eliminated via the *rotating wave approximation*. First, I will put this Hamiltonian into the interaction picture, giving it the time dependence  $H_{\text{int}}(t) = e^{iH_0 t/\hbar} H_{\text{int}} e^{-iH_0 t/\hbar}$ , where  $H_0$  is Hamiltonian without interactions. In this picture,

states of definite numbers (Fock states) of conduction electrons, valence electrons, and photons are eigenstates of  $H_0$ , such that  $e^{iH_0t/\hbar} |N_c, N_v, N_p\rangle = e^{i(N_c E_c + N_v E_v + N_p E_p)t/\hbar} |N_c, N_v, N_p\rangle$ , where the  $N$ 's refer to particle numbers, and subscripts c, v, and p refer to the conduction band, valence band, and photon, respectively. I have also made the simplification that the electron bands are flat, such that the energies are the same for the many  $\mathbf{k}$ -states within the bands, though this is not necessary for this approximation. The populations  $N_c$ ,  $N_v$ , and  $N_p$  are really the total populations, not the populations of specific  $\mathbf{k}$ -states. A general matrix element for  $H_{\text{int}}(t)$  will then look like

$$\begin{aligned} & \langle N'_c, N'_v, N'_p | H_{\text{int}}(t) | N_c, N_v, N_p \rangle \\ &= \langle N'_c, N'_v, N'_p | H_{\text{int}} | N_c, N_v, N_p \rangle \exp \left\{ (it/\hbar) \left[ (N'_c - N_c) E_c + (N'_v - N_v) E_v + (N'_p - N_p) E_p \right] \right\}. \end{aligned} \quad (1.55)$$

Clearly, only some transitions are allowed by the creation and annihilation operators in Equation 1.54, and these will give values to the  $N' - N$  terms of either 1 or -1. Since the electrons are fermions, their populations can only be 1 or 0 for a particular  $\mathbf{k}$ -state, and terms that violate this would simply be set to zero. Using the gap energy  $E_g = E_c - E_v$ , the time dependencies of each of the terms, in the order shown in Equation 1.54, are  $e^{i(E_g - E_p)t/\hbar}$ ,  $e^{i(E_g + E_p)t/\hbar}$ ,  $e^{-i(E_g + E_p)t/\hbar}$ ,  $e^{-i(E_g - E_p)t/\hbar}$ . The rotating wave approximation basically says that, for near-resonant excitation ( $E_g \approx E_p$ ), the second and third terms can be ignored because they oscillate very quickly (at a frequency of  $\approx 2E_g/\hbar$ ) on the time scales of a measurement, averaging to zero. Meanwhile, the first and fourth terms oscillate very slowly, with a frequency given by the difference in energies, which is small. This approximation should actually be very good, since for excitations at about 1.6 eV, the neglected terms would oscillate with a period of about a femtosecond. The full interaction Hamiltonian is then given by

$$H_{\text{int}} = -\frac{e}{m} \sum_{\mathbf{k}, \mathbf{q}, \lambda} \frac{A_{\mathbf{q}}}{\sqrt{V}} \boldsymbol{\epsilon}_{\mathbf{q}\lambda} \cdot (b_{c, \mathbf{k}+\mathbf{q}}^\dagger a_{\mathbf{q}\lambda} b_{v\mathbf{k}} \langle \mathbf{p} \rangle_{\mathbf{k}cv} + b_{v, \mathbf{k}-\mathbf{q}}^\dagger a_{\mathbf{q}\lambda}^\dagger b_{c\mathbf{k}} \langle \mathbf{p} \rangle_{\mathbf{k}vc}). \quad (1.56)$$

Since we are interested in excitons specifically, it makes sense to put this Hamiltonian in terms of exciton operators instead of the valence and conduction band operators. From

Ref. [12], the creation operator for a Wannier exciton in terms of valence and conduction band creation and annihilation operators is

$$X_{\mathbf{K}n}^\dagger = \sum_{\mathbf{k}} \phi_n(\mathbf{l}) b_{c,\mathbf{k}+\mathbf{K}}^\dagger b_{v,\mathbf{k}}, \quad (1.57)$$

where the envelope function  $\phi_n(\mathbf{l})$  is the Fourier transform of the hydrogen-like wavefunction  $\psi_n(\mathbf{x})$  of the excitons with band index  $n$  (see Section 1.3.1), such that

$$\phi_n(\mathbf{l}) = (1/\sqrt{V}) \int d^3\mathbf{x} \psi_n(\mathbf{x}) \exp(-i\mathbf{l} \cdot \mathbf{x}). \quad (1.58)$$

The center of mass momentum of the exciton is given by  $\mathbf{K}$ , and  $\mathbf{l} = [m_h/(m_h+m_e)](\mathbf{k} + \mathbf{K}) + [m_e/(m_h+m_e)]\mathbf{k}$ , where  $m_e$  and  $m_h$  are the conduction band electron and valence band hole masses, respectively, for isotropic bands. For the more general case, the masses could also be considered as being dependent on the directions of  $\mathbf{k}$  and  $\mathbf{k} + \mathbf{K}$ . Looking at Equation 1.56 and comparing to Equation 1.57, it is clear that this interaction (which up to this point has been generally applicable to two electron bands separated by a gap) could be specifically applied to create an exciton with momentum  $\mathbf{q}$  by destroying a photon with the same momentum, or vice versa. We now need to apply the approximation  $\mathbf{k} \pm \mathbf{q} \approx \mathbf{k}$  to the electron creation operators, which was explained above (see the discussion following Equation 1.52, and also leads to the simplification  $\mathbf{l} \approx \mathbf{k}$ . Now, if we consider the matrix element of  $H_{\text{int}}$  that couples an initial state with one photon with a final state with a single exciton with nearly zero momentum, it will have the form

$$\begin{aligned} \langle X_{0n} | H_{\text{int}} | 0 \rangle &= \sum_{\mathbf{k}'} \left[ \langle 0 | \phi_n(\mathbf{k}') b_{v,\mathbf{k}'}^\dagger b_{c,\mathbf{k}'} \right] \left( \frac{-e}{m} \right) \sum_{\mathbf{k}, \mathbf{q}, \lambda} \frac{A_{\mathbf{q}}}{\sqrt{V}} \boldsymbol{\epsilon}_{\mathbf{q}\lambda} \cdot \langle \mathbf{p} \rangle_{\mathbf{k}c v} b_{c,\mathbf{k}}^\dagger a_{\mathbf{q}\lambda} b_{v\mathbf{k}} | 0 \rangle \\ &= -\frac{e}{m} \sum_{\mathbf{k}, \mathbf{q}, \lambda} \frac{A_{\mathbf{q}}}{\sqrt{V}} \phi_n(\mathbf{k}) \boldsymbol{\epsilon}_{\mathbf{q}\lambda} \cdot \frac{1}{V_{\text{cell}}} \int_{\text{cell}} d^3\mathbf{x} u_{c\mathbf{k}}^\dagger(\mathbf{x}) \mathbf{p} u_{v\mathbf{k}}(\mathbf{x}), \end{aligned} \quad (1.59)$$

where I have written out the integral  $\langle \mathbf{p} \rangle_{\mathbf{k}c v}$  explicitly. The creation and annihilation operators, along with the orthogonality of the states, selects only those states with  $\mathbf{k} = \mathbf{k}'$ , and the anticommutation relation  $\{b_{c,\mathbf{k}}, b_{c,\mathbf{k}}^\dagger\} = 1$  gives a non-vanishing term for an initial state with zero conduction band electrons at wavevector  $\mathbf{k}$  when putting the operators in normal order. Another approximation can be made using  $\mathbf{k} \cdot \mathbf{p}$  perturbation theory, again closely following Ref. [12]. To first order, with small  $\mathbf{k}$ , the periodic  $u_{n\mathbf{k}}$  functions are simply

approximated by  $u_{n\mathbf{k}}(\mathbf{x}) \approx u_{n0}(\mathbf{x})$ . This causes only  $\phi_n$  to depend on  $\mathbf{k}$ , and the matrix element in Equation 1.59 then becomes

$$\begin{aligned}
\langle X_{0n} | H_{\text{int}} | 0 \rangle &= -\frac{e}{m} \sum_{\mathbf{q}, \lambda} \frac{A_{\mathbf{q}}}{\sqrt{V}} \left( \sum_{\mathbf{k}} \frac{1}{\sqrt{V}} \int d^3 \mathbf{x}' \psi_n(\mathbf{x}') e^{-i\mathbf{k} \cdot \mathbf{x}'} \right) \boldsymbol{\epsilon}_{\mathbf{q}\lambda} \cdot \frac{1}{V_{\text{cell}}} \int_{\text{cell}} d^3 \mathbf{x} u_{c0}^\dagger(\mathbf{x}) \mathbf{p} u_{v0}(\mathbf{x}) \\
&= -\frac{e}{m} \sum_{\mathbf{q}, \lambda} \frac{A_{\mathbf{q}}}{\sqrt{V}} \sqrt{V} \psi_n(0) \boldsymbol{\epsilon}_{\mathbf{q}\lambda} \cdot \frac{1}{V_{\text{cell}}} \int_{\text{cell}} d^3 \mathbf{x} u_{c0}^\dagger(\mathbf{x}) \mathbf{p} u_{v0}(\mathbf{x}) \\
&= -\frac{e}{m} \sum_{\mathbf{q}, \lambda} A_{\mathbf{q}} \psi_n(0) \boldsymbol{\epsilon}_{\mathbf{q}\lambda} \cdot \langle \mathbf{p} \rangle_{0cv}
\end{aligned} \tag{1.60}$$

where I have taken advantage of the  $\delta$ -function identity within the parentheses:  $\sum_{\mathbf{k}} e^{-i\mathbf{k} \cdot \mathbf{x}'} = V\delta(\mathbf{x}')$ . As long as the matrix element  $\boldsymbol{\epsilon}_{\mathbf{q}\lambda} \cdot \langle \mathbf{p} \rangle_{0cv}$  is non-zero, the higher order terms in the  $\mathbf{k} \cdot \mathbf{p}$  perturbation expansion are negligible. A very similar result can be found for the matrix elements  $\langle 0 | H_{\text{int}} | X_{0n} \rangle$ . Since I am concerned with the ground state of the exciton, which is the  $1s$  hydrogen-like state, the interaction Hamiltonian in Equation 1.56 can now be rewritten as

$$\begin{aligned}
H_{\text{int}} &= -\frac{e}{m} \psi_{1s}(0) \sum_{\mathbf{q}, \lambda} A_{\mathbf{q}} \boldsymbol{\epsilon}_{\mathbf{q}\lambda} \cdot (\langle \mathbf{p} \rangle_{0cv} X_{\mathbf{q},1s}^\dagger a_{\mathbf{q}\lambda} + \langle \mathbf{p} \rangle_{0vc} a_{\mathbf{q}\lambda}^\dagger X_{\mathbf{q},1s}) \\
&= -\frac{e}{m} \frac{1}{\sqrt{\pi a_B^3}} \sum_{\mathbf{q}, \lambda} A_{\mathbf{q}} \boldsymbol{\epsilon}_{\mathbf{q}\lambda} \cdot (\langle \mathbf{p} \rangle_{0cv} X_{\mathbf{q},1s}^\dagger a_{\mathbf{q}\lambda} + \langle \mathbf{p} \rangle_{0vc} a_{\mathbf{q}\lambda}^\dagger X_{\mathbf{q},1s}) \\
&= \sum_{\mathbf{q}, \lambda} (g_0(\mathbf{q}) X_{\mathbf{q},1s}^\dagger a_{\mathbf{q}\lambda} + g_0^*(\mathbf{q}) a_{\mathbf{q}\lambda}^\dagger X_{\mathbf{q},1s}).
\end{aligned} \tag{1.61}$$

The Bohr radius  $a_B$  in this case is the exciton Bohr radius which I discussed in Section 1.3.1. Thus, the interaction between the ground exciton state and the photons can be reduced to a simple coupling between the two states with a coupling strength given by  $g_0(\mathbf{q}) = (-e/m)(1/\sqrt{\pi a_B^3}) A_{\mathbf{q}} \boldsymbol{\epsilon}_{\mathbf{q}\lambda} \cdot \langle \mathbf{p} \rangle_{0cv} = (-e/m)(\sqrt{\hbar}/\sqrt{2\pi\epsilon_0 a_B^3 \omega_{\mathbf{q}}}) \boldsymbol{\epsilon}_{\mathbf{q}\lambda} \cdot \langle \mathbf{p} \rangle_{0cv}$ , where in the last step I have substituted the original values back in for  $A_{\mathbf{q}}$ . In Section 1.4.2, I will show that the full Hamiltonian with this interaction term can be diagonalized, leading to the new upper and lower polariton eigenstates.

In the 2D case, such as for quantum wells, the only significant difference is that the exciton creation operator Equation 1.57 will be expanded using the Fourier transforms of the 2D hydrogen-like wavefunctions. The result (as discussed in Section 1.3.4 and chapter 4

of Ref. [39]) is an increase in the exciton binding energy and a decrease in the Bohr radius, but with similarly symmetric solutions near  $k_{\parallel} = 0$ . Thus, the result is basically the same but with even stronger coupling, since  $g_0 \sim a_{\text{B}}^{-3/2}$ .

### 1.4.2 The polariton eigenstates

A typical exciton-polariton system consists of semiconductor quantum wells embedded within a DBR microcavity of the type discussed in Section 1.2.6. This allows the field of the cavity mode, which is in an in-plane direction, to couple with the unconfined in-plane component of the exciton. To keep things simple, and to be able to use the result of Section 1.4.1, I will consider a system of only two original states: the ground state exciton in a quantum well and the cavity photon. It is worth noting that there are actually two lowest-energy exciton states. These correspond to the coupling between heavy holes and conduction band electrons such that the total spin is either 1 or -1. These are the only states that can be produced by absorbing either a spin 1 or spin -1 (right or left circularly polarized) photon. So the exciton states we are considering can be written as  $|\frac{1}{2}, \frac{1}{2}\rangle_{\text{e}} |\frac{3}{2}, -\frac{3}{2}\rangle_{\text{h}}$  and  $|\frac{1}{2}, -\frac{1}{2}\rangle_{\text{e}} |\frac{3}{2}, \frac{3}{2}\rangle_{\text{h}}$ . These are typically called the “bright exciton” states since they can couple to a single photon. The other states, with total spin  $\pm 2$ , can also be reached, but only by a two-photon excitation process, and are referred to as the “dark exciton” states. Since I worked with photons at similar or higher energies than the exciton states, two-photon excitation is irrelevant and the dark states can be ignored since they are mostly only reachable by a spin-flip process. These typically require scattering events and are mostly negligible in the low-defect samples and low temperatures used for my work [49, 50] (see also Ref. [51], especially section 2.6).

The full Hamiltonian for this simple two-state system can be written as a sum of the two individual exciton and photon parts with the extra interaction term given in Equation 1.61, which is

$$H = \sum_{k_{\parallel}} \left[ E_{\text{exc}}(k_{\parallel}) X_{k_{\parallel}}^{\dagger} X_{k_{\parallel}} + E_{\text{cav}}(k_{\parallel}) a_{k_{\parallel}}^{\dagger} a_{k_{\parallel}} + \left( g_0 X_{k_{\parallel}}^{\dagger} a_{k_{\parallel}} + g_0^* a_{k_{\parallel}}^{\dagger} X_{k_{\parallel}} \right) \right]. \quad (1.62)$$

$E_{\text{exc}}$  is the energy of a single exciton and  $E_{\text{cav}}$  is the energy of single cavity photon. The operators  $X_{k_{\parallel}}^{\dagger}$  and  $X_{k_{\parallel}}$  are the exciton creation and annihilation operators, while  $a_{k_{\parallel}}^{\dagger}$  and  $a_{k_{\parallel}}$

are those of the cavity photons, so the first two terms are simply the Hamiltonians for the set of exciton and photon fields. The additional  $\frac{1}{2}$  can be ignored since they simply adjust the ground state energy. The third term, proportional to  $g_0$ , gives the approximate coupling between the excitons and photons, and is worked out in detail in Section 1.4.1. The first two terms simply give the energies of each type of particle multiplied by the number of each particle.

Since the excitons are trapped in a 2D plane by the quantum wells and the cavity photons are similarly trapped by the cavity, there is only a dependence on the in-plane momentum  $k_{\parallel}$ . This is another simplification of this approach, since there are in general higher order confined modes for both the exciton and the photon. As discussed in Section 1.2.6, for a microcavity the mode separation is typically quite large. In the samples used in my work, the cavities were designed to be resonant with the exciton at the third order mode ( $3\lambda/2n_{\text{cav}}$  cavity thickness). For these cavities, which we designed to work at about 1.6 eV, the next order mode is 0.53 eV higher, which is well outside of the range of the DBR stopband and completely negligible. The situation is similar for the next lower mode, which is 0.53 eV lower in energy. For the excitons, assuming an infinite square well potential and using the mass in the confinement direction (see Section 1.3.3) of  $0.33m_e$ , the next mode would be about 70 meV higher in energy. Interband transitions to these higher order states by the excitons are therefore very low in energy compared to the resonant mode. The same is even more true of transitions to hydrogen-like excited states of the exciton itself.

This Hamiltonian is easy to deal with in matrix representation since the operators all act on the same  $k$  state. If we work in the basis of states  $|C_{k_{\parallel}}\rangle$  and  $|X_{k_{\parallel}}\rangle$ , which represent the cavity photon and exciton with in-plane wavevector  $k_{\parallel}$ , for a general state  $|\Psi_{k_{\parallel}}\rangle = a|X_{k_{\parallel}}\rangle + b|C_{k_{\parallel}}\rangle$  Schrödinger's equation is simply  $H|\Psi_{k_{\parallel}}\rangle = E(k_{\parallel})|\Psi_{k_{\parallel}}\rangle$  with

$$H = \begin{pmatrix} \langle C_{k_{\parallel}} | H | C_{k_{\parallel}} \rangle & \langle C_{k_{\parallel}} | H | X_{k_{\parallel}} \rangle \\ \langle X_{k_{\parallel}} | H | C_{k_{\parallel}} \rangle & \langle X_{k_{\parallel}} | H | X_{k_{\parallel}} \rangle \end{pmatrix} = \begin{pmatrix} E_{\text{cav}}(k_{\parallel}) & g_0^* \\ g_0 & E_{\text{exc}}(k_{\parallel}) \end{pmatrix}. \quad (1.63)$$

This matrix is easily diagonalized, yielding the eigenvalues

$$E_{\text{LP}}^{\text{UP}}(k_{\parallel}) = \frac{1}{2} \left( E_{\text{cav}}(k_{\parallel}) + E_{\text{exc}}(k_{\parallel}) \pm \sqrt{[E_{\text{cav}}(k_{\parallel}) - E_{\text{exc}}(k_{\parallel})]^2 + |2g_0|^2} \right). \quad (1.64)$$

The two energy eigenvalues correspond to the new upper polariton (UP) and lower polariton (LP) states, with the UP given by the “+” sign and the LP given by the “−” sign. If we write the LP state as  $|\text{LP}\rangle = C_{k_{\parallel}} |C_{k_{\parallel}}\rangle + X_{k_{\parallel}} |X_{k_{\parallel}}\rangle$ , the LP eigenvector is given by

$$C_{k_{\parallel}} E_{\text{cav}}(k_{\parallel}) + X_{k_{\parallel}} g_0 = C_{k_{\parallel}} E_{\text{LP}}(k_{\parallel}). \quad (1.65)$$

With a bit of algebra and using the definitions  $\Delta_{k_{\parallel}} = E_{\text{cav}}(k_{\parallel}) - E_{\text{exc}}(k_{\parallel})$  and  $\Omega = |2g_0|$ , this becomes

$$C_{k_{\parallel}} \left( \Delta_{k_{\parallel}} + \sqrt{\Delta_{k_{\parallel}}^2 + \Omega^2} \right) = -\Omega X_{k_{\parallel}}. \quad (1.66)$$

Since the LP state should be normalized, we know that  $|C_{k_{\parallel}}|^2 + |X_{k_{\parallel}}|^2 = 1$ . We can also use the relationship  $\Omega^2 = \left( \sqrt{\Delta_{k_{\parallel}}^2 + \Omega^2} + \Delta_{k_{\parallel}} \right) \left( \sqrt{\Delta_{k_{\parallel}}^2 + \Omega^2} - \Delta_{k_{\parallel}} \right)$ . Equation 1.66 can then be solved as follows:

$$\begin{aligned} |C_{k_{\parallel}}|^2 \left( \sqrt{\Delta_{k_{\parallel}}^2 + \Omega^2} + \Delta_{k_{\parallel}} \right)^2 &= \Omega^2 (1 - |C_{k_{\parallel}}|^2) \\ \Rightarrow |C_{k_{\parallel}}|^2 \left[ \left( \sqrt{\Delta_{k_{\parallel}}^2 + \Omega^2} + \Delta_{k_{\parallel}} \right)^2 + \Omega^2 \right] &= \Omega^2 \\ \Rightarrow |C_{k_{\parallel}}|^2 \left[ \sqrt{\Delta_{k_{\parallel}}^2 + \Omega^2} + \Delta_{k_{\parallel}} + \sqrt{\Delta_{k_{\parallel}}^2 + \Omega^2} - \Delta_{k_{\parallel}} \right] &= \sqrt{\Delta_{k_{\parallel}}^2 + \Omega^2} - \Delta_{k_{\parallel}}, \end{aligned} \quad (1.67)$$

which finally yields the coefficients

$$\begin{aligned} |C_{k_{\parallel}}|^2 &= \frac{1}{2} \left( 1 - \frac{\Delta_{k_{\parallel}}}{\sqrt{\Delta_{k_{\parallel}}^2 + \Omega^2}} \right) \\ |X_{k_{\parallel}}|^2 &= \frac{1}{2} \left( 1 + \frac{\Delta_{k_{\parallel}}}{\sqrt{\Delta_{k_{\parallel}}^2 + \Omega^2}} \right). \end{aligned} \quad (1.68)$$

These are the *Hopfield coefficients*, which are named for J. J. Hopfield who first solved this problem in his famous 1958 paper [11]. A similar process using the UP eigenstate  $|\text{UP}\rangle = a |C_{k_{\parallel}}\rangle + b |X_{k_{\parallel}}\rangle$  returns a similar solution, but with the coefficients reversed such that  $|a|^2 = |X_{k_{\parallel}}|^2$  and  $|b|^2 = |C_{k_{\parallel}}|^2$ . These equations are somewhat ambiguous about the choice of sign. However, demanding orthogonality such that  $\langle C_{k_{\parallel}} | X_{k_{\parallel}} \rangle = 0$ , we easily arrive at the following equations for the UP and LP eigenstates:

$$\begin{aligned} |\text{LP}\rangle &= C_{k_{\parallel}} |C_{k_{\parallel}}\rangle + X_{k_{\parallel}} |X_{k_{\parallel}}\rangle \\ |\text{UP}\rangle &= X_{k_{\parallel}} |C_{k_{\parallel}}\rangle - C_{k_{\parallel}} |X_{k_{\parallel}}\rangle. \end{aligned} \quad (1.69)$$

The polariton states have a number of interesting properties that arise from their energy dispersion and the Hopfield coefficients. These are mostly dependent on the detuning between the exciton and photon states, as will be shown in the next section.

### 1.4.3 Detuning

Probably the most important parameters in discussing polaritons are the coupling strength  $|g_0| = \Omega/2$  ( $\Omega$  is called the *Rabi splitting*) and the detuning  $\Delta = E_{\text{cav}} - E_{\text{exc}}$ . There are actually two ways that the detuning is defined. The first is the general expression  $\Delta_{k_{\parallel}} = E_{\text{cav}}(k_{\parallel}) - E_{\text{exc}}(k_{\parallel})$ , already introduced in Section 1.4.2, which includes the dependence on in-plane wavenumber and is usually the most important parameter for determining specific polariton properties. Another definition is  $\delta_0 = E_{\text{cav}}(0) - E_{\text{exc}}(0)$ , which uses only the energy difference at  $k_{\parallel} = 0$ . This parameter is useful for characterizing the polariton energy dispersions in general terms, since it applies to the entire curve at all  $k_{\parallel}$ . For clarity, I will stick to these definitions throughout this thesis.

Equation 1.64 provides the polariton energies as functions of the cavity photon and exciton energies. The exciton energy given in Equation 1.29 can be rewritten, with the addition of the quantum well confinement energy  $E_{\text{QW}}$ , as

$$E_{\text{exc}}(k_{\parallel}) = E_{\text{gap}} + E_{\text{QW}} - E_{\text{B}} + \frac{\hbar^2 k_{\parallel}^2}{2(m_{\text{c}} + m_{\text{v}})} = E_{\text{exc}}(0) + \frac{\hbar^2 k_{\parallel}^2}{2(m_{\text{exc}})}, \quad (1.70)$$

where  $E_{\text{B}}$  is the exciton binding energy, and the effective exciton mass is given by  $m_{\text{exc}} = m_{\text{c}} + m_{\text{v}}$ . As discussed in Section 1.3.3, the masses of the hole bands actually depend on the direction. However, an overall approximate mass can be found as the average of the in-plane mass, which is about  $0.2m_{\text{e}}$  [52, 53]. The free electron has a lower effective mass of  $0.067m_{\text{e}}$ , which is not dependent on direction due to the *s*-like conduction band. The exciton energy can then be treated as a ground state energy with the addition of a kinetic energy of a free particle with a mass of  $0.2m_{\text{e}}$ .

The cavity photon energy is given in Equation 1.5, which can be rewritten as

$$E_{\text{cav}}(k_{\parallel}) = \frac{\hbar c}{n_{\text{cav}}} \sqrt{k_{\perp}^2 + k_{\parallel}^2} = \frac{\hbar c}{n_{\text{cav}}} \sqrt{\left( \frac{E_{\text{cav}}(0) n_{\text{cav}}}{\hbar c} \right)^2 + k_{\parallel}^2}. \quad (1.71)$$

Putting the exciton and cavity photon dispersion into Equation 1.64 gives the characteristic dispersion of the lower and upper polaritons, shown in the top row of Figure 1.12. The resonant case is shown in panel (b), where the exciton and photon energies are equal at  $k_{\parallel} = 0$ . At low  $k_{\parallel}$ , the LP is quite low in mass, while it flattens out at higher wavenumber.

This leads to an inflection point in the dispersion above which the curvature is actually negative. At negative detuning, shown in panel (a), the cavity photon energy is lower than the exciton energy at  $k_{\parallel} = 0$ , lending a more photonic-looking dispersion to the LP at low  $k_{\parallel}$ , while the UP has a flatter excitonic dispersion at low  $k_{\parallel}$ . The opposite case is apparent at the positive detuning shown in (c). This leads to another way of referring to the polariton detuning which reflects the overall character of the LP mode at low  $k_{\parallel}$ : Negative detunings are said to be “photonic” while positive detunings are said to be “excitonic.” Since the LP is the lower of the two branches, and therefore the overall ground state of the system, it is usually the focus of experimental study and it makes sense that the detunings are named with respect to it. However, in all cases, at large  $k_{\parallel}$ , the LP becomes increasingly exciton-like while the UP becomes increasingly photon-like.

Looking at Equation 1.69, in elementary terms, the probability of measuring a polariton to be in the cavity mode is simply the square of its projection into that mode, i.e.  $|\langle C_{k_{\parallel}} | LP \rangle|^2 = |C_{k_{\parallel}}|^2$ . Probably a better way to say it is that the fraction of the LP that is photon is given by  $|C_{k_{\parallel}}|^2$ . Similarly, the fraction of the LP that is exciton is given by  $|X_{k_{\parallel}}|^2$ . Therefore, we can define two parameters

$$\begin{aligned} f_{\text{cav}}(k_{\parallel}) &= |C_{k_{\parallel}}|^2 \\ f_{\text{exc}}(k_{\parallel}) &= |X_{k_{\parallel}}|^2, \end{aligned} \tag{1.72}$$

where  $f_{\text{cav}}$  is the *cavity fraction* and  $f_{\text{exc}}$  is the *exciton fraction*, again named with respect to the LP mode. These are shown in the second row of Figure 1.12, as given by Equation 1.68. This also justifies referring to the negative detunings as “photonic,” since the photon fraction (of the LP) is larger overall for  $\delta_0 < 0$  [Figure 1.12(d)], while the opposite is true for  $\delta_0 > 0$  [Figure 1.12(f)]. At  $\delta = 0$  and  $k_{\parallel} = 0$ , the exciton and photon fractions are both exactly 1/2 at  $k_{\parallel} = 0$ , corresponding to UP and LP modes that are both exactly half photon and half exciton.

One of the most striking features of a system of two coupled states appears in Figure 1.12(a): an *anticrossing* of the cavity and exciton bands. Thinking in terms of the original cavity photon and exciton, wherever they would cross each other in energy (the resonant points), they instead avoid the crossing in the presence of coupling. This anticrossing

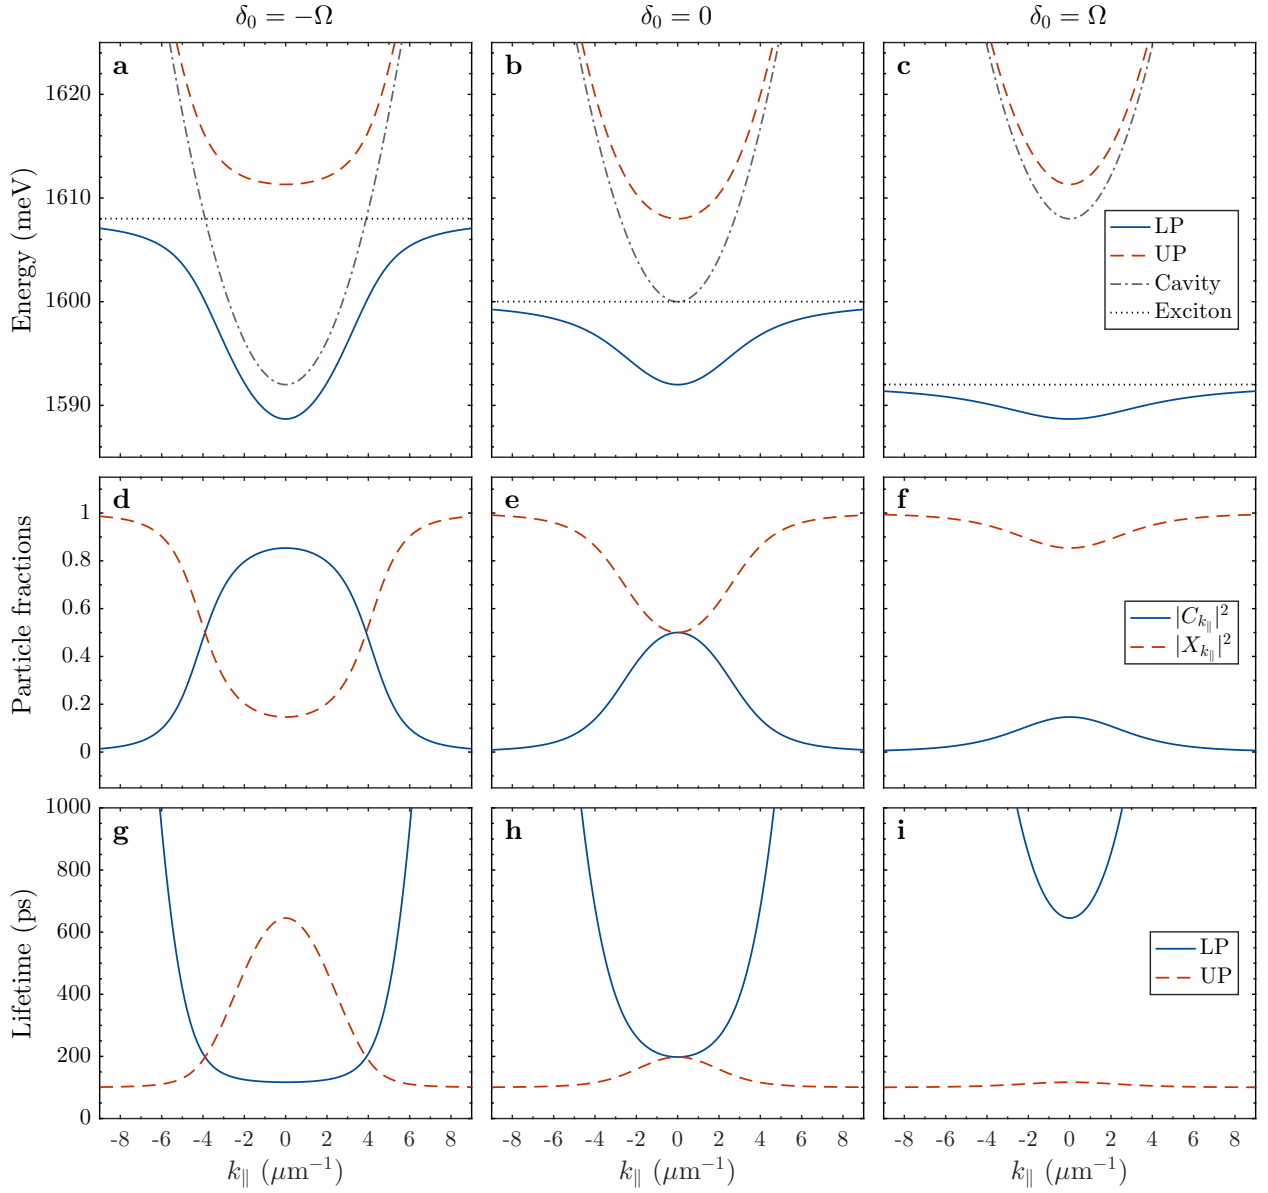


Figure 1.12: (a-c) The lower polariton, upper polariton, bare cavity, and bare exciton energies as functions of in-plane wavenumber, for an average  $k_{\parallel} = 0$  energy of 1.6 eV, a cavity refractive index of 3.4, and an exciton mass of  $0.2m_e$ . (d-e) The square of the Hopfield coefficients, corresponding the exciton and photon fractions, for the same parameters as (a-c). (g-i) The LP and UP lifetimes given by Equation 1.85, with a cavity lifetime of 100 ps and an exciton lifetime of 10 ns. The same detuning is used for all panels in the same column.

of the original bands leads to the formation of the new LP and UP modes. Since this feature is strongly dependent on the detuning, only truly appearing where the original modes are resonant [though it does show up as a “bump” in the dispersion when they are not resonant, as seen in Figure 1.12(c)], it makes sense to plot the energies of all four modes as a function of detuning, which is shown in Figure 1.13(a). In this case, I have simply used a fixed exciton energy of 1600 meV, and a coupling strength of  $|2g_0| = \Omega = 16$  meV. The cavity energy is set by simply adding the detuning  $\delta_0$  to the exciton energy. The anticrossing clearly emerges at resonance, with the original cavity and exciton modes avoiding each other to instead form the new polariton modes. The splitting is exactly  $\Omega$ , which is easy to show for the resonant condition in Equation 1.64. It should be noted that this is a plot of the energies at  $k_{\parallel} = 0$  for different  $\delta_0$ , but the exact same plot could be produced for a fixed exciton energy and detunings  $\Delta_{k_{\parallel}}$  not at  $k_{\parallel} = 0$ . This is also obvious from Equation 1.64, since for any given exciton energy and  $k_{\parallel}$ , the detuning already sets the cavity energy. A similar plot could be made by using a fixed cavity energy and varying the exciton energy with detuning, just with generally negative slopes to the curves. This is important experimentally, because it means that the same basic characteristic anticrossing plot can be produced by varying the  $k_{\parallel} = 0$  cavity energy (such as with a movable mirror [54, 55] or a cavity wedge [Section 2.1]), the  $k_{\parallel} = 0$  exciton energy (with stress [14, 23] or applied electric fields [56, 57]), or by simply going to a highly negatively detuned part of the sample and resolving wavevector and energy [such as is shown in Figure 1.12(a)].

It should be clear now that some of the most important features of the polariton modes are dependent on just the detuning and the coupling strength, as summed up in the cavity and exciton fractions. Two of the most important parameters of the polaritons, the effective mass and lifetime, can be expressed in terms of these fractions (or, identically, the Hopfield coefficients), as will be discussed in the next sections.

#### 1.4.4 Polariton effective mass

While the curvature of the polariton modes (top row of Figure 1.12) clearly changes with in-plane momentum, an expression for the effective mass at  $k_{\parallel} \approx 0$  is very useful in describing

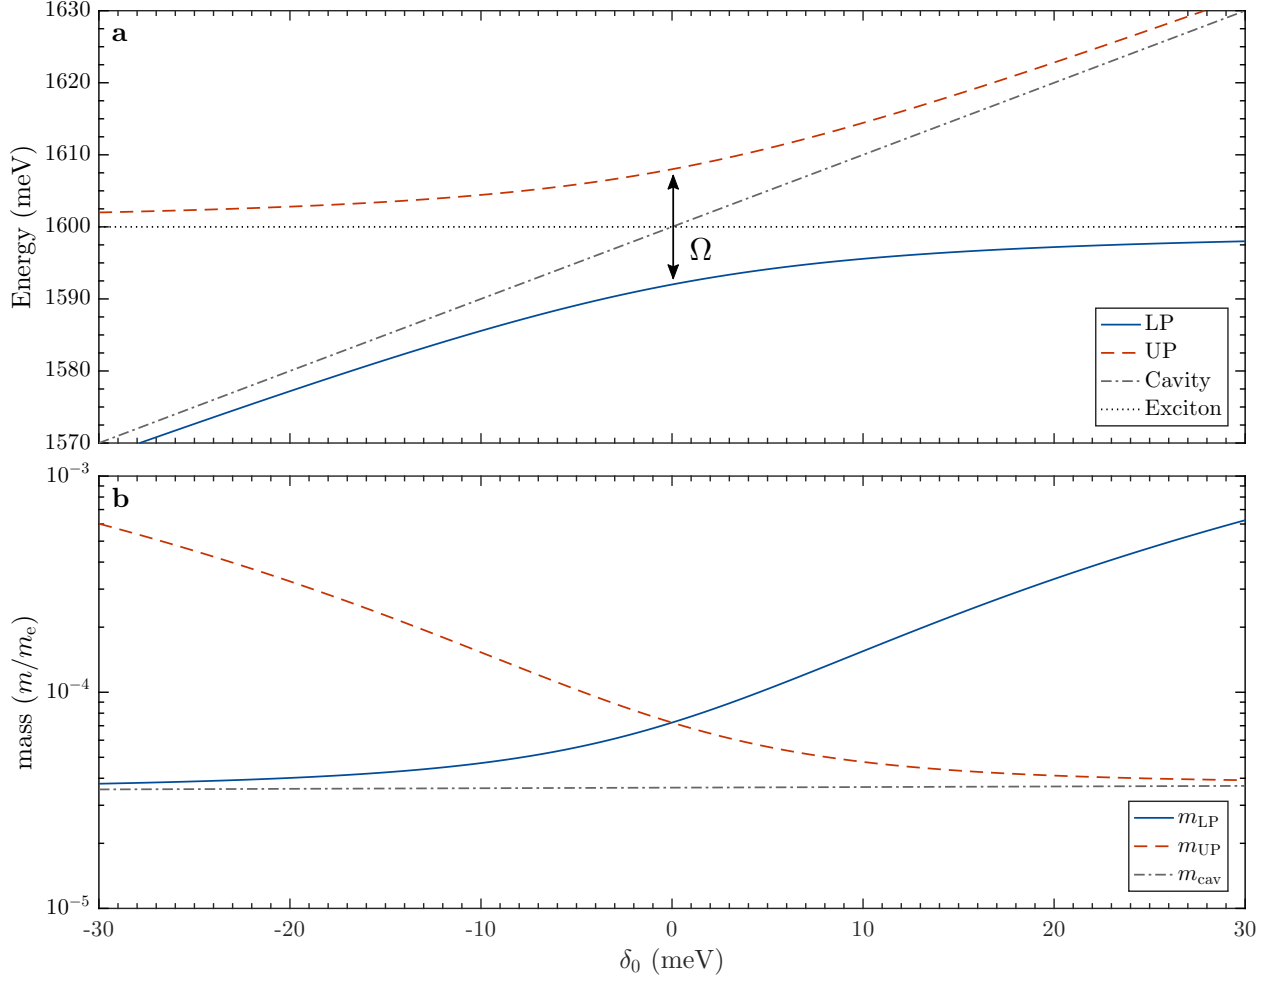


Figure 1.13: The polariton mode anticrossing. (a) The  $k_{\parallel} = 0$  energies for the LP, UP, exciton, and cavity modes, as functions of detuning  $\delta_0$ , for  $\Omega = 16$  meV. The closest the LP and UP mode get in energy is the Rabi splitting  $\Omega$ , which occurs exactly at resonance. (b) The LP, UP, and cavity masses near  $k_{\parallel} = 0$  as functions of  $\delta_0$ . The exciton mass is much higher, close to the free electron mass, so it does not appear on this plot. The LP mass quickly approaches the cavity mass at negative (photonic) detunings, while more slowly approaching the exciton mass at positive (excitonic) detunings. The UP does the same thing, but at opposite detunings.

them. Starting from the general expression of the polariton dispersions in Equation 1.64 and making use of the effective mass approximation for the cavity mode in Equation 1.6 we can write the LP energy as

$$E_{\text{LP}}(k_{\parallel}) \approx \frac{1}{2} \left( E_{\text{c},0} + \frac{\hbar^2 k_{\parallel}^2}{2m_{\text{cav}}} + E_{\text{x},0} + \frac{\hbar^2 k_{\parallel}^2}{2m_{\text{exc}}} - \sqrt{\left( E_{\text{c},0} + \frac{\hbar^2 k_{\parallel}^2}{2m_{\text{cav}}} - E_{\text{x},0} - \frac{\hbar^2 k_{\parallel}^2}{2m_{\text{exc}}} \right)^2 + \Omega^2} \right), \quad (1.73)$$

where I have used the abbreviations  $E_{\text{cav}}(0) = E_{\text{c},0}$  and  $E_{\text{exc}}(0) = E_{\text{x},0}$ . Making the substitution  $\xi = \hbar^2 k_{\parallel}^2 (1/m_{\text{cav}} - 1/m_{\text{exc}})/2$ , the terms in the square root can be rewritten as

$$\begin{aligned} \sqrt{(\delta_0 + \xi)^2 + \Omega^2} &\approx \sqrt{\delta_0^2 + \Omega^2 + 2\delta_0\xi} \\ &\approx \sqrt{\delta_0^2 + \Omega^2} \left( 1 + \frac{\delta_0\xi}{\delta_0^2 + \Omega^2} \right), \end{aligned} \quad (1.74)$$

where the first approximation was made by ignoring the terms that go as  $\xi^2 \sim k_{\parallel}^4$ , and the second was made by assuming  $\delta_0\xi \ll \delta_0^2 + \Omega^2$ , both of which are valid at small  $k_{\parallel}$ . Putting this back into Equation 1.73 gives

$$\begin{aligned} E_{\text{LP}}(k_{\parallel}) &\approx \frac{1}{2} \left( E_{\text{c},0} + E_{\text{x},0} + \frac{\hbar^2 k_{\parallel}^2}{2m_{\text{cav}}} + \frac{\hbar^2 k_{\parallel}^2}{2m_{\text{exc}}} - \sqrt{\delta_0^2 + \Omega^2} - \frac{\delta_0\xi}{\sqrt{\delta_0^2 + \Omega^2}} \right) \\ &\approx E_{\text{LP}}(0) + \frac{\hbar^2 k_{\parallel}^2}{2} \frac{1}{2} \left[ \frac{1}{m_{\text{cav}}} + \frac{1}{m_{\text{exc}}} - \frac{\delta_0}{\sqrt{\delta_0^2 + \Omega^2}} \left( \frac{1}{m_{\text{cav}}} - \frac{1}{m_{\text{exc}}} \right) \right] \\ &\approx E_{\text{LP}}(0) + \frac{\hbar^2 k_{\parallel}^2}{2} \left( \frac{|C_0|^2}{m_{\text{cav}}} + \frac{|X_0|^2}{m_{\text{exc}}} \right). \end{aligned} \quad (1.75)$$

A similar process can be done for the UP, giving the final expressions for the polariton effective masses:

$$\begin{aligned} \frac{1}{m_{\text{LP}}} &= \frac{|C_0|^2}{m_{\text{cav}}} + \frac{|X_0|^2}{m_{\text{exc}}} = \frac{f_{\text{cav}}(0)}{m_{\text{cav}}} + \frac{f_{\text{exc}}(0)}{m_{\text{exc}}} \\ \frac{1}{m_{\text{UP}}} &= \frac{|X_0|^2}{m_{\text{cav}}} + \frac{|C_0|^2}{m_{\text{exc}}} = \frac{f_{\text{exc}}(0)}{m_{\text{cav}}} + \frac{f_{\text{cav}}(0)}{m_{\text{exc}}}. \end{aligned} \quad (1.76)$$

These masses are plotted in Figure 1.13(b), for an exciton mass of  $0.2m_e$  and a cavity mass given by  $m_{\text{cav}} = (n^2/c^2)E_{\text{cav}}(0)$ , which is about  $3.6 \times 10^{-5}m_e$  for  $n_{\text{cav}} = 3.4$ , but of course varies slightly with  $E_{\text{cav}}(0)$ . The result is as expected for a composite particle: when each is larger part cavity photon, their masses are close the to cavity mass, while the opposite

is true when they are larger part exciton. At resonance, the effective masses are extremely small for a particle that is half exciton, at  $< 10^{-4}m_e$ . It is interesting to note that, while the approach to the cavity mass is rather abrupt, the effective masses are still extremely small compared to the exciton mass (about 30 times smaller, in fact) even when they have exciton fractions of about 0.94. The coupling actually results in a particle that is almost entirely excitonic but with an extremely reduced mass.

Is important to remember that a single number for a polariton mass is only valid at low  $k_{\parallel}$ , in the obviously parabolic parts of the dispersion curves. At larger  $k_{\parallel}$ , a more careful definition of the mass is necessary for a study of the dynamics of these systems (see, e.g. Ref. [58]).

#### 1.4.5 Rate equations and polariton lifetime

The typical starting point for discussing the lifetime of any particle population is first order time-dependent perturbation theory, and especially Fermi's Golden Rule (For helpful discussions of this topic, see chapter 12 of [42] and chapter 4 of [38]). This can be expressed as

$$\Gamma_{i \rightarrow f} = \frac{2\pi}{\hbar} |\langle f | H_{\text{int}} | i \rangle|^2 \delta(E_f - E_i), \quad (1.77)$$

where  $\Gamma_{i \rightarrow f}$  gives the rate of transition from an initial state  $|i\rangle$  to a final state  $|f\rangle$  and  $H_{\text{int}}$  is the interaction part of the Hamiltonian (using the “interaction representation” of Schrödinger's equation). The Dirac  $\delta$ -function simply limits this consideration to energy-conserving transitions. The states in this case are Fock states of definite particle number, so the square of the *matrix element*  $\langle f | H_{\text{int}} | i \rangle$  represents the amplitude of the transfer of particles from one state to another by the interaction. This is the complicated part, because it typically returns occupation numbers of the initial and final states when using the typical  $H_{\text{int}}$  made up of creation and annihilation operators. For number conserving interactions, the matrix element always gives a factor of  $N_i$ , since you must annihilate a particle in the initial state. For a system of bosons, creation in the final state leads to terms of  $(1 + N_f)$ , where  $N_f$  is the occupation number of the final state, resulting in stimulated scattering. For fermions, factors of  $(1 - N_f)$  appear, giving Pauli exclusion. This means that, for a

typical transition involving one particle simply moving from an initial state to a final state,  $\Gamma_{i \rightarrow f} = \gamma N_i (1 \pm N_f)$ , where  $\gamma$  is a constant that holds all the little details, and the plus sign is for Bosons and the minus sign is for fermions.

For a two-state system that has a relatively large initial population, and a very low density of particles in the final state, this results in the very simple rate equations

$$\begin{aligned}\frac{dN_i}{dt} &= -\gamma N_i \\ \frac{dN_f}{dt} &= \gamma N_i.\end{aligned}\tag{1.78}$$

The first of these equation is easy to solve, giving simple exponential decay [ $N_i(t) = N_i(0)e^{-\gamma t}$ ]. Thus, the lifetime of the state  $i$  is given by  $\tau_i^{-1} = 1/\gamma$ . If the population of the final state (or, in terms of a single particle, the probability of finding a particle in the final state) gets too high, then these equations no longer work, both because of the factors of  $(1 \pm N_f)$  and because of possible transitions *back* into the initial state.

Now if the system has more than one final state, then there will be a rate of transition from the initial state into all of the final states (which I will label  $n$ ), with varying amplitudes  $\gamma_n$ . This gives a new total rate of transitions  $\Gamma_T$  from the initial state to final states:

$$\Gamma_T = \sum_n \Gamma_{i \rightarrow n} = N_i \sum_n \gamma_n\tag{1.79}$$

This is an important result, since it states that *rates* of loss from an initial state simply add together. Thus, the new rate equation for the initial state is given by

$$\frac{dN_i}{dt} = -N_i \sum_n \gamma_n = -N_i \sum_n \frac{1}{\tau_n}.\tag{1.80}$$

The overall lifetime of the initial state is given by the sum  $\tau_T = \sum_n \tau_n^{-1}$ , where  $\tau_n$  is the individual loss time into state  $n$ . The rates for the final states are all exactly the same as in the two-state case, with each using their respective  $\gamma_n$ . For a system with a continuum of states, the same general principle applies but with integration of Equation 1.77 over energy with a density of states thrown in. The result is still basically the same, with most of the details able to be tucked away into a  $\gamma$  term. The total set of final states can be divided up at will, with different  $\gamma$ 's assigned to each subset, but the overall result is still an overall rate given by the sums. However, the ability to divide into subsets becomes especially helpful

when final state population densities get higher, as some subsets may have factors of  $(1 \pm N_f)$  while others may not.

In considering polaritons, the loss rates are due to the transition rates of the underlying cavity photon and exciton states. In general, these equations are  $k_{\parallel}$ -dependent, but I will leave off the explicit dependence since the cavity photon and exciton only couple (to good approximation) at the same  $k_{\parallel}$  anyway. Following the general process above, the rate equation for a cavity photon can then be written as

$$\frac{dN_{\text{cav}}}{dt} = -\frac{N_{\text{cav}}}{\tau_{\text{cav}}} - \frac{N_{\text{cav}}}{\tau_{\text{abs}}}. \quad (1.81)$$

The normal loss rate of the cavity is given by  $\tau_{\text{cav}}$ , which emits into the continuum of external photon states. I have also included an absorption time  $\tau_{\text{abs}}$  which corresponds to the material of the cavity or the mirrors absorbing photons. In our polariton structures the DBR and cavity spacer layers are chosen to have bandgaps well above the resonant energy of the polaritons, so absorption from those materials is negligible. Absorption from the quantum wells has already been accounted for by the coupling term that leads to polaritons in the first place (see Section 1.4.1). I have also ignored re-emission into the cavity photon state and two-photon processes (which go as  $N_{\text{cav}}^2$ ). The validity of this expression rests on a large and unfilled set of external photon states (not being in another larger cavity, for example). The exciton rate equation can be written as

$$\frac{dN_{\text{exc}}}{dt} = -\frac{N_{\text{exc}}}{\tau_{\text{nr}}}, \quad (1.82)$$

where only non-radiative ( $\tau_{\text{nr}}$ ) emission contributes a loss. Radiative emission, which just adds to the cavity photon state, has already been accounted for in the coupling in the polariton Hamiltonian. Defect and phonon scattering also do not contribute to exciton losses, since they simply change the momentum distribution while maintaining the number of particles. The nonradiative loss is mostly due to inelastic processes that drop excitons down into defect states lower in the bandgap. The exciton-photon coupling term does not contribute losses to either state since it is number-conserving, and instead leads to a sort of “steady-state” population of both the exciton and photon states, which are already expressed in the exciton and photon fractions at any given  $k_{\parallel}$ . Therefore, in a polariton system, only

the terms giving radiative loss of the cavity ( $\tau_{\text{cav}}$ ) and nonradiative loss of the excitons ( $\tau_{\text{nr}}$ ) actually contribute to loss of polaritons. Since the polaritons are made up of both excitons and cavity photons, they will suffer from losses of either state. For the lower polariton, this can be expressed as

$$\frac{dN_{\text{LP}}}{dt} = -\frac{N'_{\text{cav}}}{\tau_{\text{cav}}} - \frac{N'_{\text{exc}}}{\tau_{\text{nr}}} = -\frac{N_{\text{LP}}}{\tau_{\text{LP}}}, \quad (1.83)$$

where I have used primes to specify that these are the exciton and photon states that are actually coupled (and therefore “participating” in the lower polariton state). A similar equation can be written for the UP state, but with terms that allow for conversion into lower polaritons and dark exciton states. Since the cavity photon and exciton populations are given simply by the projection of the polariton mode onto each of those states times the population of polaritons, this can be nicely rewritten as

$$\frac{dN_{\text{LP}}}{dt} = -\frac{N_{\text{LP}}|C|^2}{\tau_{\text{cav}}} - \frac{N_{\text{LP}}|X|^2}{\tau_{\text{nr}}} = -\frac{N_{\text{LP}}}{\tau_{\text{LP}}}, \quad (1.84)$$

which immediately leads to the final result for the UP and LP lifetimes (with the  $k_{\parallel}$  dependence put back in explicitly):

$$\begin{aligned} \frac{1}{\tau_{\text{LP}}(k_{\parallel})} &= \frac{|C_{k_{\parallel}}|^2}{\tau_{\text{cav}}} + \frac{|X_{k_{\parallel}}|^2}{\tau_{\text{nr}}} \\ \frac{1}{\tau_{\text{UP}}(k_{\parallel})} &= \frac{|X_{k_{\parallel}}|^2}{\tau_{\text{cav}}} + \frac{|C_{k_{\parallel}}|^2}{\tau_{\text{nr}}}, \end{aligned} \quad (1.85)$$

where I have ignored the additional decay terms for the UP mentioned above. In general, the cavity lifetime is usually the dominant factor since it is much smaller (usually  $1 - 100$  ps) at lower temperature than the nonradiative exciton lifetime (typically  $\sim 1$  ns). These lifetimes are plotted on the bottom row Figure 1.12 as a function of in-plane momentum, clearly showing how they change as exciton and cavity fractions change (shown in the middle row of the same figure). A value of 100 ps was used for the cavity lifetime, which is close to the measured values in our samples ([25, 59]). A lifetime of 10 ns was used for the excitons, which should also be close to the correct values for the high quality, low defect quantum wells in our samples.

### 1.4.6 Strong coupling

If photons are trapped in the cavity forever and the excitons never lose their coherence from nonradiative processes, the picture of two new polariton eigenstates applies to systems with even the weakest coupling strengths. However, real systems have decay, which causes line-broadening in both the exciton and photon states. Practically speaking, this means that in order to actually observe the splitting associated with coupling, as shown in Figure 1.13, the total splitting must be large enough to distinguish two separate lines at the anticrossing. Otherwise the two lines simply overlap and look like a normal uncoupled crossing of modes.

As discussed in Section 1.2.3, the cavity energy linewidth is simply inversely proportional to the cavity lifetime. A similar treatment can be done for the exciton. These can be included in Equation 1.64 by adding an *imaginary self-energy* to each of the original mode energies [60,61]. The energies are then changed such that  $E_{\text{cav}} \rightarrow E_{\text{cav}} + i\Gamma_{\text{cav}}$  and  $E_{\text{exc}} \rightarrow E_{\text{exc}} + i\Gamma_{\text{exc}}$ , where  $\Gamma_{\text{cav}} = \hbar/\tau_{\text{cav}}$  and  $\Gamma_{\text{exc}} = \hbar/\tau_{\text{nr}}$ , as given in Section 1.4.5. As mentioned in that section, only the nonradiative exciton lifetime should be included since the radiative coupling to the photon field has been explicitly included in the Hamiltonian. At  $k_{\parallel} = 0$ , this gives

$$E_{\text{LP}}^{\text{UP}} = \frac{1}{2} \left( E_{\text{cav}} + E_{\text{exc}} + i(\Gamma_{\text{cav}} + \Gamma_{\text{exc}}) \pm \sqrt{[E_{\text{cav}} - E_{\text{exc}} + i(\Gamma_{\text{cav}} - \Gamma_{\text{exc}})]^2 + |2g_0|^2} \right) \quad (1.86)$$

At resonance, when the exciton and cavity energies are equal, two effects become immediately obvious. The first is that the two modes now have equal broadenings, given as the average of the broadening of the original modes. The second is that the splitting is effectively reduced, no longer being given by twice the coupling strength. In fact, it disappears unless  $2g_0 > \Gamma_{\text{cav}} - \Gamma_{\text{exc}}$ , becoming another imaginary term that affects the line broadening (though also making the linewidths of each mode different from the other). While this condition technically leads to splitting, it still is not necessarily strong enough for the practical observation mentioned above. The condition for strong coupling should then be  $\Omega' = \sqrt{|2g_0|^2 - (\Gamma_{\text{cav}} - \Gamma_{\text{exc}})^2} > (\Gamma_{\text{cav}} + \Gamma_{\text{exc}})/2$ , where  $\Omega'$  now gives the actual splitting at resonance.

It is interesting to note that the splitting is affected by the difference in the decay rates of each original mode, being greatest when the decay rates are equal. Thus, for strong coupling, we need to have both a small difference in decay rates and overall small decay

rates. However, if the latter condition is sufficiently met, then the former will be met as well. The decay rate of one of the modes could be zero and the splitting still very large as long as the decay rate of the other mode is very small compared to the coupling strength. For example, considering the samples used in the work presented in this thesis (which are described in Section 2.1), the cavity lifetime is  $\sim 100$  ps, giving  $\Gamma_{\text{cav}} \sim 10 \mu\text{eV}$  (There is, of course, also inhomogeneous broadening apart from the exponential decay rates, but this is also small in our samples. See e.g. chapter 2 of Ref. [39]). The nonradiative exciton lifetime is  $\sim 1 - 10$  ns, so  $\Gamma_{\text{exc}}$  is negligible in comparison. Since  $\Omega$  has been measured at about 16 meV, it is clear that, despite the *lifetimes* of the exciton and photon being vastly different, their decays both contribute a negligible change to the splitting, and only the cavity decay rate significantly affects the linewidths at resonance.

#### 1.4.7 Interactions

Since photons are only extremely weakly interacting, the polaritons mostly interact through their exciton component. Therefore, for any interaction strength  $g_{\text{x}-i}$  between an exciton and a general particle  $i$ , the lower polariton interaction strength with the same particle will be given by  $g_{\text{p}-i} = |X_{k_{\parallel}}|^2 g_{\text{x}-i}$ , where  $|X_{k_{\parallel}}|^2$  is the exciton fraction of the polariton. There are three important types to consider for my work: exciton-exciton, exciton-phonon, and exciton-electron interactions.

The exciton-exciton interaction is overall repulsive. This is because the direct Coulomb dipole-dipole interaction is actually very weak compared to the exchange between the electrons and holes that make up the exciton [32]. It is this interaction that leads to blue shifting of the polariton mode in the presence of both other polaritons and uncoupled or high- $k_{\parallel}$  excitons. This is a very useful effect, since it allows manipulation of the ground state energy of the polaritons, as discussed in Section 2.2.1. It also leads to a “smoothing out” of the ground state, as polaritons fill in shallow local minima, eventually producing extensive level regions in energy (as shown in Section 3.3). The exciton-phonon interaction is also important, but only indirectly as it is an avenue by which excitons are able to thermalize. It is, however, very inefficient at the “bottleneck” region of the lower polariton dispersion,

which is the region of large slope near and above the inflection point in Figure 1.12(a-c), and is discussed further in Chapter 6. Since the density of states is greatly reduced, scattering into states near and below the inflection point is highly inefficient, causing a build-up of excitons (technically high- $k_{\parallel}$  lower polaritons) just above the bottleneck region. As discussed in Ref. [33], exciton-phonon interactions seem insufficient for thermalizing below that point. However, exciton-electron (and exciton-hole) interactions can make up for this, because the free electrons are much more strongly coupled to the phonon bath.

The exciton-electron interaction is especially important to the work in this thesis, since much of it was focused on looking for effects of polaritons on free charge carriers and vice versa. This has been worked out previously [33, 62], though Ref. [33] provides a simple picture of collisional scattering between the free electrons and the electron components of the excitons. Another view of this interaction is through the polaron effect, in which a free electron interacts with a background of dipoles, or the other way around [63, 64]. In this picture, excitons (and therefore also polaritons) and electrons induce a drag on each other. Such drag effects have also been considered in work on excitons and free electrons in separate layers [65–67], though the separation is not necessary in principle.

#### 1.4.8 Strain

The effect of strain on the polaritons is similar to interactions in that it mostly only affects the exciton part, leaving the cavity photon component much the same as without strain. It does so by shifting the underlying electron bands. For the valence bands of GaAs and similar materials, the standard approach for dealing with strain is with the Pikus-Bir Hamiltonian [68] (see chapter 4 of Ref. [40] and chapter 3 of Ref. [43] for general discussions of this topic). This can be written as

$$H_{\text{PB}} = a(\epsilon_{xx} + \epsilon_{yy} + \epsilon_{zz}) + b[(J_x^2 - J^2/3)\epsilon_{xx} + c.p.] + \frac{2d}{\sqrt{3}}\left(\frac{1}{2}(J_x J_y + J_y J_x)\epsilon_{xy} + c.p.\right), \quad (1.87)$$

where the  $\epsilon$  terms are the components of the strain tensor corresponding to deformation of the original lattice,  $a$ ,  $b$ , and  $d$  are deformation potentials, and the  $J$ 's are the angular momentum operators that act on the spin states of the bands. The “ $c.p.$ ” stands for cyclic

permutations of  $x$ ,  $y$ , and  $z$ . This results in a Hamiltonian identical in form to the Luttinger-Kohn Hamiltonian discussed in Section 1.3.3. An applied stress will shift the energies of the electron bands, effectively changing the band-gaps between various states. This was used to great effect in earlier work, where stress was applied to create a low energy point in the valence-conduction band gap (see Refs. [14, 23, 69, 70], or Refs. [71, 72] for more details). In general, a fairly complex analysis of the stress applied to a given structure must be carried out in order to determine the components of the strain tensor. For my purposes, it is enough to consider that a strain on the material will cause a shift of the electron bands, subsequently shifting the exciton energy given by Equation 1.29 by changes to the energy gap and also possibly the individual masses.

The result for a given location on a sample is a change in detuning while the cavity photon energy remains the same. This can obviously change both the polariton energies as well as the Hopfield coefficients.

#### 1.4.9 Bose-Einstein condensation

Bose-Einstein condensation (BEC) is definitely one of the more fascinating phenomena that appears in bosonic systems. As discussed in Section 1.1, it was first worked out by Einstein following the statistical approach of Bose in the 1920s, and first seen conclusively in cold atoms in 1995. The fundamental requirement is that the particles involved be bosons. Excitons are approximate bosons, fulfilling the standard bosonic commutation relation  $[X_{q'n'}, X_{q,n}^\dagger] = \delta_{q,q'}\delta_{n,n'}$  very well as long as  $n_{\text{exc}}a_{\text{B}}^3 \ll 1$ , where  $n_{\text{exc}}$  is the exciton density and  $a_{\text{B}}$  is the exciton Bohr radius. This basically means that excitons can be considered bosons as long as they do not overlap in space, at which point the fermionic effects of the underlying electrons and holes become important [12].

BEC follows pretty easily from the Bose-Einstein distribution,

$$N(k) = \frac{1}{e^{\beta(E_k - \mu)} - 1}, \quad (1.88)$$

where  $\beta = k_{\text{B}}T$ ,  $k_{\text{B}}$  is the Boltzmann constant,  $T$  is the temperature, and  $\mu$  is the chemical potential. Following Ref. [73], the total number of particles can be found by summing this

over all  $\mathbf{k}$ , which can be split into two parts:

$$\begin{aligned} N &= \sum_{\mathbf{k}} \frac{1}{e^{\beta(E_{\mathbf{k}}-\mu)} - 1} = \frac{1}{e^{\beta(E_0-\mu)} - 1} + \sum_{\mathbf{k} \neq 0} \frac{1}{e^{\beta(E_{\mathbf{k}}-\mu)} - 1} \\ &= N_0 + N_T. \end{aligned} \quad (1.89)$$

The ground state occupation is given by  $N_0$ , while the total population of everything outside of the ground state is given by  $N_T$ .  $N_0$  is negligible compared to  $N_T$  except when  $\mu \rightarrow E_0$ , where the ground state occupation diverges. This is also the point at which  $N_T$  is maximized for a given temperature. Therefore, when  $\mu = E_0$ , the total number of particles in excited states reaches a maximum critical value  $N_C$ . This means that, at a set temperature, as particles are added to the system, the occupancy of the excited states is limited such that the population in the ground state must increase in order to accommodate additional particles and satisfy Equation 1.89. Since  $N_T$  decreases with temperature,  $N_C$  also decreases with temperature. This means that for any total number of particles  $N$ , there will be a critical temperature  $T_C$  at which particles must begin to pile up in the ground state. The point at which this happens is given by the condition

$$N_T(T_C, \mu = E_0) = N. \quad (1.90)$$

Below the critical temperature for a given population, or above the critical density for a given temperature, the ground state becomes highly occupied, which is called Bose-Einstein condensation.

For a 3D system of an ideal Bose gas, with no interactions and confined in a box of volume  $V$ , the critical temperature is given by [73]

$$T_C = \frac{2\pi\hbar^2}{mk_B} \left( \frac{n}{2.612} \right)^{2/3}, \quad (1.91)$$

where  $n$  is the particle density and  $m$  is the particle mass. Other than increasing the density, the only way to significantly increase the critical temperature is by decreasing the particle mass, which is one of the major advantages of the extremely light mass of polaritons. In two dimensions, the integral for  $N_T$  actually diverges, allowing unlimited occupation of excited states [74]. However, this can be overcome by including a trapping potential, which modifies

the density of states and causes the integral to converge (see chapter 17 of Ref. [73], and also Refs. [75, 76]). Finally, Equation 1.91 can be rearranged to give the critical density as

$$n_C = 2.612 \left( \frac{2\pi\hbar^2}{mk_B T} \right)^{-\frac{3}{2}} = 2.612 \cdot \lambda_T^{-3}, \quad (1.92)$$

where I have substituted in the thermal de Broglie wavelength  $\lambda_T$ . This means that the BEC occurs when the density reaches a point where the particle spacing becomes comparable to  $\lambda_T$ . This is equivalent to saying that condensation occurs when the wavefunctions of individual particles begin to significantly overlap. Using this argument, we can deduce the basic behaviour in the 2D case without knowing anything about the trapping potential. Since in 2D the spatial extent of a single particle is given by its area which goes as  $\lambda_T^{-2}$ ,

$$n_{C,2D} \sim \lambda_T^{-2} = \left( \frac{2\pi\hbar^2}{mk_B T} \right)^{-1}. \quad (1.93)$$

For polaritons with mass  $\sim 10^{-4}m_e$  and a temperature of 10 K, this gives a critical density of about  $0.2 \mu\text{m}^{-2}$ , which is easily achievable in our systems.

## 2.0 FABRICATION AND EXPERIMENTAL METHODS

### 2.1 MICROCAVITY SAMPLE DESIGN

The samples I used in my work were all of the same design basic design, which was worked out by my predecessors. Much of that work, including the specific use of the transfer matrix method in that process, is detailed in Refs. [72, 77]. In addition to the work in this thesis, these samples were also used in many other projects [23, 25, 59, 78, 79]. I will give the relevant information here as well.

The overall design is shown in Figure 2.1. The structure was grown on a GaAs substrate using molecular beam epitaxy (MBE) by our collaborators Loren Pfeiffer and Ken West. A  $3\lambda/2n$  cavity was formed by two  $\text{Al}_{0.2}\text{Ga}_{0.8}\text{As}/\text{AlAs}$  distributed Bragg reflectors (DBRs) with 32 periods in the top DBR and 40 periods in the bottom DBR. Within the cavity are a total of 12 GaAs quantum wells with AlAs barriers. The wells themselves are nominally 7 nm thick. The cavity was designed to work in the 3<sup>rd</sup> order mode (that was the mode resonant with the excitons), so the quantum wells are arranged in three sets with one set at each antinode of the cavity mode. Since these are the locations with the largest amplitudes of the field  $\mathbf{A}$ , and the coupling increases with the amplitude of  $\mathbf{A}$  (see Section 1.4.1), these locations give the best exciton-photon coupling. To further increase the coupling, each set has four coupled quantum wells (coupled via the basic tunneling described in Section 1.3.4), with separation barriers nominally 3 nm thick. Due to the high barriers ( $\sim 100$  meV), the coupling is weak and the splitting of the coupled states is negligible compared to the exciton linewidth (and therefore not detectable with spectroscopic measurements). This means that the multiple coupled states are essentially degenerate in our samples, with excitons rarely tunneling between wells. The purpose of the four closely-spaced quantum wells is to increase

the probability of absorption of a photon as it makes a pass through the cavity, similarly to having quantum wells at each antinode, and thereby increase the coupling. The large number of carefully placed quantum wells leads to a measured Rabi splitting of about 15.9 meV (see Section E.1).

The large number of DBR periods in the mirrors gives the cavity a very high  $Q$ -factor, resulting in a cavity lifetime of  $\sim 100$  ps [25] (for more discussion, see Sections 1.2.3 and 1.2.6). The long cavity lifetime, combined with the very low defect density in the quantum wells, leads to very long transport distances with few scattering events [25].

The cavity itself is actually made up of all three materials. Obviously, significant parts of it are pure GaAs and AlAs because of the quantum wells, but the spacers are made of layers of both AlAs and the alloy  $\text{Al}_{0.2}\text{Ga}_{0.8}\text{As}$ /AlAs. This causes a slight angle dependence on the effective cavity refractive index, though this is negligible as discussed in Appendix A. The average cavity refractive index has been measured as  $n_{\text{cav}} \simeq 3.4$  using the characterization method outlined in Section E.1.

### 2.1.1 The cavity wedge

Because of the MBE growth process, our samples have a significant change in thickness radially from the center of the wafer to the edges. This causes a change in the thickness of each layer in the structure, with the thickest part at the wafer center and the thinnest at the edge. Since both the cavity and DBR stop band energies are dependent on the layer thicknesses, these change across the sample. Thankfully, the dependence scales equally for both parts, so the cavity mode remains in generally the same location of the stopband as they both shift in energy. The exciton also shifts slightly, but to a lesser extent, due to the quantum well confinement energy dependence on thickness ( $\sim 1/L^2$ ). While the cavity mode is only linear with thickness, the exciton energy changes much less across the sample because it is only partially composed of the confinement energy, while its largest portion comes from the constant bandgap energy. This effect is shown in Figure 2.2 for a similar sample but with about half the number of DBR layers [69].

The most significant result of this wedge is a gradient in the polariton energy. This leads

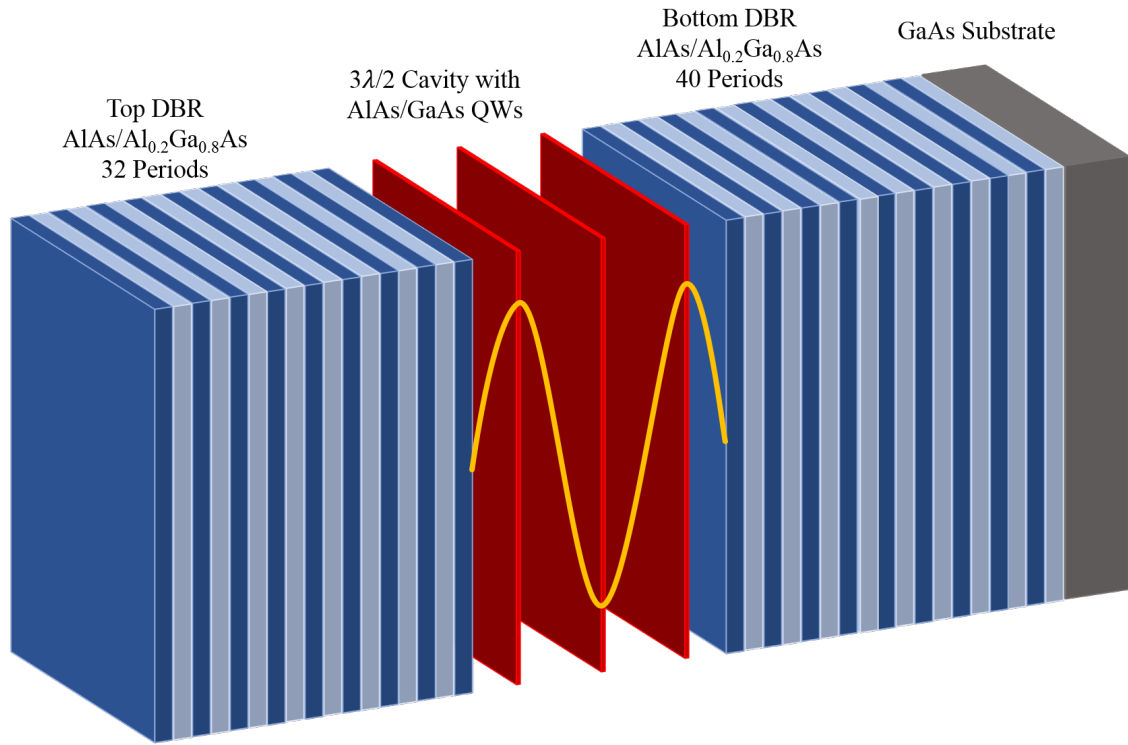


Figure 2.1: Diagram of the microcavity sample design. Three sets of four quantum wells were situated with one set at each of the antinodes of the  $3\lambda/2n$  DBR cavity.

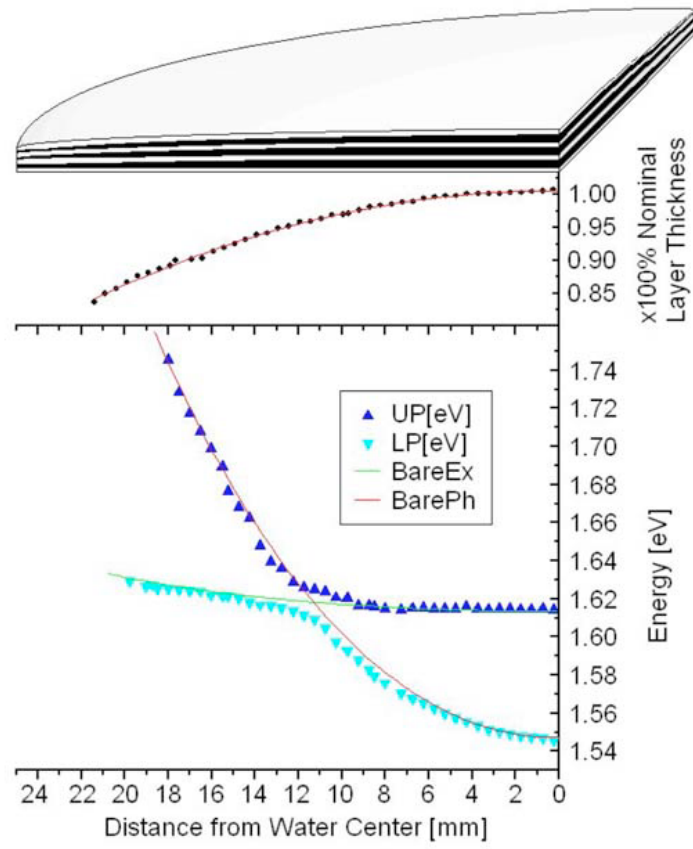


Figure 2.2: The change in mode energies across the sample due to changing layer thicknesses. Reprinted from Ref. [69], with the permission of AIP Publishing.

to an overall force on the polaritons from the standard  $\mathbf{F} = -\nabla U$ . In any small region of sample ( $\sim 100 \mu\text{m}$ ), the gradient is approximately constant, leading to a gravity-like force in the direction of the center of the sample wafer.

## 2.2 FABRICATION METHODS

The main idea for our devices was to produce polaritons in a trap, and push current through that trap region in order measure some interaction between the charge carriers and the polaritons. While a trap is not necessary to produce polaritons, we were primarily interested in what would happen when the polaritons formed a condensate, due to the stimulated scattering effects that arise when bosons are in a macroscopically occupied ( $N \gg 1$ ) state (see Section 1.4.5). The trap then serves the purpose of allowing a condensate to form in a well defined and spatially confined ground state (see Section 1.4.9 and also [80]).

### 2.2.1 Review of confinement and fabrication methods

There have been a number of methods demonstrated over the years of producing traps or otherwise manipulating the potential energy profile for polaritons. Many of these are discussed at length in Ref. [57], but I will briefly mention several of them here. Most of the work discussed in this review section was not done by our group, but it helps to set the context in which we chose our particular fabrication methods. The first demonstration of polaritons in an intentionally produced trap (not a trap due to a sample defect) used stress to shift the exciton energy [14]. By first thinning the sample to  $\approx 100 \mu\text{m}$  using abrasion and bromine to remove most of the substrate, and then pressing a sharp needle (just a regular sewing pin is fine) into the back side, the sample can be stressed enough to shift the exciton energy by several meV. Due to the radial symmetry, a harmonic bowl-shaped trap is produced in the center [14, 70]. Polaritons produced nearby flow into the trap and form a condensate. As discussed in Section 1.4.7, polaritons interact via their excitonic part (somewhat obvious since photons are essentially non-interacting), and are therefore capable

of interactions with each other and with uncoupled or weakly coupled excitons. The use of a nonresonant pump (see Section 2.3.1.2) gives a population of hot (high  $k_{\parallel}$ ) excitons at or near the location of the pump spot. Since polaritons are repelled by excitons, a simple adaptation of the harmonic bowl trap is the addition of a nonresonant pump in the middle. This puts a barrier in the center of the trap, producing a ring, and also pumps the polariton state, forming a ring condensate [23].

The barrier produced by the nonresonant pump spot can also be used by itself without stress. One pump spot by itself is actually enough to produce a condensate at high enough densities due a local minimum created between the pump spot and the natural slope from the cavity wedge (see Section 2.1.1). This produced a trapped polariton condensate with very long coherence times [78]). Two pump spots together can be used in a similar way, but with the trapping now between the two spots [81]. The more obvious geometries involve more than two spots arranged in triangles, squares, and higher order polygons with the trap in the middle [82], or a ring-shaped pump [26, 83–85] (Ref. [26] is particularly interesting because they actually measured an equilibrium Bose-Einstein distribution).

Both applied stress and pump-induced trapping work by shifting the exciton energy. Polaritons can also be confined by shifting the cavity photon energy. Technically, the case mentioned above of a trap formed in the “nook” between a pump spot and the natural slope of the sample already does this, since the slope is mostly caused by the changing cavity energy. That gradient is rather shallow ( $\sim 10$  meV/mm), and the concept of changing the cavity energy by changing the spacing of the mirrors can be carefully engineered during the original sample fabrication. This is typically done by pulling the sample out of the growth chamber after depositing the bottom mirror, quantum wells and top cavity spacer, etching the cavity spacer slightly in a pattern, and finally putting it back into the growth chamber and depositing the top mirror [86–88]. The etched areas will have a smaller cavity thickness than the unetched areas, and therefore also a higher energy. Since the mode antinodes are also shifted by this process, while the quantum wells stay in the same place relative to the bottom mirror, the coupling should also be affected (reduced if the unetched design was idealized), further increasing the trapping potential since weaker coupling increases the lower polariton energy. A different overall fabrication method could be used, of course, as

long as some areas have a thicker cavity than others. Pretty much any trapping geometry in two dimensions can be realized, with demonstrated potential well depths of several meV. Another related method is depositing metal strips onto the top mirror [89]. This changes the boundary conditions at the top surface, raising the cavity photon energy slightly by shifting the mode. This only gives trap depths  $\sim 100 \mu\text{eV}$ , so its usefulness is limited.

Another method is to simply not put a top mirror everywhere on the sample. While selectively growing a DBR structure in a pattern is not particularly feasible (they are usually grown at high temperature and are very thick), top reflectors have been formed out of subwavelength gratings [90,91]. This type of grating is designed such that the only allowed diffraction mode is the 0<sup>th</sup> order (hence the name “subwavelength”), giving high reflectivity due simply to periodicity of the material. In the works cited, this was done by placing both a sacrificial layer and a thin capping layer on top of the cavity. The capping layer was etched into a subwavelength grating, and the sacrificial layer was removed via acid etching, resulting in an air gap. Since the grating can be produced in different sizes and shapes, small trap regions can be formed. The immediate downsides are that the reflectivity of such a structure is not nearly as high as a DBR (which can be increased by continuing to add layers), and is also highly polarization dependent (though this could be an advantage if polarization selection is desired).

Finally, the polaritons can be confined by etching away the DBR microcavity structure to form wires [92–94], pillars [95–98], coupled pillar arrays [99], and 2D lattice arrays [100]. This is the method that I used for my work. When the entire top DBR is removed, this has the advantage of completely eliminating the cavity photon state. It can, however, lead to edge recombination due to the etched and exposed quantum wells. As discussed in Chapter 3, this is not always an issue due to repulsion from the edges.

**2.2.1.1 Advantages of etched pillars** Each method has its own advantages and disadvantages. Probably the greatest advantage of using the repulsion from the pump spot as a trapping method is that it is temporary (unless the sample is damaged by extremely high pump intensities, of course), but it does link the trapping to the pumping. Even stress traps are not completely removable, as residual strain is detectable after the needle is re-

moved, but they do allow the pump to be independent of the trapping, as do all of the other methods. Etched traps are obviously more permanent, so they cannot be easily changed or removed after fabrication, but they produce very deep traps (essentially infinitely high boundaries because the cavity state does not exist in the etched regions). This is their primary advantage over the patterned deposition of metal, which produces very shallow traps. Patterned etching of the cavity spacers can produce relatively deep traps, but removing a sample mid-growth tends to reduce the final quality, especially the quality of the top mirror in this case. Subwavelength gratings have similar problems by starting off with lower-quality top reflectors by design. Our goal was to produce traps in samples of the quality discussed in Section 2.1, so these methods were not an option.

Additionally, our goal was to push current through the traps. As discussed in Section 4.1, many studies have already been done on vertical (perpendicular to the structure) injection of current in polariton structures. However, since the most obvious measurements that can be made of the polaritons are their energy and in-plane momentum, it makes sense to look for interactions that directly affect either of these. Some of most obvious types of interactions are of the collisional scattering variety, which involve some exchange of momentum. Thus, we wanted to be able to give charge carriers in-plane momentum, and have some control over that momentum, so in-plane injection made the most sense. Since the excitons exist in the quantum wells, the most direct method of injecting charge carriers was to make contact directly onto (or close to) the quantum well layers, which for our structure required etching away the top DBR. Therefore, the compatibility of etching with our current sample growth methods, the inherent advantages of etched traps (deep confinement, independence of the pump from the trap), and the need to etch away the top DBR and expose the quantum wells for contacts all combined to make etched trapping the best candidate for the project.

### 2.2.2 Etching

The overall idea of the design of our devices was to have a pillar region confining the polaritons, with contacts on the quantum wells on both sides of the pillar. This would allow us to produce a polariton condensate within the pillar and then push current through quantum

wells within the same pillar. In order to be sure that the current in the quantum wells flows through the pillar region, we added the additional step of etching the quantum wells such that the current would be confined to flow through the pillar region. For more details on all of the general processes discussed in this section, see Ref. [101]. Figure 2.3 shows an overview of the entire fabrication process that we use.

The main difficulty in etching our devices is the thickness of the top DBR, which is  $\approx 3.2 \mu\text{m}$  thick. In order to fully remove the cavity photon mode and access the quantum wells, we needed to etch away the entire top DBR. For an isotropic etch (such as a wet etch) the minimum attainable feature size is on the order of the etch depth, and such an etch produces roughly quarter-circle sidewall profiles. We therefore required an anisotropic dry etch, which gives a mostly vertical sidewall profile, encouraging total internal reflection at the edges of the pillars. Furthermore, anisotropic etches allow for feature sizes much smaller than the etch depth, and generally more consistent etch profiles for deep etches.

The main disadvantage of dry etches is that they have much lower selectivity between the masking material and the sample materials than purely chemical wet or plasma etches. For example, an acid can quickly etch away metal or oxide, but usually has almost no effect on photo-resist (or e-beam resist) since it is a polymer. Thus, a very thin layer of resist can be used to mask for a wet etch. This is important because, similar to etching, the minimum feature size of lithography is limited by the resist thickness (it is ultimately limited by the wavelength of the particles used to expose the resist). A dry etch process, such as reactive ion etching (RIE), achieves anisotropy by combining the chemical reactivity of the etchants with the sputtering effect of bombarding the material with high-energy ions. The sputtering, while highly directional, is not as material-dependent, and is somewhat effective on all materials. This means that much thicker or more complicated etch masks must be used, since the mask itself is more effectively etched by the process. Since thicker masks reduce the resolution of the lithography, while anisotropy increases it, a balance must be found between etch selectivity (between the mask and the sample material) and the mask thickness.

Earlier work on etched polariton structures [92, 102] used relatively complicated electron beam lithography (EBL) and a RIE based on HBr and O<sub>2</sub>. Many of the details of the etch

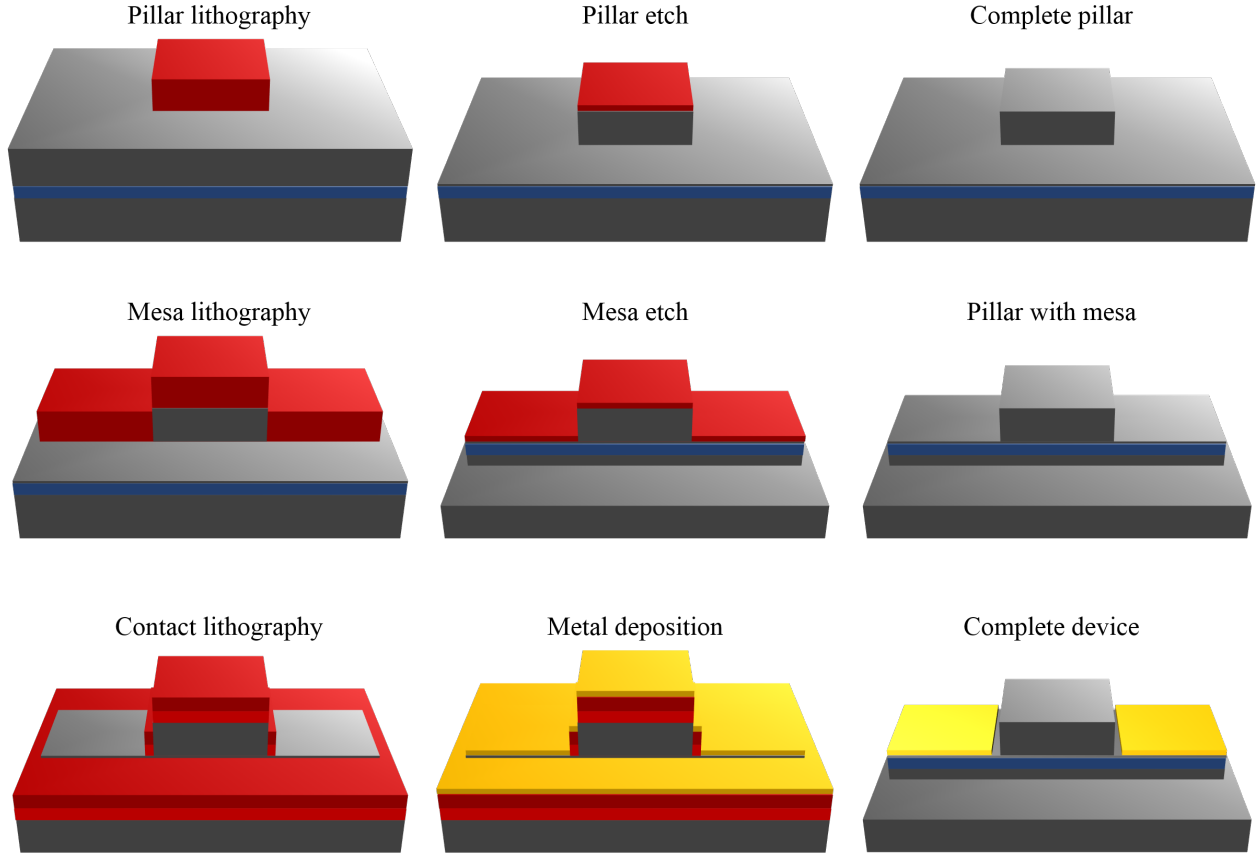


Figure 2.3: Overview of the fabrication process, for devices with sides  $\sim 10 - 100 \mu\text{m}$  long. Beginning with a cavity containing quantum wells (blue) and top and bottom DBRs (gray), an etch mask is formed out of photoresist (red). The sample is etched, removing most of the top DBR wherever there is no mask, resulting in a microcavity pillar. A new mask is formed by the same general technique, protecting the pillar and large regions on both sides from an etch that removes the cavity (and quantum well) layers, resulting in a pillar structure within a larger shallow mesa. A third mask is formed, this time with two layers: a thinner photoresist layer on top of a lift-off resist layer. The mesa regions beside the pillar are left unmasked. A four layer Ni/AuGe/Ni/Au metal stack is evaporated onto the sample. After lift-off, the metal stack is left just above the quantum wells, forming two contacts, one on each side of the pillar.

process used in those works are in Ref. [97]. EBL, which can achieve extremely small feature sizes, was necessary in their case because of the small size of the devices. To get the high selectivity needed for the deep etch (they etched through the top DBR, cavity, and most of the bottom DBR), the EBL step was used to produce a hard mask out of  $\text{Si}_3\text{N}_4$ , which had been deposited on top of the sample. This hard mask then protected the sample during the deep HBr etch.

In our case, we were mostly interested in fabricating pillars on a larger scale (lateral sizes of  $\sim 10\text{ }\mu\text{m}$ ), so I settled on using a more basic photolithography process. The photoresist itself is used to mask the actual etch of the sample. We also did not have easy access to a HBr etching system, so we used a chlorine-based inductively coupled plasma (ICP) RIE at a chemical ratio of 20:7  $\text{BCl}_3/\text{Cl}_2$ . This ratio was chosen because it is mostly non-selective (i.e. etches at the same rate) for AlAs and GaAs and yet also has fairly high selectivity over the photoresist mask. The mixture of the two also produces smoother etched surfaces and high anisotropy (see chapter 5 of Ref. [101]). This is because the heavier  $\text{BCl}_3$  ions increase the sputtering effect of the etch, which helps to remove the oxides for both materials and especially the more readily-oxidized aluminum. The non-selective feature allows for easier timing of the etching to obtain a desired etch depth, since the etch rate is fairly constant through the many DBR and cavity layers. The other parameters of the ICP RIE can be used to further tune the selectivity by changing the balance of the chemical etching with the sputtering. Increasing the ICP power essentially increases the plasma density, which increases the chemical rate. The chamber pressure has a similar effect. The RF power between the capacitor plates controls the overall kinetic energy of the ions, so increasing it will increase the sputtering rate. These parameters were similarly chosen to achieve the results of equal rates between all of the sample materials, high anisotropy, and high etch selectivity over the mask.

The etch selectivity of the sample over the photoresist mask was about 3:1. We therefore needed a photoresist mask about one-third the thickness of the desired etch depth. The actual etch mask thickness was about 2700 nm. This was partially to allow for longer etch times without worrying about the mask being completely removed in thin places. Additionally, the actual profile of the edge of a photoresist mask is often curved, leaving thinner regions near

the edges. If those thinner regions are etched through, they will allow the material beneath to begin to be etched, resulting in similarly curved sidewall profiles in the sample. By using a thicker mask, this effect is prevented while sacrificing overall minimum feature-size.

As shown in Figure 2.3, this etch process was used to produce pillars, with the top DBR almost entirely removed in the unmasked regions. The same method was used for a shallower etch to remove the quantum wells, leaving them only in a narrow linear region with nearly-exposed quantum wells on both sides of the pillar. While a thinner photoresist could have been used for this etch, it was overall not a critical step in terms of feature size so the same process was more than adequate. The result was a linear “mesa” region of accessible quantum wells with a microcavity pillar in the middle. The precise details of the etch are given in Appendix C.

### 2.2.3 Electrical contacts

The next step was to place electrical contacts on the regions of the structure with exposed quantum wells next to each pillar. The overall recipe was based on ohmic contacts described in Appendix B of Ref. [103]. The basic steps are shown in the bottom row of Figure 2.3. This lithography uses an additional layer of lift-off resist (LOR) below the photoresist layer. The LOR is undercut under the photoresist layer during development, causing an overhang. This encourages a break to form in a deposited thin film, allowing consistent lift-off (see Section C.3). Furthermore, since no etching was necessary, a thinner  $\approx 1200$  nm layer of photoresist is used for better control of the feature size.

After lithography, a metal stack, consisting of 5 nm of Ni, 120 nm of  $\text{Au}_{0.88}\text{Ge}_{0.12}$ , 25 nm of Ni, and 150 nm of Au, is deposited on the sample in that order. The first thin layer of nickel assists with bonding the stack to the material surface, and also allows penetration of the material during annealing. The AuGe forms an ohmic contact with GaAs, with  $n$ -type doping of the GaAs via the germanium [101]. The second layer of nickel separates the alloy from the pure gold layer, preventing mixing during annealing since the annealing temperature is close to the alloy’s melting point. The final gold layer provides a thick and clean contact surface for wirebonding. After deposition, a lift-off step removes the metal

everywhere except in the desired location on the mesa structure. Annealing at 410°C allows the AuGe to penetrate the remaining thin layers above the quantum wells and make contact with the quantum wells themselves. Appendix C contains further details.

## 2.3 EXPERIMENTAL METHODS

### 2.3.1 Optical pumping

There are essentially two methods of optical pumping: resonant and nonresonant. Resonant pumping simply means that the photons used to excite a particular mode match the mode in energy and momentum. Nonresonant pumping is simply when the photon does not match the mode in energy and momentum. Each type of pumping has distinct advantages and disadvantages, and the type of pumping is usually determined by the goals of the experiment. I used nonresonant pumping for all of the work reported in this thesis, but I will describe both to explain why that was the correct choice.

**2.3.1.1 Resonant pumping** Resonant pumping is overall the most simple. As discussed in Section 1.2, cavity photon modes can decay by coupling into the external photon modes. That process conserves energy and momentum, such that the in-plane components of the wavevectors of the external photon and the cavity photon are equal (classically, this is from the boundary conditions of Maxwell’s equations discussed in Appendix B). Since the cavity photons are clearly able to decay into external photons, the reverse process must also be possible. Thus, resonant pumping is simply done by matching the energy and in-plane momentum of the external pump photons to the desired cavity photon mode. This is shown in Figure 2.4(a), which gives a reflectivity curve for a DBR cavity, in this case with a much lower  $Q$ -factor than our actual devices in order to clearly show the cavity mode. In high  $Q$  cavities, the dip is equally as deep at resonance, but with a much narrower energy linewidth since the lifetime is much higher (see Section 1.2.3). While the case shown is for a cavity without a quantum well, the same principle holds for a polariton sample. This makes sense

since the polariton modes effectively shift the original modes to higher or lower energy, and both the lower polariton and upper polariton are partially photonic.

Figure 2.4(b) illustrates the different pumping methods from the perspective of the dispersions. Resonant pumping is simply matched to overlap with a particular point on the LP branch. This also allows the creation of polaritons at a particular momentum state, preserving the in-plane momentum of the incoming photons (which can be changed by simply altering the incident angle of the pump). Since this pumping corresponds to a direct interaction that creates a polariton and destroys a photon, it also preserves the spin state of the photon. Thus, the spin of the created polaritons can be controlled by switching between different directions (and degrees) of circular polarization. This allows for studies involving very specific initial states of the polaritons.

One downside of resonant pumping is that it is difficult to get a population of particles moving all in the same direction at the same velocity to thermalize, since collisions between polaritons are minimized. This can make it very difficult to reach any state even remotely resembling equilibrium. Thermalization can still happen via collisions with intrinsic free electrons, but this is highly dependent on the exciton fraction since it is fundamentally an exciton-electron interaction (see Section 1.4.7 and Ref. [33]). Another disadvantage is that the photoluminescence from the polaritons is difficult to distinguish from the pump. Since there is always some amount of the scattering, and usually lots of reflection, from the sample surface and the optics in the experiment, resonant pumping can make detection of the polaritons difficult. In fact, since the in-plane momentum is conserved, the majority of the emission from the polaritons will even match the reflected pump in emission angle, often eliminating the possibility of removing the pump with Fourier plane ( $k$ -space) filtering (see Section 2.3.2). This can be overcome with real space filtering, as long as the expected emission is far from the pumping location. However, this also limits the experiment to using large incoming angle, corresponding to large in-plane momentum which carries the polaritons to a different area of the sample. This eliminates the possibility of producing a quasi-equilibrium population by pumping directly into the  $k_{\parallel} = 0$  state. Furthermore, coupling a pump laser into the lower polariton mode can be very difficult simply because of the extremely narrow transmission line, which near resonance or at photonic detunings

is similar in width to the pure cavity line (See Section 1.2.3. For an example of a very narrow cavity line, see Figure 1.6, though our samples have even narrower transmission lines than the one shown since we use twice as many DBR periods with similar refractive index contrast.). This means that much of the input light from the pump, which always has some spread in energy, will see very high reflectance, with only the narrow part that overlaps exactly with the resonant transmission line being coupled into the sample. This greatly reduces the efficiency of the pumping, especially for very high- $Q$  cavities. That being said, there are a lot of very interesting things that can be done with resonant pumping, such as observing ballistic flow [25], producing solitons [104–106], and observing spin-dependent polariton-polariton interactions [107, 108].

**2.3.1.2 Nonresonant pumping** Some of the disadvantages of resonant pumping can be overcome with nonresonant pumping. Technically, this can be done by pumping any state at higher energy than the mode of interest. Practically, there is still the necessity of efficiently coupling into the device. As shown in Figure 2.4(a), nonresonant pumping usually aims for the dips in reflectance of the DBRs; anywhere else, especially in the DBR stopband, results in most of the pump light being reflected by the sample and few polaritons are produced. Nonresonant pumping in these systems is often described by which stopband dip is targeted (e.g. the first or second DBR dip above the stopband).

The fine details of how a photon goes into the sample at much higher energy and eventually becomes a polariton are complicated, as they are in most solid-state many-body processes. A simplistic picture is provided in Figure 2.4(b). Above the coupled lowest-energy exciton mode, there are many other modes of the sample (the area of the figure shaded in gray). These include the higher order quantum well states, the higher order hydrogen-like exciton states (which eventually become a continuum), and the many free electron conduction band states. The upper polariton can itself be considered an available nonresonant state. These states can couple to each other and to the phonon modes, allowing efficient and fast ( $\sim 1 - 10$  ps [35]) settling into the LP mode. The primary advantage of this process is that most of the original information about the pump is lost, due to many scatterings with the phonon modes of the lattice (The illustration in Figure 2.4(b) is obviously simplified since the

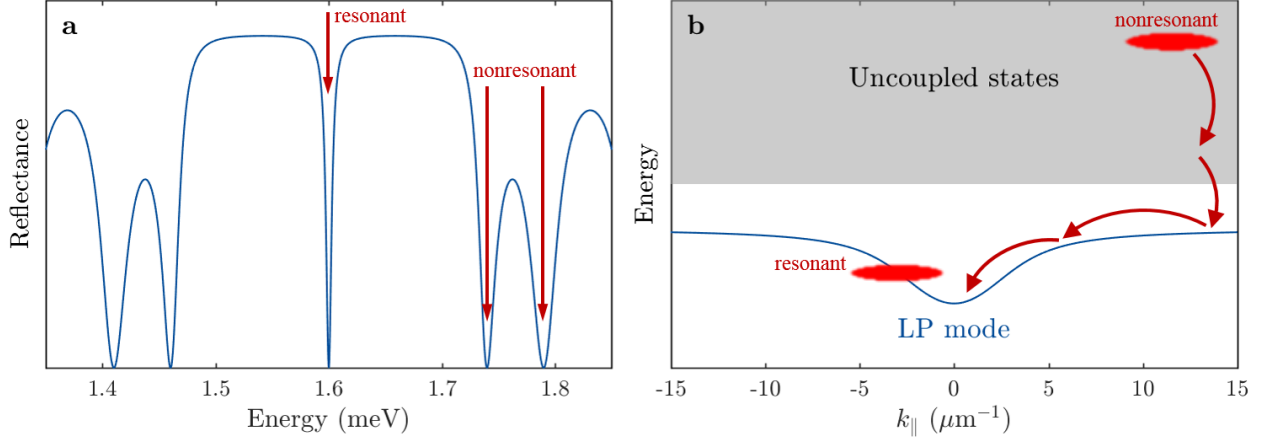


Figure 2.4: Comparison of resonant and nonresonant pumping. (a) A typical plot of reflectance as a function of energy for a DBR cavity, produced using the transfer matrix method, similar to those depicted in Figure 1.6. The cavity mode is clearly visible in the center, and is rather wide in this case due to the low number of DBR periods used for the simulation. Resonant pumping would be matched to the reflectance dip at the cavity mode, while nonresonant pumping is done at higher energy. A polariton cavity would look essentially the same, but with dips in reflectance at the LP and UP modes. The plot was made using 8.5 DBR periods on both the top and bottom, with a cavity refractive index of 3.4, and DBR refractive indices of 3.0 and 3.5. (b) A schematic of the pumping process. Exciton and free electron states not significantly coupled to the cavity mode would fill the shaded region, somewhere above the coupled lowest-energy exciton state. Nonresonant pumping targets these states, while resonant pumping is matched to the LP mode.

cooling will not follow a single path in  $k$ -space.). This allows for formation of a population at and near  $k_{\parallel} = 0$ , approaching something like equilibrium distributions (see Ref. [26] for a particularly nice example of measurements of these distributions). The polarization/spin and momentum information is also mostly lost in the process. Thus, nonresonant pumping is great for creating a population of particles with somewhat randomized initial states, much like normal thermodynamic systems, and is therefore ideal for studying condensation (as was done in the first observations of condensates in polaritons [13, 14]), or the spontaneous formation of vortices [21, 23].

Experimentally, nonresonant pumping avoids the “blinding” effect of pumping at the same energy and angle as the desired signal. The imaging optics can easily contain a long-pass filter to remove the scattered and reflected pump light, the angle of the reflected light can be filtered out, and the collection of the reflected pump light can be avoided entirely by using an incident angle outside the numerical aperture of the objective lens. Nonresonant pumping also tends to lead to more efficient coupling of the pump light into the system. As discussed above for resonant pumping, the narrow transmission line at the cavity resonance leads to high reflectance for much of the input light, but the dips in the reflectivity away from the DBR stopband are much broader, allowing most of the light from the pump to enter the system. The inability to set the initial state can also be a disadvantage in some cases. An important effect that can be either an advantage or disadvantage is the build-up of a population of hot excitons. Since the settling process involves first coupling to the exciton states (and LP states that are essentially excitons), there will always be a population of heavy, mostly immobile excitons at the location of the pump spot. This is further exacerbated by the “bottleneck” effect, which is due to the very low density of states (given by  $\frac{dE}{dk}$ ) near the inflection point in the LP band (see Ref. [62] for an early discussion of the difficulties this caused). These excitons are obviously capable of interacting with the exciton part of the polaritons (see Section 1.4.7), which has two effects. The first is that it blueshifts the polariton mode [35] (due to blueshifting of the underlying exciton mode by exciton-exciton interactions as discussed in Section 1.4.7), which can greatly alter the polariton state. The second is that it strongly repels the polaritons from the pumpspot. As discussed in Section 2.2.1, this can be harnessed to trap or otherwise manipulate the flow of the

polaritons, but it is undesirable in certain situations, such as attempts to measure polariton-polariton interaction strength (see Chapter 6 and especially Refs. [59, 85].

Since most of the experiments I performed were intended to see either the effects of a thermalized condensate of polaritons, or to specifically study the excitons produced during nonresonant pumping, most of the work reported in this thesis was done with a nonresonant pump.

### 2.3.2 Optics and imaging

The primary means of gathering data in my experiments were conventional optics, so I will describe the specific methods in this section. There are essentially two types of imaging that I used: real-space and  $k$ -space. Very simple optical arrangements for each of these are shown in Figure 2.5. In the case of the real space imaging in panel (a), the object will be imaged into the camera with magnification  $M_R = f_2/f_1$ . For an ideal lens, there is always a  $k$ -space plane (also called Fourier, far-field, or angle-resolved plane) exactly one focal length from the back side onto which the light incident on the front of the lens is sorted according to incident angle. Thus, a lens effectively performs a Fourier transform on incident light, mapping angle information onto real-space (this is discussed in many optics and photonics textbooks, e.g. Ref. [109]). This effect can be utilized to gather angle information directly into an image. This is shown in Figure 2.5(b), where two other lenses are arranged to image the  $k$ -space plane of the objective lens, with magnification  $M_k = f_3/f_2'$ .

A more useful arrangement is shown in Figure 2.6. In this case, both the real- and  $k$ -space planes are imaged at an intermediate point for both real- and  $k$ -space final imaging. This allows both position and angle filtering to be done in either imaging setup at an intermediate plane. Since the image at  $R'$  is often magnified by  $M_R = f_2/f_1$ , this allows better resolution for filtering in real-space, and is often the only method when the object (i.e. the sample) is deep within the cryostat. It also allows for position selection without affecting the  $k$ -space optics, which is essential for some experiments, including the work reported in Chapter 6.

The secondary  $k$ -space plane at  $K'$  is usually even more useful. First of all, imaging this plane allows real-space filtering at  $R'$  without changing the  $k$ -space imaging optics. Also,

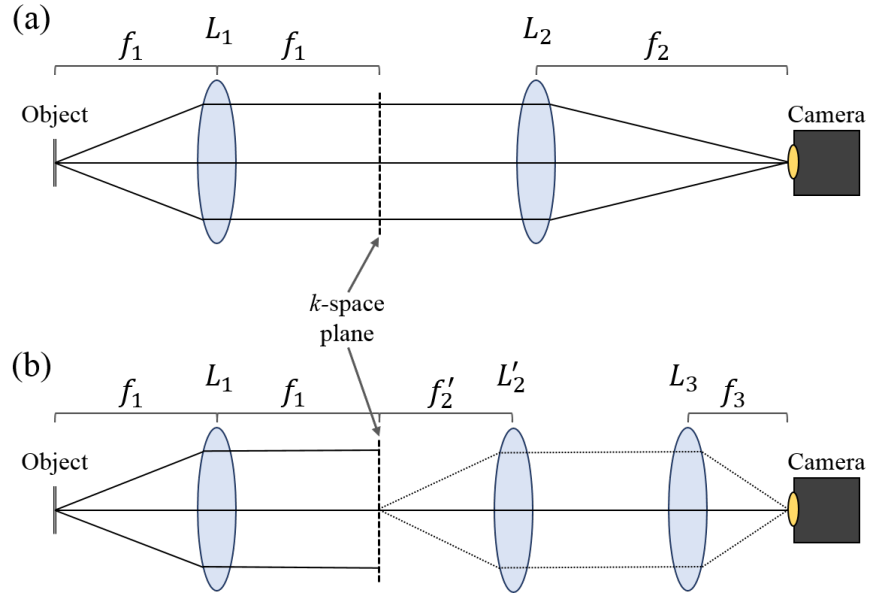


Figure 2.5: Basic optical setups. (a) A real space imaging setup, with an objective lens  $L_1$  and an imaging lens  $L_2$ . (b) A  $k$ -space imaging setup, with a different set of lenses  $L'_2$  and  $L_3$  imaging the  $k$ -space plane of the objective lens  $L_1$ .

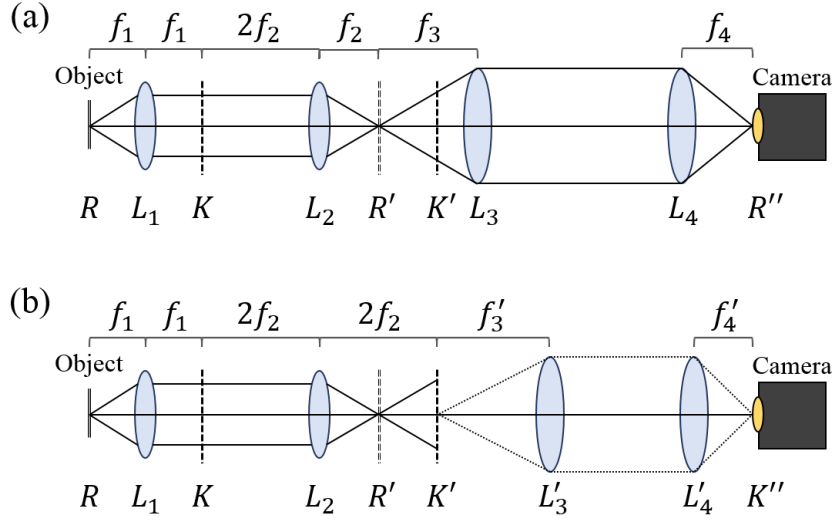


Figure 2.6: More typical optical arrangements for experiments using  $k$ -space imaging or filtering. (a) Real space imaging, with the original real-space plane  $R$  imaged twice, at both an intermediate real-space plane  $R'$  and at the camera at  $R''$ . (b)  $k$ -space imaging, with the original  $k$ -space plane  $K$  image twice, at both an intermediate  $k$ -space plane  $K'$  and at the camera, which is now  $K''$ . The distance between  $L_2$  and  $K$  need not be  $2f_2$ , but that arrangement gives a nice amount of space between  $R'$  and  $K'$  while also avoiding magnifying at  $K'$ .

while Figure 2.6 depicts the use of simple singlet lenses, the lens  $L_1$  is often a microscope objective. These usually have their original  $k$ -space plane  $K$  embedded within the body of the lens itself, making it inaccessible for filtering. Thus,  $k$ -space filtering for real-space imaging must be done at the plane  $K'$ . If a singlet is used for the objective,  $k$ -space filtering can be done at either plane.

Furthermore, these overall arrangements offer more flexibility in choosing magnifications or inserting other optics. In terms of magnification, for a given choice of  $L_1$  and  $L_2$ , I usually need to have large real-space magnification while also getting significant  $k$ -space demagnification. This is because the scale of the devices I have studied is  $\sim 10 \mu\text{m}$ , while the  $k$ -space plane of the objective is a rather large object, typically about the size of the objective lens diameter ( $\sim 10 - 100 \text{ mm}$ ). The ability to choose lenses  $L_3$  and  $L_4$  independently of  $L'_3$  and  $L'_4$  offers greater flexibility. As an added advantage, if the spaces between each of those pairs of lenses is large enough and has enough overlap, a large optic that is ideally placed in a region where the light rays are collimated, such as a large Dove prism, can be used in the same location for both setups. For diagrams and details of the optical setups for actual experiments, see Appendix D.

### 2.3.3 Spectroscopy and detection

The two primary devices I used for optical measurement, other than basic optical components, were a charge-coupled device (CCD) camera (mostly a Roper 512b or Roper 1k by Cascade Microtech) and spectrometer (a Princeton Instruments SpectraPro 2500i).

A CCD is essentially just a photosensitive material (typically  $p$ -doped silicon) backed with an array of gates separated from the photosensitive material by an insulating layer (much like an array of MOSFETS). With the gates set to a positive voltage (creating a depletion region in the  $p$ -type semiconductor), incoming photons excite electrons in the photosensitive material, which fall into the potential wells produced by the gates. By using more than one gate per pixel, the “clumps” of stored charge can be shifted within the semiconductor while maintaining their relative locations in the array. The charge values can therefore be systematically shifted to a readout circuit, and the total stored charge determined for each

pixel location. Since the stored charge is proportional to the number of input photons (with the proportionality given by the quantum efficiency), the relative number of photons at each image point can be measured. If the quantum efficiency is known, then the actual number of photons is also known. Of course, for an actual experiment, there are many other optics between the sample or device and the CCD, so photon counting must also consider the efficiency of the entire optical setup.

A spectrometer is basically a diffraction grating with some means of carefully measuring the diffraction angle. Many older spectrometers were just a fixed grating with a telescope mounted to a swinging arm. Most modern spectrometers consist of five essential components: An entrance slit, a collimating concave mirror, a reflective diffraction grating, a rotating turret on which the grating is mounted, and a focusing concave mirror. An image is formed on the entrance slit, which selects a narrow slice (parallel to the grating lines), and then is collimated again by the first concave mirror. With the photons traveling parallel to each other, they intersect the diffraction grating and are diffracted according to the standard equation

$$\sin(\theta_i - \theta_m) = m\lambda\rho, \quad (2.1)$$

where  $\rho$  is the grating line density,  $\lambda$  is the wavelength of the photons,  $m$  is an integer,  $\theta_i$  is the incident angle on the grating, and  $\theta_m$  is the angle of the  $m^{\text{th}}$  diffraction order. The image is then re-formed by the final concave mirror, maintaining the original image information along the direction of the entrance slit and spread in the orthogonal direction in energy by the grating. By changing the angle of the grating using the turret, the angle corresponding to a specific photon energy can be shifted, shifting its location on the imaging plane as well.

A spectrometer therefore exchanges the original image information along one direction (usually either real-space position or emission angle) for photon energy. A CCD can be placed at the final imaging plane to record the image, which is the setup used for most of my work. Other options include sending the light into a streak camera to get temporal information, or using another spectrometer or series of spectrometers (e.g. to get better signal-to-noise if there is noise at energies close to the signal).

Most spectrometer setups utilize the first-order diffraction. A blazed grating is therefore typically used in order to collect as much light as possible into the first-order. This design

essentially shapes the periodic pattern on the grating such that the individual surfaces are close to the regular spectral reflection condition (reflected angle equals incident angle) on the microscopic level for the first-order diffraction (which is dependent on the angle of the entire periodic structure). This is obviously wavelength dependent, so the “blaze angle” is always associated with a particular wavelength. Additionally, a spectrometer operating in the 0<sup>th</sup> order mode acts effectively like a mirror. Thus, an optical setup can be easily switched between spectrally-resolving (i.e. energy-resolving) and collecting the entire original image by simply turning the grating to the zeroth-order position and opening the entrance slit. This is useful for alignment of the spectrometer slit in real-space, and can be used for full real-space or  $k$ -space imaging without making any changes to the optics. The disadvantage of this is that a grating in this mode is really more like a low quality reflecting polarizer, which also reflects very poorly because most of the light is sent into the first-order diffraction.

#### 2.3.4 Electrical measurements

The general goal of the electrical measurements was to look for a response in the current-voltage characteristics to the presence of a polariton condensate. For a truly continuous wave (CW) experiment, this would be as easy as pumping optically and measuring the current while sweeping the voltage. However, due to heating at the pump powers necessary to see the full desired range of polariton density, chopping of the laser was required. I therefore used a mechanical chopper at 400 Hz and 2.4% duty cycle in most cases, producing a 60  $\mu$ s pulse every 2.5 ms. As discussed in Section 3.2, the optical pumping was still essentially CW due to the very fast time scales relevant to polariton thermalization and lifetime.

As discussed in Chapter 4, the current through the device was extremely low without illumination. Therefore, any reasonable signal would essentially be scaled by a factor of 0.024 simply because the sample was not illuminated for most of the time. To prevent this, I synchronized the electrical measurement with the quasi-CW laser pulses. The scheme I used is depicted in Figure 2.7. A pick-off was placed in the beam path (with about 3% reflectivity) following the chopper and before any filtering optics. This was sent to an amplified photodiode. The photodiode signal was sent to an inverting amplifier (using a

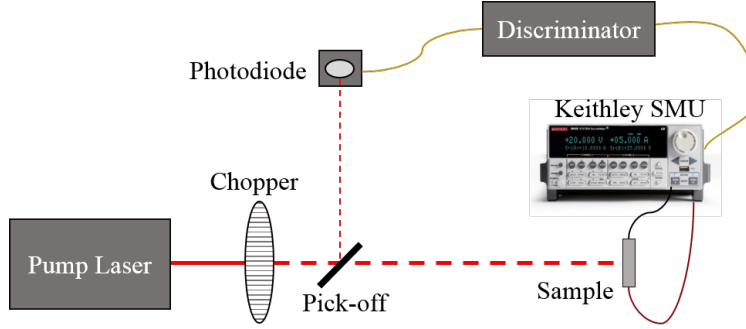


Figure 2.7: Diagram of the synchronization setup for optical and electrical measurements.

standard 741 op amp) to further amplify the signal and switch to negative-voltage pulses. A discriminator was then used to select the rising edge of the signal, returning a corresponding output pulse which (after conversion to +5 V), was sent to a Keithley 2636B source-meter unit. No delay line was necessary for the microsecond time scales of the pulses.

With an input signal sent to the Keithley for every laser pulse, I was able to time the measurements with the illumination. This is shown in Figure 2.8. For most of my experiments, the voltage was initially set to some value. Then, at least partially simultaneously (depending on the required integration time of the camera) emission was measured with the optics and CCD while the Keithley collected current measurements. The current-measurement integration time was set to begin as soon as an input signal was received and end before the pump pulse ended. The Keithley integration time is set in units of number of power line cycles (NPLC), with a minimum of 0.001 NPLC. Since a power line cycle corresponds to 0.0167 seconds (for 60 Hz), the minimum was  $16.67 \mu\text{s}$ . I therefore set the integration time to three of these, or  $50 \mu\text{s}$ , which fits within the  $60 \mu\text{s}$  illumination periods.

One problem with this method is that electrical measurements with a meter of the type used are very unstable over periods under a single power line cycle. Obvious oscillations at 60 Hz can be observed for consecutive measurements at 2.5 ms intervals. To avoid this problem, at least 250, and often many more, current measurements were taken for each voltage setting, corresponding to a total actual integration time of 12.5 ms and spanning a total period of 37.5 power line cycles. After a measurement at a given voltage, the voltage

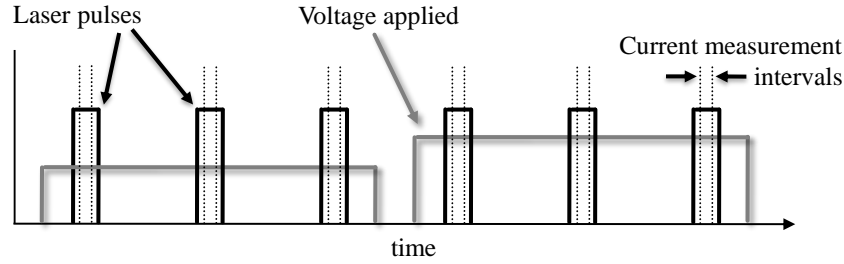


Figure 2.8: Diagram of the electrical measurement timing. The measurement intervals were timed to occur simultaneously with the illumination by each pulse. A DC voltage was applied continuously during each measurement set.

was progressed to the next setting and the measurement repeated.

During some of my early measurements, I also applied voltage pulses rather than using a constant set voltage. This was out of concern for heating with the applied current and voltage, but was ultimately unnecessary due to the very low current and voltage levels used in the devices. Those measurements typically turned on the voltage about 0.5 ms before the arrival of the optical pump pulse, and then turned it off again about 0.5 ms afterward. Through comparative studies, I determined that this had no noticeable effect on the actual measurements, so I dropped that method in favor of the one described above due to its simplicity in programming and lower rate of producing errors in the timing circuit of the Keithley unit.

## 3.0 TRAPPING AND CONDENSATION IN ETCHED PILLARS

### 3.1 MOTIVATION

The work in this chapter came about as part of the larger goal of injecting electrical current into a trapped polariton condensate, the results of which are given in Chapters 4 and 5. As discussed in Section 2.2, since direct contact to the quantum wells was already desired, trapping the polaritons by etching away the top mirror seemed to be the best option. The first step, then, was to produce etched “pillars” and characterize the effect of the etch upon the polaritons. As will be shown in this chapter, the result was indeed an overall trapping within the pillar and the production of a polariton condensate, but with the additional effect of a narrow trap a few microns from any etched edge

The work described in this chapter was previously published in Ref. [110].

### 3.2 THE EXPERIMENT

For details on the methods mentioned here, see Section 2.3. Since the eventual goal was to look for electrical responses due to polariton condensates, I used a chopped, continuous-wave (CW), nonresonant pump. The chopping was to prevent heating of the sample, which was necessary since the CW powers required to achieve condensation were  $\approx 100$  mW. The CW pump was necessary since a quasi-stable system (which these quickly interacting, short lifetime polaritons are on any scale  $\gtrsim 100$  ps) should be helpful in looking for a potentially slow electrical response. Finally, a nonresonant pump is necessary to produce a true thermalized condensate.

Except where noted, the optical pumping was done at  $E_{\text{pump}} = 1754$  meV, at an incident angle of approximately  $18^\circ$  with Ti:sapphire pump laser. The sample was kept at constant temperature in a liquid helium cryostat ( $\approx 5$  K). The chopping was done mechanically with a 1.3% duty cycle at 400 Hz. This corresponded to pulse widths of about  $30 \mu\text{s}$ , which effectively gave short bursts of CW pumping since  $30 \mu\text{s}$  is  $10^6$  times greater than the relevant time scales for polariton thermalization and decay. The powers are reported as the peak of each short pulse, i.e. the power if the measurement were taken only during the  $30 \mu\text{s}$  that the pulse was on (average power divided by duty cycle). The primary data were extracted from energy-resolved images, taken using a CCD camera on the output of a spectrometer (0.5 m path length, 1200 gr/mm grating). We also obtained in-plane momentum information using angle-resolved imaging. Except in the case of the angle-resolved imaging, the collection angle was limited to  $\approx 4^\circ$  from normal emission. This meant that the PL gathered in most of the images corresponded to the low momentum thermalized states, and especially to the region in momentum populated by a condensate.

The imaging optics were mostly constrained by two requirements: long working distance for the objective and high magnification. Because I was working in a large cryostat (a Janis SVT-300), the sample was already about 40 mm from the outside surface of the window. Since I wanted to be able to push the pump power well over threshold, I needed as much power as possible to reach the sample. This is most easily accomplished by reflecting the pump off of a mirror placed in front of the objective (or pumping at a steep angle past the objective), rather than reflecting the pump off of a beam splitter and sending it through the objective to the sample. This further increased the space requirements between the sample and the objective. I therefore settled on a plano-convex singlet with a 75 mm focal length and 2 inch diameter. Behind this was placed a similar lens with 150 mm focal length, creating a secondary real-space plane (a regular image). In order to get high magnification, a single 60 mm lens of the same type was used in a single-lens imaging arrangement, giving a total magnification of about 25. This allowed a spatial resolution of  $\sim 1 \mu\text{m}$  (The CCD pixel size was  $16 \mu\text{m}$ , but there was also blurring from aberrations.). Emission-angle ( $k$ -space, Fourier-plane) filtering was done using an iris placed directly at the backside focal plane of the objective. The angle-resolved images were taken by imaging the same Fourier

plane directly with a second collimating-focusing lens pair that could be switched out in place of the final 60 mm real-space imaging lens. One other important optical element was a Dove prism, placed between the objective and the final imaging lens. The Dove prism allowed rotation of the image plane through any angle, which allowed any desired slice of the sample (or slice of  $k$ -space in the case of angle resolved imaging) to be aligned with the spectrometer slit. Overall, the lateral (in-plane) position of the real space imaging could be adjusted by small lateral movements of the objective, while the *orientation angle* of the real space imaging could be adjusted with the Dove prism.

The sample used in this study is described in Section 2.1. The post-growth fabrication in this case was rather simple, and is shown in Figure 3.1(b). An scanning electron microscope (SEM) image of a similar pillar (identical other than not being the exact pillar studied) is shown in panel (a) of the same figure. It consisted of an  $\approx 100 \mu\text{m} \times 100 \mu\text{m}$  region where the top DBR was removed, eliminating the cavity photon mode in the etched regions. Though the pillar was relatively large (in terms of microscopy), spatially resolving the PL from different regions of the pillar did require the fairly high magnification discussed above. To form a condensate, the pillars were pumped by a laser spot of  $\approx 20 \mu\text{m}$  FWHM directly in the center of the square. The detuning of the pillar, unless otherwise specified, was about  $-3 \text{ meV}$ .

### 3.3 RESULTS

The first step was to establish an overall potential landscape of the pillars. Since this corresponds to the ground-state energy of the LP mode, it was done by using a defocused pump and collecting energy-resolved slices of the pillar at many locations. Parallel slices could be gathered by shifting the position of the objective, while the slice angle could be adjusted with the Dove prism. The interpolated results are shown in Figure 3.1(c), showing the overall effect of the etch process on the ground-state energy of the lower polaritons. A low energy region is apparent at all of the edges, with the minima  $\approx 7 \mu\text{m}$  in from the outer etched edge, and overall energy minima appearing at the corners. The depth relative to the

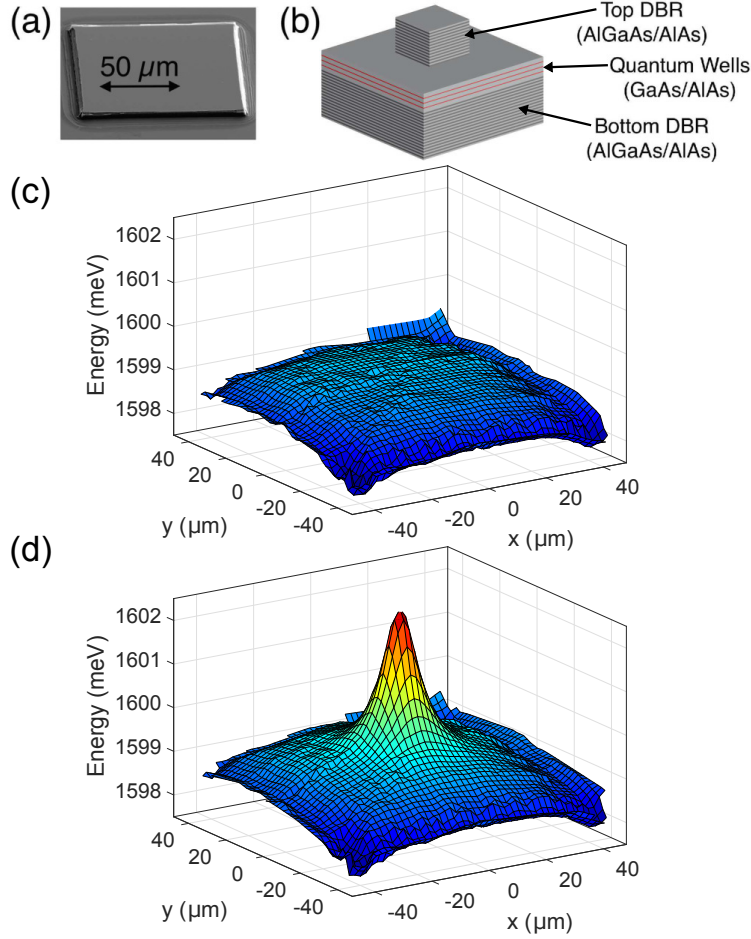


Figure 3.1: (a) SEM image of an etched pillar. (b) The etched pillar structure, with the top DBR removed except in a square area. The bottom DBR and QW layers are un-etched, leaving an open cavity except at the location of the pillar. (c) Linearly interpolated data showing the resulting low density potential landscape. Since the pillar was a square with  $\approx 100\ \mu\text{m}$  long sides, these data come from entirely within the pillar and out to about  $5\ \mu\text{m}$  from the edges. A mostly flat region in the middle gives way to low energy regions at the edges, with minima at the corners. The overall tilt of the landscape is due to the sample gradient. (d) The same surface shown in (c) with a fit of the effective potential from a focused cw pump laser at 51 mW added in the center.

flat region in the middle is  $0.5 - 1.5$  meV. This in itself is much shallower than the trapping attained from the smaller scale etching or the cavity spacing methods described in Section 2.2.1, but comparable to the depths attained from stress traps. The overall tilt shows the direction of the cavity gradient (approximately  $75^\circ$  from the  $+x$ -axis) resulting from the wedge in the cavity thickness (see Section 2.1).

The effect of lower energy near the edges was actually observed in other independent work [93, 94]. However, those works only mentioned the effect and did not study it in detail, as I do here.

The potential landscape changed with the addition of a focused pump. The repulsive exciton-exciton interactions formed a new effective potential, with an energy “hill” which is highest at the location of the pump spot [Figures 3.1(d) and 3.2(a)]. This hill became higher at high pump power due to the increase in the exciton density. This is expected given the density-dependent blue shift due to interactions (see Section 1.4.7, and Refs. [78, 92]). Figure 3.2 shows the polariton luminescence of a real space slice across the middle of the pillar containing the pump spot, which is apparent in the center, at several powers. This particular slice follows the  $y$ -axis and intersects the origin, using the axes defined in Figure 3.1(c). At low pump power, the polariton luminescence was clearly visible at all points across the pillar, but the highest density was at the pump spot [Figure 3.2(a)]. The pump spot was even more dominant at higher powers [Figure 3.2(b)]. The density distribution underwent a dramatic shift at even higher powers, with the edge regions emitting comparable intensities to the central pump spot [Figure 3.2(c)]. The narrow spectral width of the emission from these regions indicates that they contained Bose condensates. This is also consistent with the dramatically increased transport distance, which can be associated with superfluidity that allowed the polaritons to find the local minima of the potential landscape. At the highest powers, the condensates in the edge regions were blue-shifted to higher energies, eventually becoming equal in energy [Figure 3.2(d)], consistent with a phase-locked single condensate.

A similar story emerges from the full real-space imaging of the entire pillar at various pump powers (Figure 3.3). The pump spot was prominent at low powers, but polariton luminescence was visible across the whole pillar [Figure 3.3(a)]. The relatively bright luminescence at the edges and around the pillar was from the bare excitons in the open cavity.

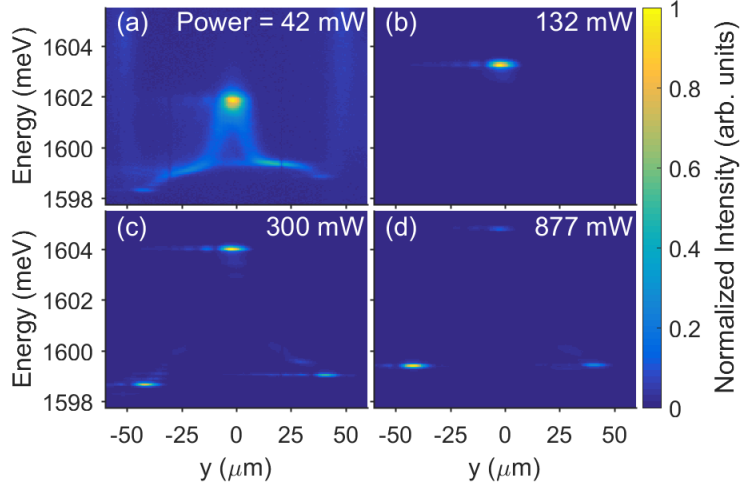


Figure 3.2: Energy-resolved polariton luminescence intensity near  $k_{\parallel} = 0$  for a slice across the center of the pillar [along the  $y$ -axis defined in Figure 3.1(c)]. The peak pump power (the power of each quasi-cw pulse) for each image was (a) 42 mW, (b) 132 mW, (c) 300 mW, and (d) 877 mW. The vertical streaks at both edges in (a) are from bare exciton luminescence. The horizontal streaks near bright regions are an artifact of the imaging optics. The intensity is normalized separately for each image, so the values are not comparable between images.

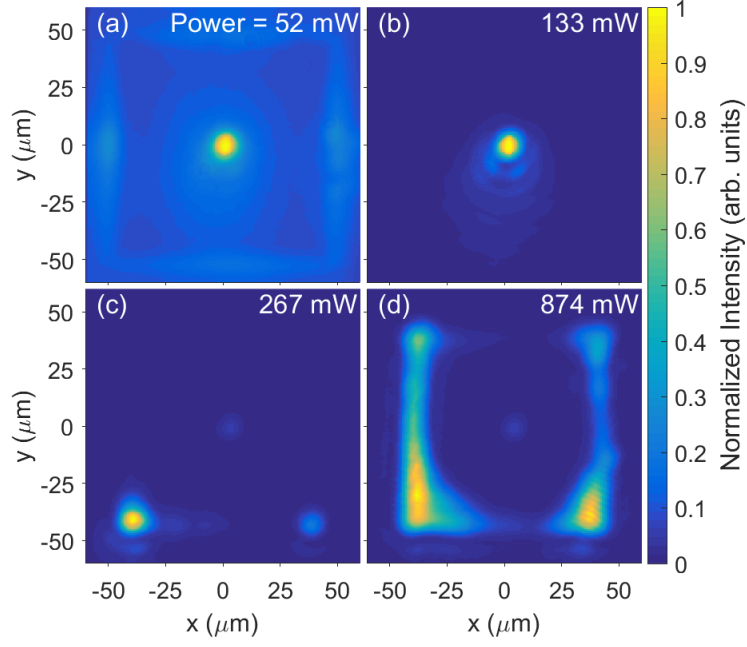


Figure 3.3: Full real space polariton luminescence intensity near  $k_{\parallel} = 0$  at peak pump powers (the power of each quasi-cw pulse) of (a) 52 mW, (b) 133 mW, (c) 267 mW, and (d) 874 mW. The coordinate axes are the same as those defined in Figure 3.1(c). The intensity was normalized separately for each image, so the values are not comparable between images.

As the power was increased [Figure 3.3(b)] the pump spot became even more dominant, but the polaritons remained in the non-condensate regime and their diffusion length remained low. At higher powers [Figure 3.3(c)], the corners became the dominant features as polaritons flowed long distances to fill in the locations of the energy minima. As the power was increased further [Figure 3.3(d)], the polaritons filled a connected region extending along all the edges. The asymmetry relative to the pillar shape was caused by the cavity gradient, with the side closest to  $y = -50 \mu\text{m}$  at lower overall energy, and the lower left corner ( $x = -50 \mu\text{m}$ ) at the lowest point.

Figure 3.4 shows the angle-resolved polariton luminescence, with the measured angle corresponding to momentum along the direction close to the  $x$ -axis defined in Figure 3.1(c). These data were collected from a different but nearly identical pillar (in terms of etch parameters, size, detuning, and source sample). The pump in this case was at  $\approx 1759 \text{ meV}$  and at normal incidence (the pump was sent through the imaging objective). This resulted in considerably different power levels for the same luminescence characteristics, but the overall characteristics were the same. At low pump power [Figure 3.4(a)] the polaritons were emitted at many angles fairly evenly, corresponding to the filling of many momentum states. At higher power, high-energy emission [near and above  $1601 \text{ meV}$  in panel (b)] emerged from the top of the repulsive “hill” at the pump spot, but the dispersion remained generally the same. Above the condensate threshold power [Figure 3.4(c)] the polaritons moved overwhelmingly into a narrow low-energy state near  $k_{\parallel} = 0$ , which is characteristic of Bose-Einstein condensation. The gap between the condensate at lower energy and the parabolic dispersion at higher energy is due to collecting from the entire pillar: The condensate is located in the low energy traps at the edges, while the parabolic dispersion is emitted primarily from the higher energy region in the middle. The multiple parabolas are due to emission from different parts of the middle region that are also at slightly different energies. At even higher pump powers [Figure 3.4(d)] the parabolic polariton dispersion became nearly undetectable compared to the luminescence at  $k_{\parallel} = 0$ , indicating a large condensate fraction. The condensate also exhibited a slight blue shift in energy due to the high polariton density along the edges of the pillar.

To determine the precise threshold power, we plotted the polariton luminescence intensity

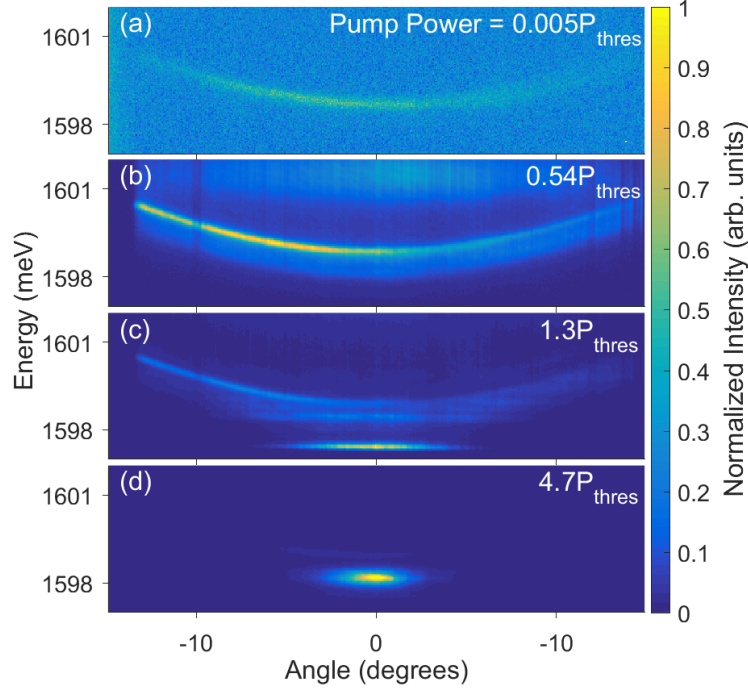


Figure 3.4: Energy- and angle-resolved polariton emission at various pump powers. The collection angle was along the direction close to the x-axis defined in Figure 3.1(c), but the luminescence was collected from the entire pillar. The pillar was pumped at  $\approx 1759$  meV and at normal incidence, with a threshold power ( $P_{\text{thres}}$ ) of  $\approx 86$  mW, as defined similarly to those in Figure 3.5. The pump powers are (a)  $0.005P_{\text{thres}}$ , (b)  $0.54P_{\text{thres}}$ , (c)  $1.3P_{\text{thres}}$ , and (d)  $4.7P_{\text{thres}}$ . The intensity was normalized separately for each image, so the values are not comparable between images.

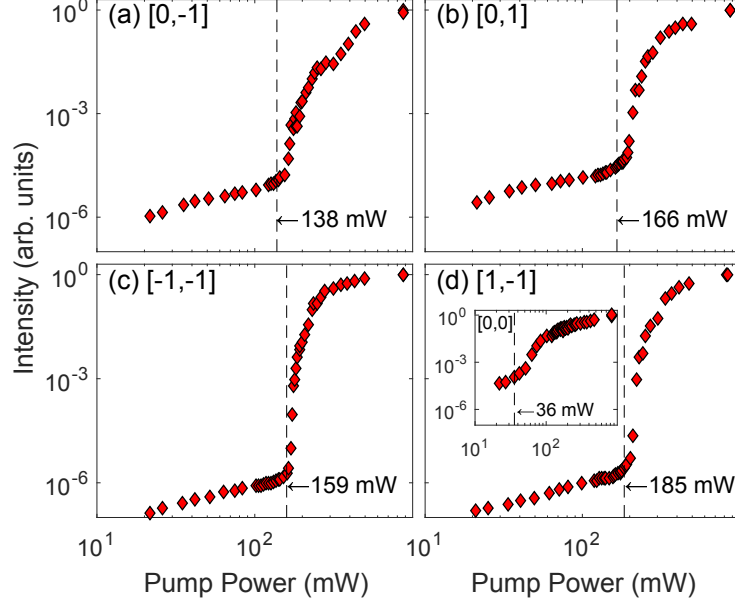


Figure 3.5: Log-log plot of the intensity of the polariton luminescence at the pillar edges as a function of the pump power at various positions. Intensity values were taken by integrating over a small area in spectrally resolved images at the lower polariton energy at the pillar edge in images similar to those in Figure 3.2. The vertical lines mark the threshold powers, which are defined by the onset of nonlinearity. Positions are defined according to Figure 3.6(b). The threshold powers and positions for each plot are (a)  $P_{\text{thres}} = 138 \text{ mW}$  at point  $[0, -1]$ , (b)  $P_{\text{thres}} = 166 \text{ mW}$  at point  $[0, 1]$ , (c)  $P_{\text{thres}} = 159 \text{ mW}$  at point  $[-1, -1]$ , (d)  $P_{\text{thres}} = 185 \text{ mW}$  at point  $[1, -1]$ , and (inset)  $P_{\text{thres}} = 36 \text{ mW}$  at point  $[0, 0]$ .

vs. power at various points along the pillar (Figure 3.5). The process of getting energy-resolved data for each point was similar to that used for mapping the potential landscape, discussed above. The positions of each of these points are defined using the same coordinate axes as Figure 3.1(c), but in units of the shortest distance from the origin to the edge trap ( $\approx 41 \mu\text{m}$ ) [see Figure 3.6(b)]. At each position, an approximately linear dependence on intensity was maintained up to the condensate threshold, where a strong nonlinearity emerged. The threshold power was generally lower along the edges [Figure 3.5(a-b)] than at the corners [Figure 3.5(c-d)], with the lowest threshold on the middle of the edge downhill along the gradient from the pump spot, consistent with the picture that the majority of polaritons flow downhill [see also Figure 3.6(a)]. The increase in emission from the edges was more than five orders of magnitude at the threshold in some cases. The emission eventually saturated with increasing pump power and returned to increasing linearly with pump power. This is due the limited number of polaritons available in the system at a given pump power, which puts an upper bound on the possible condensate density even with extremely efficient stimulated scattering into the condensate.

The emission from the pump spot region showed a much less sharp and lower threshold ( $\approx 36 \text{ mW}$ ) [Figure 3.5(d)(inset)]. As discussed in Section 3.6 and Ref. [78], we interpret this behavior at the pump spot as the onset of a quasi-condensate at densities just below true condensation; the momentum distribution is strongly altered, but there is no evidence of superfluidity. At higher density, a true condensate emerges in whatever local energy minimum is available.

A compilation of the threshold powers at various points along the edge is shown in Figure 3.6(a). The line of symmetry and overall tilt is oriented along the sample gradient. In general, higher energy points along the gradient had higher threshold powers, as expected, since polaritons from the pump region will generally run downhill and away from these regions. The corners are the obvious exception, since the two shown had the lowest energies of any points on the pillar, but noticeably higher threshold powers than the neighboring points on the edges.

The merging of the separated condensates into one, single condensate can be seen clearly in the data of Figure 3.7. In this figure we compare the energy vs. pump power at various

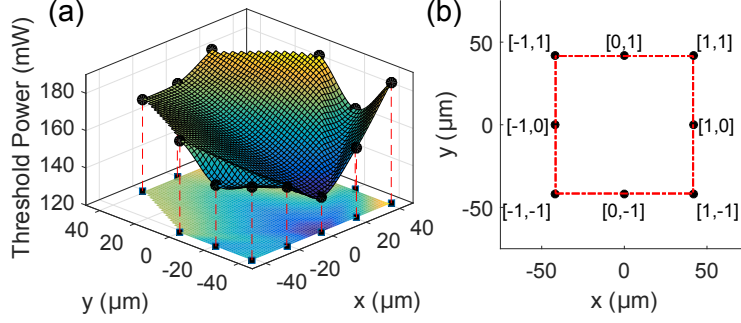


Figure 3.6: (a) Threshold powers (as defined in Figure 3.5) at various points on the pillar. The black circles give the actual threshold powers at each point. The surface is there to help guide the eye, and is a linear interpolation of the data points. The axes are the same as those in Figure 3.1(c). The two corners at  $y = 41 \mu\text{m}$  did not emit brightly enough at low pump powers to establish a threshold power. (b) Diagram defining a simple designation for each point, given as the actual position in Figure 3.1(c), but in units of the shortest distance from the center to the bottom of the edge trap ( $\approx 41 \mu\text{m}$ ).

separated locations on the pillar. The data were taken at two spatial locations simultaneously by imaging a full slice across the pillar. This removes deviations in the measurements due to pump laser pointing instabilities or small shifts of the pump location. Mostly, those would cause slight changes from measurement to measurement in the actual LP density, which makes it impossible to directly compare the precise energy values between two points in separate images. For each location, energy blue-shifting was seen above the threshold power, with the lower energy locations showing larger blue shifts due to higher polariton densities. In all cases, the lower energy locations eventually shifted to the same energy as the higher energy locations, and then their energies were locked together above that power to within our spectral uncertainty ( $\pm 0.05 \text{ meV}$ ). This same feature is apparent in Figure 3.2(c-d), which correspond to two pairs of data points in Figure 3.7(a). This indicates that the apparently interconnected polariton luminescence at powers well above threshold shown in Figure 3.3(d) was from a single mono-energetic polariton condensate.

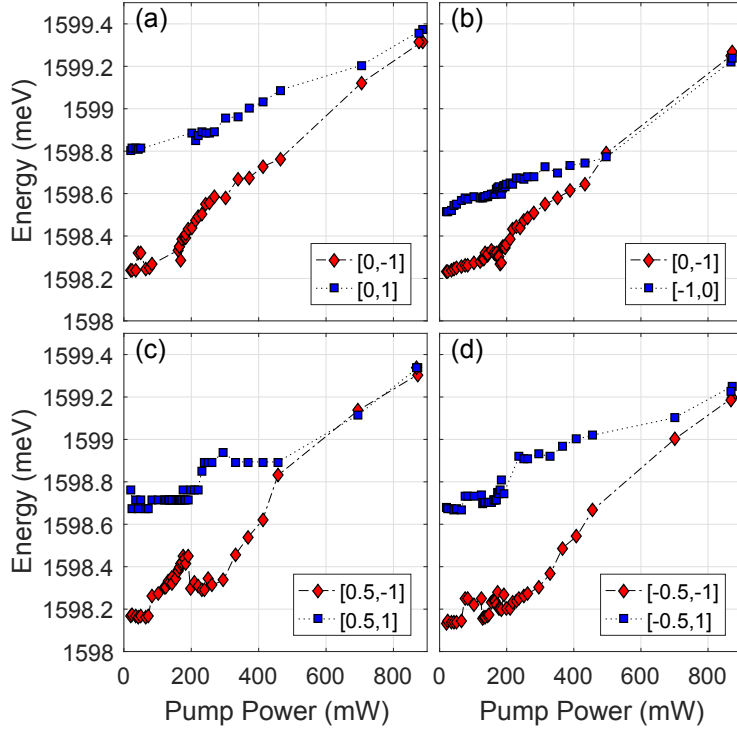


Figure 3.7: Energy vs. pump power at separate locations on the pillar. Data for each pair of locations were taken from a full energy resolved slice across the pillar, similar to and including those in Figure 3.2, allowing simultaneous measurement of both locations. The locations are defined according to Figure 3.6(b) and are (a)  $[0, -1]$ ,  $[0, 1]$ , (b)  $[0, -1]$ ,  $[-1, 0]$ , (c)  $[0.5, -1]$ ,  $[0.5, 1]$ , and (d)  $[-0.5, -1]$ ,  $[-0.5, 1]$ , shown as red diamonds and blue squares, respectively. The uncertainty is  $\pm 0.05$  meV.

### 3.4 EFFECT ON THE POLARITON LIFETIME

The lifetime of the polaritons was not affected by the etch process to a degree detectable in the energy linewidth using our equipment. As mentioned in Ref. [79], which used the same spectroscopy system and a similar sample, our spectral resolution (0.1 meV) gives a lower bound on the polariton lifetime that is much smaller than the measured lifetimes around 200 ps (see Ref. [25], and also Section 2.1). In fact, in a more recent experiment [59], the same etch process was shown in time-resolved measurements to have no measurable effect on the cavity lifetime. However, that experiment was not at equilibrium, and there were many dynamical process that affected the decay of the LP population over time. It could not be used to establish an accurate lifetime measurement, but it did confirm that there certainly was not a large increase in the cavity decay rate due to the etching.

That said, there certainly should be some reduction of the lifetime as indicated by the PL emitted from just outside the pillar edges in Figure 3.2(a) and Figure 3.3(a). That luminescence is similar in linewidth to the bare exciton ( $\sim 1$  meV) and centered close to the expected bare exciton energy. Since the pump was placed in the center of the pillar for those measurements, it is clear that the polaritons must be flowing to the edge and coupling into the bare or nearly-bare excitonic modes. This indicates that there is another decay channel for the polaritons, which should reduce their lifetime; the effect has just not been measurable on the scales used in our experiments up to this point.

### 3.5 STRAIN AS AN EXPLANATION OF THE TRAPPING

As discussed in Section 1.4.8, the polariton energy (via the exciton energy) can be changed by strain in the structure. Etching obviously changes the strain since it causes an abrupt end to the lattice at the etched edge. It is no surprise then that an energy shift appears in the polariton mode near the edges. The actual change in the exciton energy depends on the deformation, and abruptly terminating the lattice could, in general, either induce or relieve strain. More than one type of strain could also be involved, each with different healing

lengths. Such a picture can explain the observed result of a gradual decrease in polariton energy toward the edge with an abrupt energy increase closer to the edge, creating a trap set slightly away from the edge. If one strain type reduces the exciton energy and has a long healing length, and one increases the exciton energy and has a short healing length, the sum of the two would produce just such a result.

### 3.6 THREE-THRESHOLD BEHAVIOR

In this section, I will discuss the multiple thresholds that we observed in these pillars. This data comes from a different experiment than that shown above, and also from a different but similar device (the same as those of Chapter 4). However, this behavior seems to be general to the square pillars. The only major experimental difference is that the pumping was done at a duty cycle of about 2.4%, but that should not have a significant effect.

Figure 3.8 shows the condensate intensity in the traps at the corners and edges of the pillar as a function of optical pump power, for one of the devices. There are three thresholds for nonlinear increase of the intensity clearly visible here; in some devices only two of these thresholds are visible.

This threshold behavior can be understood as follows. Below the lowest threshold, there is a thermal distribution of polaritons concentrated at the pump spot, diffusing throughout the pillar [see Figure 3.8(b), 10 mW data]. The emission at the corners and edges increases linearly with the pump power. The first threshold occurs when condensation occurs at the optical excitation spot in the middle of the pillar. As discussed in previous work on these samples [78, 79], in our long-lifetime structures, as the excitation density is increased, a “quasicondensate” first appears at the excitation spot, characterized by increased occupancy of low energy states, but without long-range coherence or superfluidity. This is apparent in the fact that it does not flow away from the pump spot to a low-energy region. It is therefore also blue-shifted by the high exciton density at the pump spot. The emission of the quasicondensate is spectrally narrow but still has measurable spectral width, and is far more intense than all other regions [see Figure 3.8(b), 24 mW data]. Polaritons from this

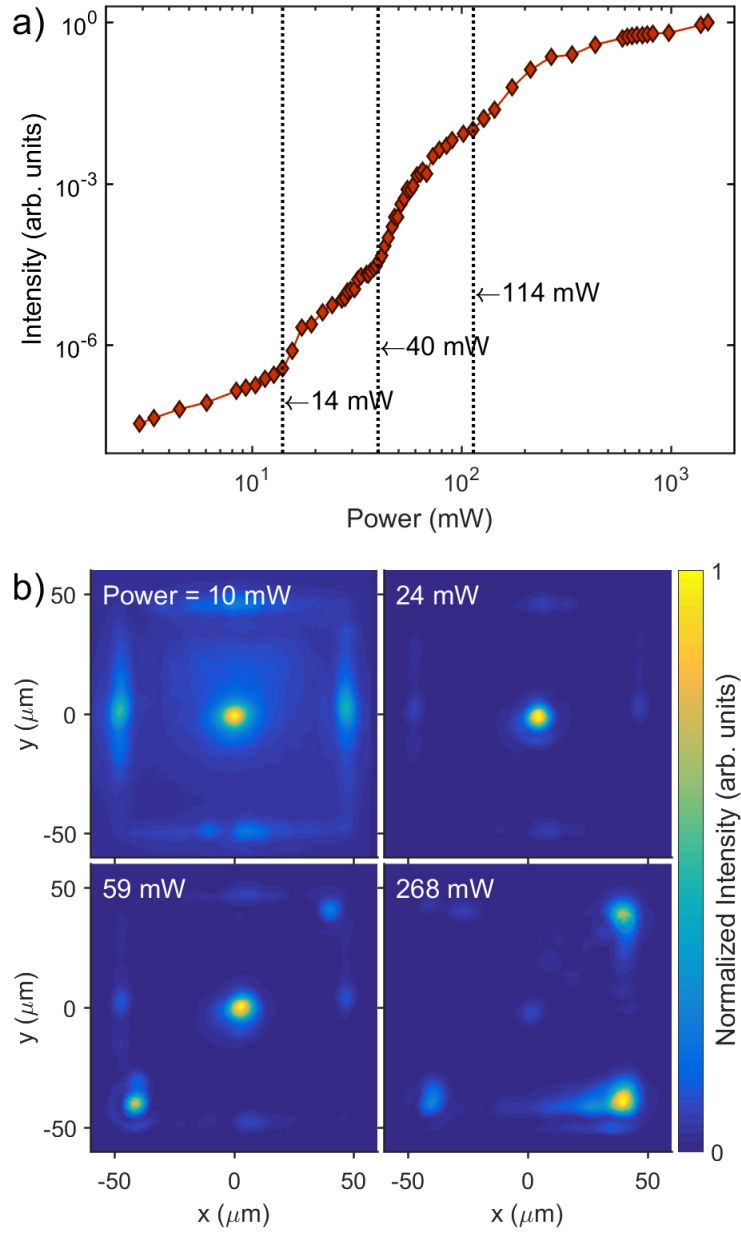


Figure 3.8: (a) Total intensity from the pillar corners as a function of pump power, showing three separate thresholds. (b) Real-space images showing the density distribution within the pillar at pump powers below, between, and above each of the thresholds. These data were from a different device from the one used for the data of Figures 3.2-3.7. This device had a detuning of  $\approx -9$  meV.

quasicondensate stream away ballistically, repelled by the high exciton density at the pump spot, with some reaching the low-energy traps in densities still too low to form a condensate.

A second threshold occurs when there is a “true” condensate that appears in the traps at the edges of the pillar, namely the corners and edges of the pillar, which are the lowest energy points as discussed above. This state is characterized by emission which is spectrally very sharp and which is localized to the ground states of the traps [see Figure 3.8(b), 59 mW data]. This behavior has also been seen in other trapping geometries [23, 78]. This drop in the photon emission energy, extreme spectral narrowing, and long-distance motion is quite dramatic in the experiments. This is possibly due to the onset of superfluid flow from the quasicondensate at the pump spot. The ballistically streaming polaritons may reach high enough densities and collision rates to scatter into the low energy states, forming condensates in the traps that then stimulate the scattering from the rest of the population at the pump spot. This initial onset due to collisions of the streaming polaritons could explain the minima in threshold power that appear in line with the cavity gradient and perpendicular to the etched edge shown in Figure 3.6. Ballistic polaritons forming a stream will be partially reflected from the edge, and only directly back-scattering at those points, allowing for efficient scattering to low energy. No matter the mechanism, the condensates formed at the second threshold in these traps are well behaved, with well-defined energy and a peak in momentum near zero.

The third threshold has been seen only in the large square pillars studied in this work, and occurs when the various trapped condensates in the corners of the pillar begin to spread out of their traps into the regions between the corner traps [see Figure 3.8(b), 268 mW data]. As discussed above, this eventually results in a single, global, monoenergetic condensate. Surprisingly, this third threshold also leads to a nonlinear increase of the total emission intensity, possible due to the now even more macroscopically occupied single state.

We examined nine devices with detunings ranging from  $\delta_0 = -10$  meV to around  $\delta_0 = -3$  meV. In several of the devices, the lower two thresholds are so close together that they appear as a single threshold, leading to only two clear thresholds, as in Figure 4.5(a). In the case of Figure 3.5, the lower two thresholds appear to have merged together, and the third one is not apparent. As can be seen by the relatively high threshold powers the coupling of

the pump into the structure was not as ideal as in the case of Figure 3.8. This meant that the region of density that typically showed the third threshold was possibly not reached, or the density of points was too small to show the effect.

### 3.7 CONCLUSIONS

What we found from this experiment was that the etching successfully produced confinement without significantly altering the lifetime of the LP mode. This allowed the formation of spatially extensive condensates in the low energy regions near the edges. In fact, the trapping of the polaritons near the edge, but with a sharp barrier directly at the edge, is one possible mechanism for the lack of significant coupling into bare exciton modes or defects along the etched edges. The first direct application that stands out is to use regions far from corners (which cause local traps) to produce shallow ( $\approx 1$  meV) 1D wells in relatively wide structures ( $> 20 \mu\text{m}$ ). The deeper trapping that emerges at corners could be implemented for shallow harmonic-like traps. However, the most promising feature for my purposes was the emergence of the energy minima  $\approx 7 \mu\text{m}$  away from the etched edge. This allows for even deeper trapping with edges  $\lesssim 15 \mu\text{m}$  apart without the need for complicated and costly fabrication on the  $\sim 1 \mu\text{m}$  scale. This is exactly what we did for the work discussed in Chapter 5, and in the more recent work reported in Ref. [59].

## 4.0 CURRENT INJECTION INTO A POLARITON CONDENSATE

### 4.1 MOTIVATION

The goal of this experiment was pretty simple: Form a polariton condensate, drive a DC current in-plane through the quantum wells in the same region as the condensate, and look for any observable effect. Since polaritons are charge neutral, it may be surprising that we should expect any significant effect other than upon the underlying exciton states. However, polaritons are partially made up of conduction band electrons and valence band holes, both of which are charge carriers in the uncoupled states. It follows that a free charge carrier should be able to interact with the exciton part of the polariton. If interactions are possible between a free charge carrier and a group of polaritons in a macroscopic state, then it is possible that macroscopic effects could be observed in the charge carriers.

To that end, we injected in-plane electrical current into a trapped polariton condensate and looked for nonlinear effects from each on the other. As I will discuss in this chapter, there were two overall effects: (1) An increase in the polariton condensate density with injected current and (2) a superlinear increase in conductivity in the presence of a condensate.

It should be noted that there has been previous work studying the effects of injecting current into a polariton condensate or polariton device. This includes work on the effect of a polariton condensate on photocurrent (e.g. Refs. [102,111,112]) and electrical pumping of polaritons (e.g. Ref. [27,28]) but almost all cases used a vertical injection arrangement. This typically involves oppositely doping both the top and bottom DBR layers and placing a metal contact directly on top. This does not allow any direct effect on the in-plane momentum, from either particle on the other. Ref. [113] did study lateral injection in a very similar geometry to ours. However, as will be discussed in Section 4.5, they reported a similar

measurement but did not observe the same effect that we did.

The work described in this chapter was previously published in Refs. [114, 115].

## 4.2 THE EXPERIMENT

The microcavity samples were of the same design as the other studies in this thesis, which is discussed in Section 2.1. The structure used in this experiment was similar to that of Chapter 3, but with a couple important additions. The first addition was a second etch, which produced a linear “mesa” region outside of which the quantum wells were removed. Centered on the mesa was an  $\approx 100 \mu\text{m} \times 100 \mu\text{m}$  pillar identical to those studied in Chapter 3, and on each end was placed a NiAuGe contact of the type discussed in Section 2.2.3, which was the second addition to the fabrication process. This resulted in the structure shown in Figure 4.1(a), a real example of which is shown in Figure 4.1(b). This allowed for formation of a trapped condensate within the central square pillar, while current could be driven in-plane through the quantum wells. As discussed in Chapter 3, low-energy regions are formed along the etched edges by strain, which trap the condensate near the edge. This is shown for one of the actual devices used in this study in Figure 4.1(d), and the ground state energy profile is given for reference in Figure 4.1(c). Thus, the behavior discussed for the case of the pillars alone, including the locking into a monoenergetic condensate at high powers and the three separate thresholds (quasicondensation, “true” local condensation, monoenergetic global condensation), also applied to these devices (see Chapter 3 for discussion of these effects).

While the experimental setup was generally the same as that discussed in Chapter 3, there was also the addition of the application of a voltage and measurement of current. The details of the electrical measurement, especially the synchronization with the chopped laser pulses, are given in Section 2.3.4. The pump energy was again nonresonant, but it was varied for each individual device to give optimal absorption at the lowest point in reflectance of the stop band. Since the stop band varies with the cavity gradient, this meant that the laser needed to be tuned differently for each detuning to hit the center of one of the dips in

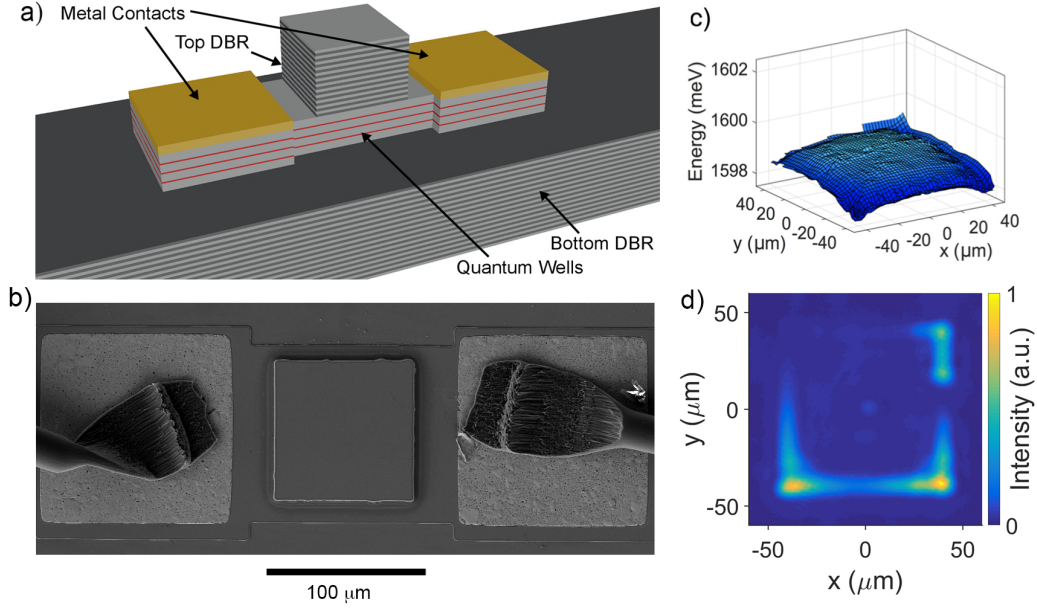


Figure 4.1: a) The structure of the polariton microcavity with symmetric  $n-i-n$  current injection. b) Scanning electron microscope (SEM) image of a typical structure. c) Typical trapping profile of the polaritons in the central pillar structure (see Section 3.3) d) Photoluminescence spatial profile of a condensate in the pillar at high pump intensity (462 mW peak power into a 20  $\mu\text{m}$  focal spot, chopped with a 2.4% duty cycle).

reflectance. This careful tuning also explains the relatively low threshold power compared to the majority of the work reported in Chapter 3 (with the exception of Section 3.6, which actually came from the same set of experiments as reported in this chapter). One other difference was the use of a 2.4% duty cycle for chopping the pump, corresponding to 60  $\mu\text{s}$  of pumping per pulse. This allowed for longer electrical measurements within the pulse window while still keeping the sample cool enough at high peak pump powers. The overall process was to set the pumping to a specific power, and then sweep the voltage through the various settings. At each voltage setting, a PL measurement was taken with the CCD and a current measurement was taken with the Keithley 2636B source-meter. Once these were collected at each voltage, the pump power was increased to the next setting and the process repeated. The result was a 2D matrix of data, corresponding to varying voltage and pump power.

### 4.3 BASIC ELECTRICAL CHARACTERISTICS: THE DEVICE AS A PHOTOTRANSISTOR

As discussed in Section 2.2.3, the NiAuGe contacts produce an effectively *n*-doped region under each contact after annealing. The result is a mostly ohmic contact into an *n*-type quantum well. However, the region between the contacts, where the condensate is formed, retains only its intrinsic doping. This gives the device an overall *n-i-n* structure. To get a better sense of the bands, we performed Hall measurements for all three regions. The results were that the contacts were *n*-type with area density  $6.7 \times 10^{15} \text{ cm}^{-2}$ , while the nominally undoped central region was also *n*-type, with area density  $1.8 \times 10^{10} \text{ cm}^{-2}$ . Since the doping concentration measurement was done on an unetched sample, the depth of the doping was estimated to be the entire top DBR and cavity thickness (about 4.2  $\mu\text{m}$ ). This gave doping concentrations of  $1.6 \times 10^{19} \text{ cm}^{-3}$  in the contact regions and  $4.3 \times 10^{13} \text{ cm}^{-3}$  in the central region. We used the “1D Poisson Beta 8j” program by Dr. Gregory Snider at Notre Dame University [116], which solves one-dimensional Poisson equations for semiconductor structures, to model the bands. We assumed ohmic contacts, no applied bias, a depletion length of 0.1  $\mu\text{m}$ , and that the dopant states are fully ionized (giving a free carrier density

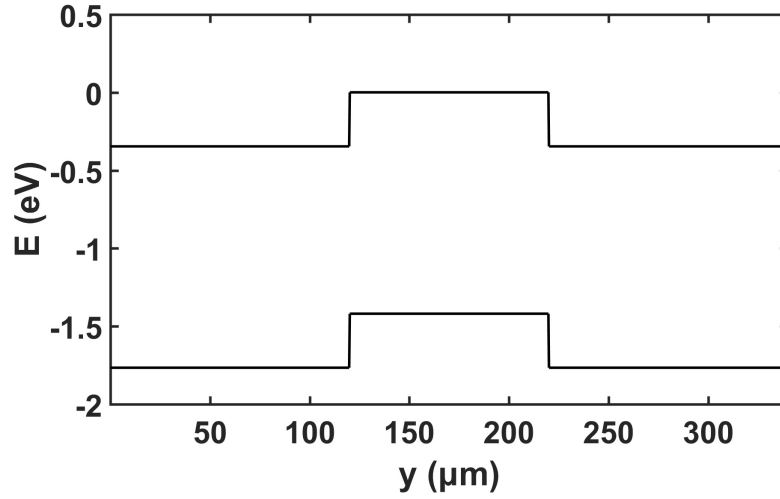


Figure 4.2: The device band diagram, produced using the “1D Poisson Beta 8j” program by Dr. Gregory Snider [116], as discussed in the text. [This figure was produced by Burcu Ozden.]

equal to the measured dopant density), and used  $100\ \mu\text{m}$  for the central region thickness and  $120\ \mu\text{m}$  for the more heavily doped regions. The result is given in Figure 4.2, showing a band offset of 0.34 eV.

The result is effectively a symmetric phototransistor with a very long base region separating the collector and emitter. This means that with no illumination to excite electrons into the conduction band in the central region, and at low temperature, essentially no current should flow through the device. Also, even with illumination, at zero or low voltage bias there should still be virtually no current. Because of the large space between the junctions, and with most of the illumination in the center, most electrons excited into the conduction band by photon absorption will simply recombine before reaching the junction.

To analyze this overall structure, we used the Ebers-Moll model for the phototransistor

[117]. This is given by the equations

$$\begin{aligned} I_B(G) &= \frac{I_s}{\beta} (e^{e(V_B-V_E)/k_B T} - 1) + \frac{I_s}{\beta} (e^{e(V_B-V_C)/k_B T} - 1) \\ I_C &= I_s (e^{e(V_B-V_E)/k_B T} - 1) \left(1 - \frac{1}{\beta}\right) - I_s (e^{e(V_B-V_C)/k_B T} - 1) \end{aligned} \quad (4.1)$$

where  $V_E$  is the emitter voltage,  $V_C$  is the collector voltage,  $V_B$  is the base voltage,  $\beta$  is the gain factor,  $I_C$  is the collector current,  $I_B$  is the base current, and  $I_s$  is the saturation current. The photogenerated carriers give a base current that depends on the voltage drop which pulls the electrons and holes away from each other and toward the contacts. We model this simply as  $I_B = G(1 - e^{-V_{CE}^2/\sigma^2})$ , where  $G$  is the pumping term and  $\sigma$  is a free parameter. Accounting for series resistance  $R$ , the measured current will be given by the solution

$$I_{\text{meas}} = \frac{V - V_C}{R} = I_C(V_C), \quad (4.2)$$

where  $V$  is the applied voltage.

Figure 4.3 shows the current-voltage data for Device 4 (which is the primary device studied for the results in this chapter) for several different pump powers. Also plotted are solutions of the Ebers-Moll model for the phototransistor given above. As seen in this figure, although the model is symmetric with bias direction, the measured current is not, presumably because the photocurrent is somewhat asymmetric, larger in the negative direction, for this particular case. One explanation of this asymmetry is that the pump spot was slightly closer to one side than the other, forcing the photogenerated carries to travel further in one direction than the other. With no optical pumping, the device has no measurable current flow until very high voltage.

We used the same model and parameters to see the expected dependence of current upon pump power, which is shown in Figure 4.4. As expected for a device of this type, this shows linear growth of the current with pump power at low power, then sublinear behavior and eventually saturation. The saturation is due to approaching the  $I = V/R$  limit, where  $R$  is the series resistance added in; it therefore corresponds to the resistance of the device itself becoming negligible.

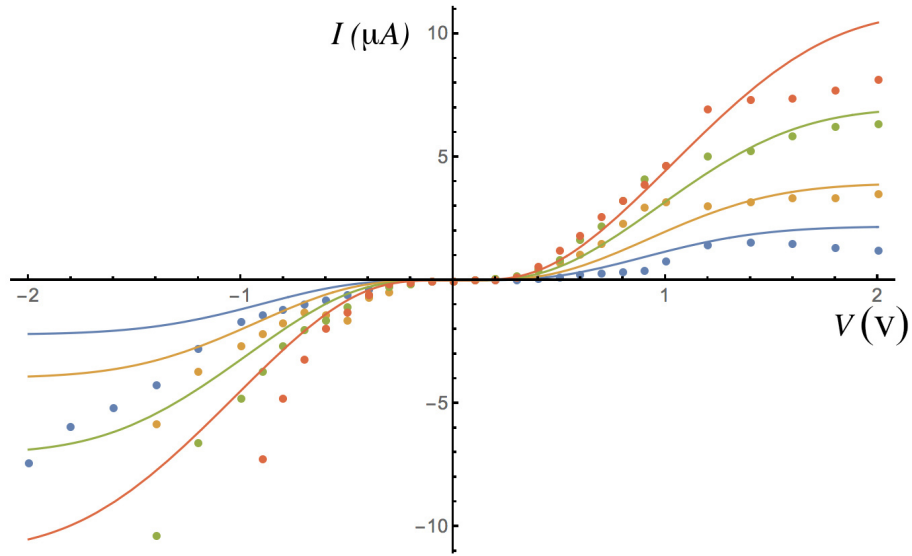


Figure 4.3: Dots: Measured current vs. voltage for Device 4 at various fixed optical pump powers: 5 mW (blue), 9 mW (yellow), 16 mW (green), and 25 mW (red). Lines: solutions of the Ebers-Moll model, for four illumination intensities with the same ratios as the four powers used for the experimental data. [This figure was produced by David Snoke.]

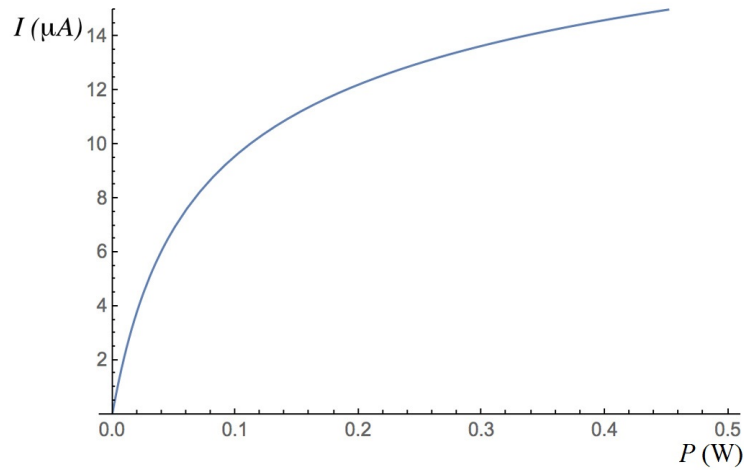


Figure 4.4: Current as a function of pump power for the same Ebers-Moll model and parameters for Device 4 used for Figure 4.3. [This figure was produced by David Snoke.]

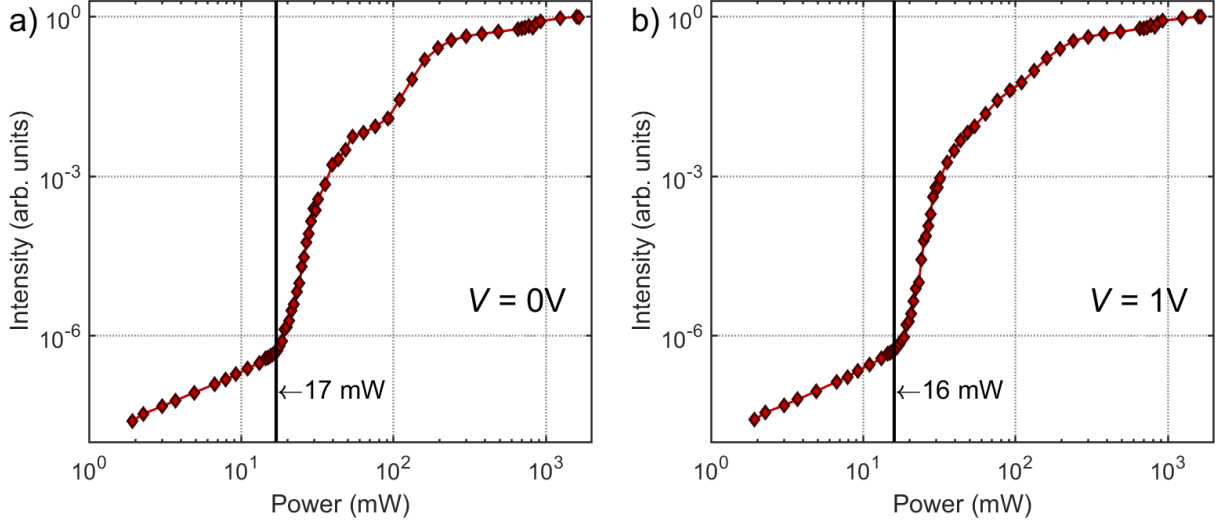


Figure 4.5: Photoluminescence intensity as a function of pump power at (a) 0 V and (b) 1 V.

#### 4.4 STIMULATED INJECTION INTO A CONDENSATE

The first effect that I will discuss is that of the current on the polariton condensate density. One of the most direct ways to see this is in a shift of the threshold power with applied voltage. This is shown in Figure 4.5. In these plots, the lower threshold is the actual condensate threshold, while the higher threshold is identified with the forming of a single condensate in the pillar (see Section 3.6). In panel (a), there is no applied voltage, so this is just the intrinsic curve for the pillar. In panel (b), which shows the same curve for an applied voltage of 1 V, an important feature is the slight shift of the curve toward lower powers. While this shift is small at lower powers, including a very small shift of the threshold, it becomes more significant in the nonlinear region of the data above the threshold. Obviously, something changed, but this particular comparison is not very revealing because of the similarities of the plots.

Instead, it is helpful to plot the data in the other direction, with the photoluminescence intensity shown as a function of current or voltage at fixed pump power. This is shown in

Figure 4.6. In panel (a), there is clearly a sharp increase in photoluminescence intensity (by a factor of about 20) at certain pump powers, while at higher and lower pump powers there is very little going on. This is especially interesting when comparing to Figure 4.5(a): Below the threshold, there is no effect, above the threshold but before the first saturation (near 100 mW) there is a strong effect, while above saturation there is again no effect.

It is very interesting that, in Figure 4.6(b), there are flat regions in the intensity at low voltage before the onset. This indicates that it is indeed an increase due to current and not somehow caused by the applied voltage by itself. Essentially nothing happens until a “turn on” voltage (notably very close to the barrier voltage shown in Figure 4.2) is reached, and then the photoluminescence increases sharply with current.

From this, it seems reasonable that the injected current is, at least partially, turning into polaritons and increasing the condensate density. If we imagine a plot of condensate density as a function of total polariton density, it should look something like Figure 4.5. The total density would simply be the addition of all the input terms, in this case optical pumping and current injection. This means that increasing the current would be equivalent to moving to the right along a curve much like the curves in Figure 4.5, and the effect of that increase would be greatest where the dependence is nonlinear. This is exactly the case shown in Figure 4.6, since the strongest effect is between the threshold and the saturation.

However, there is one problem. The electrical injection of polaritons requires the injection of both free electrons and free holes into the microcavity region (the area with the top DBR). Otherwise, the free carriers will not have partners with which to couple and form excitons. Since both contacts are *n*-type, only free electrons can be injected in significant numbers. There is possibly another explanation for the presence of free holes, even without a *p*-type region nearby.

There will always be some population of free holes in the pillar region due to ionization of excitons. Those free holes will tend to migrate to the negative side of the device, and there combine with free electrons to form new polaritons. The polaritons can then freely migrate across the pillar where they can again ionize and form a new free hole and an electron that carries current to the contact. This is depicted in Figure 4.7(a).

Since the polaritons in a condensate are in a macroscopically occupied state, the rate

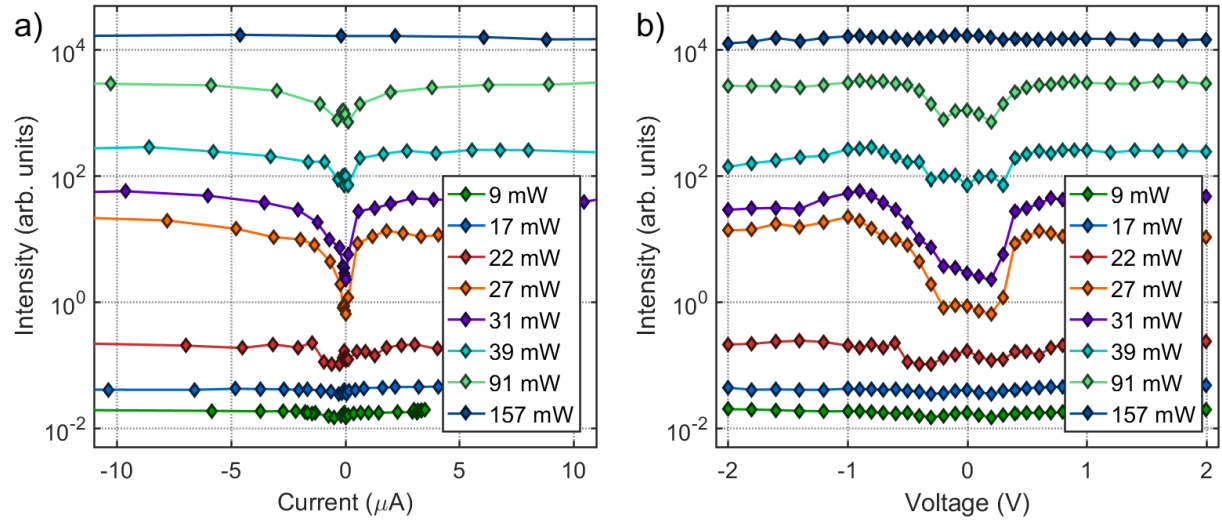


Figure 4.6: The effect of the applied voltage on the PL intensity. (a) PL intensity as a function of injected current, showing a sharp increase with current in the highly nonlinear regime just above the threshold. (b) The corresponding PL intensity vs. applied voltage plots, showing very little action until about 0.3 V is applied.

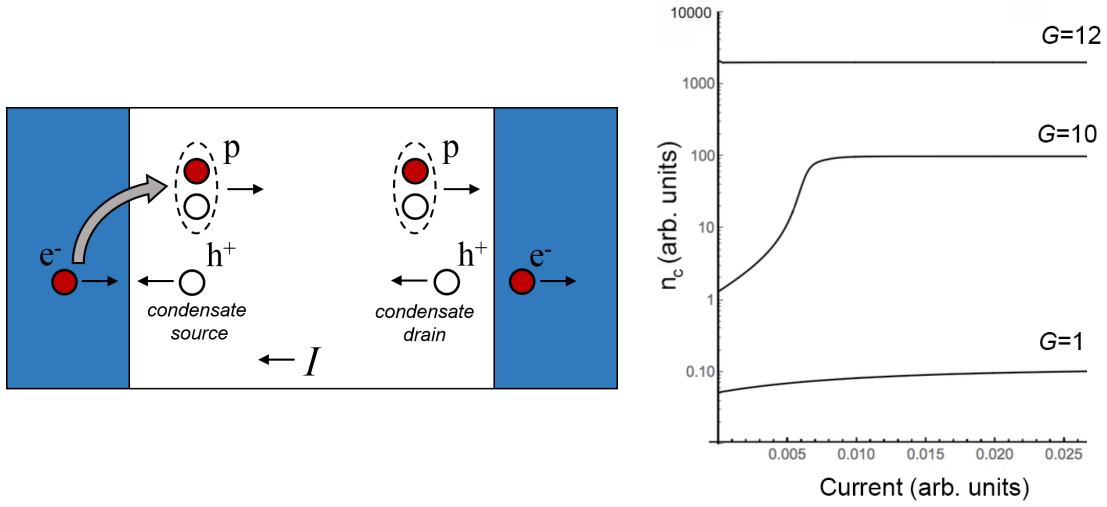


Figure 4.7: A simple model for the stimulated injection of free electrons into a polariton condensate. (a) A diagram of the process, showing electrons and holes combining into polaritons on one side, and polaritons ionizing into free electrons and holes on the other. The applied voltage will cause an overall flow of current in one direction. (b) The results of the model given by Equations 4.3, for  $A = 10^{-9}$ ,  $B = 0.5A$ ,  $\tau_r = \tau_{\text{ion}} = 1000$ ,  $\tau_{\text{ex}} = 100\tau_r$ .  $G = 10$  corresponds to the threshold for condensation. [Panel (b) was produced by David Snoke.]

of transfer into that state by any other process that can couple into it will be increased by a factor of  $(1 + n_c)$ , where  $n_c$  is the condensate population and is potentially much larger than unity (see Section 1.4.5 for a discussion of stimulated scattering and rate equations). This leads to a set of coupled rate equations, one each for the populations of the polariton condensate ( $n_c$ ), the exciton reservoir ( $n_{\text{ex}}$ ), the free electrons ( $n_e$ ), and the free holes ( $n_h$ ):

$$\begin{aligned}
\frac{\partial n_c}{\partial t} &= -\frac{n_c}{\tau_r} + An_{\text{ex}}(1 + n_c) + Bn_en_h(1 + n_c) - \frac{n_c}{\tau_{\text{ion}}} \\
\frac{\partial n_{\text{ex}}}{\partial t} &= G - An_{\text{ex}}(1 + n_c) - \frac{n_{\text{ex}}}{\tau_{\text{nr}}} \\
\frac{\partial n_e}{\partial t} &= J - Bn_en_h(1 + n_c) + \frac{n_c}{\tau_{\text{ion}}} - \frac{n_e}{\tau_{\text{esc}}} \\
\frac{\partial n_h}{\partial t} &= -Bn_en_h(1 + n_c) + \frac{n_c}{\tau_{\text{ion}}} - \frac{n_h}{\tau'_{\text{esc}}}.
\end{aligned} \tag{4.3}$$

There are really only a few processes going on here. The term with  $\tau_r$  corresponds to regular radiative decay out of the system. The  $A$  term gives stimulated scattering of the exciton reservoir into the polariton condensate, which is a fairly standard treatment (see e.g. Ref. [118]) that results in loss of excitons and gain of polaritons. The  $B$  terms are similar but correspond to the process outlined above: stimulated scattering of electron-hole pairs directly into the coupled polariton condensate. The  $\tau_{\text{ion}}$  corresponds to loss of polaritons through ionization of the excitons (which also goes as  $n_c$  without any fancy stimulation), and also feeds directly back into the free electron and hole states. The  $\tau_{\text{nr}}$  term gives the nonradiative loss rates of excitons, as introduced in Section 1.4.5. The  $\tau_{\text{esc}}$  and  $\tau'_{\text{esc}}$  terms correspond to the escape rate of the electrons and holes across the barriers, giving actual current flow. The rate of escape for the holes is taken to be much lower than for the more mobile electrons in this case by setting  $\tau'_{\text{esc}} \gg \tau_{\text{esc}}$ . The optical pumping is accounted for by  $G$ , while the electrical injection is accounted for by  $J$ , which, as discussed above, only adds electrons. This model does not directly take into account the spatial distributions of the particles or the size of the device, though these should only affect the particulars of the result and not the overall effect. Both electrons and holes can move through the device with a rate dependent on their mobilities and the applied voltage, which will contribute to the rates at which they escape ( $1/\tau_{\text{esc}}$  and  $1/\tau'_{\text{esc}}$ ). Additionally, superfluid flow of the condensate could allow even higher mobility for the electrons: electrons can be injected into the condensate at the high voltage

side, the polaritons can then migrate freely through the device, and they then ionize on the low voltage side where the electrons escape. No matter the mechanisms involved, the spatial details will simply affect the rates and thereby the steady-state densities for each population, but the general effect will still be preserved. Overall, the model is basically giving a new process by which additional particles can be snatched out of “nearby” states and added to the condensate, increasing the condensate density.

This set of equations can be solved for various choices of the pump rate  $G$  and current injection  $J$ , some of the results of which are shown in Figure 4.7(b). The general trend of the data is clearly reproduced, with a sharp dependence on current at the threshold of the pumping rate, followed by a saturation, while there is almost no response both well below threshold and above it.

An alternative explanation of the condensate density increase is an increased thermalization rate of the reservoir excitons due to interactions with the free electrons. This was predicted [62] and experimentally observed [119] in early work on exciton-polaritons, and suggested as a means of overcoming the early difficulties in producing a thermalized condensate. There is, however, one major difference between those works and this one: they were considering the interactions of the free electrons with the exciton reservoir directly at the location of the pump spot. This would correspond to looking at the photoluminescence intensity dependence of the polaritons at the pump spot on injected current. This is shown in Figure 4.8(b). The pump powers shown correspond to the same region with respect to the threshold power (now the threshold of the quasicondensate), as shown in 4.8(a). Clearly, there is no great effect of the injected current on the density of polaritons at the pillar center, at least on the scale of the effect on the thermalized condensate at the edges. Therefore, while the electrons may indeed be increasing thermalization, it is not a strong effect, and the stimulated capture of free electrons into an already-formed condensate is a better explanation, especially since it reproduces the general trend of the results.

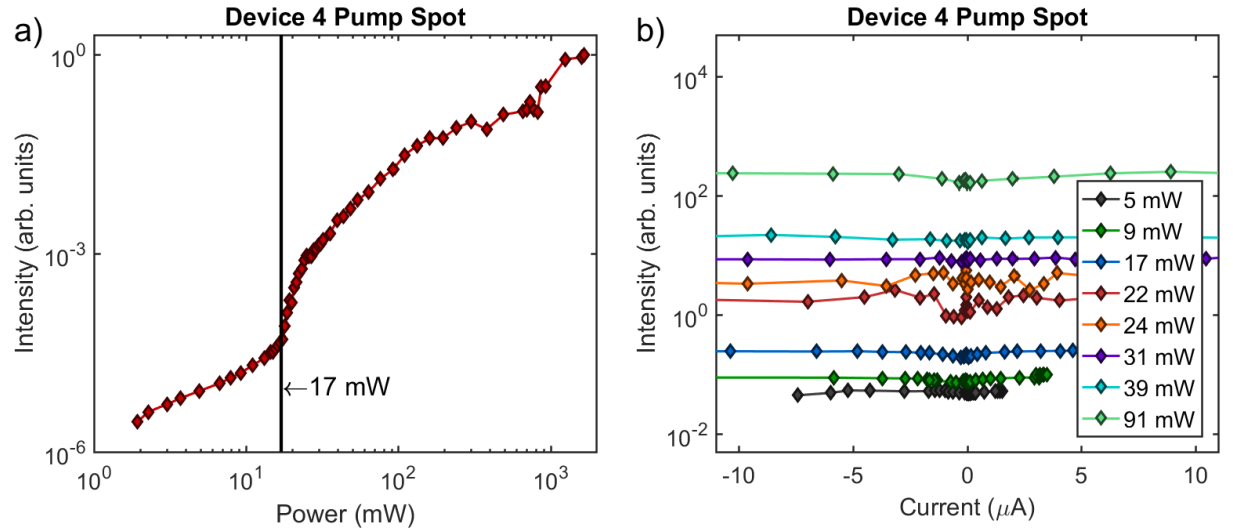


Figure 4.8: The photoluminescence intensity at the pump spot as a function of power and current. (a) Photoluminescence as a function of pump power with no applied voltage, showing the quasicondensate threshold at 17 mW. (b) The photoluminescence intensity as a function of injected current for various powers near and above the threshold, showing no significant effect.

## 4.5 SUPERLINEAR INCREASE OF PHOTOCURRENT

In Section 4.4, I focused mostly on the effect that the current has on the condensate. In this section, I will discuss the effect that the condensate has on the current.

As mentioned in Section 4.3, there is essentially no current through the device without excitation with a pump laser. This is also predicted by the Ebers-Moll model for a phototransistor, as shown in Figure 4.4. There is also an expectation from that model of a linear increase with pump power and eventually a slow saturation. The data shows something very different.

Figure 4.9(a) shows the dependence of both current and photoluminescence intensity on the pump power, at a set voltage of 1 V. At low power, there is indeed a linear increase in current with pump power, and even an obviously sublinear region in the case of Device 1. This trend looks very similar to the prediction of the Ebers-Moll model. However, as the polariton density increases sharply at and above the condensate threshold, the current also increases with it.

The clearly sublinear behavior in Device 1 is very similar to the observation in [113]. In that work, a very similar device was used, with contacts applied directly to exposed quantum wells. They also observed linear dependence followed by sublinear dependence above the quasicondensate threshold. However, they did not form a thermalized condensate, only achieving a quasicondensate at the pump spot. The sublinear region in our case also corresponds to powers above the quasicondensate threshold. It is not surprising, then, that the earlier work did not show a sharp increase in current, since they did not achieve the second threshold and form a “true” condensate. The trend in the data was clear, two examples of which are shown in Figure 4.9(a): The superlinear increase always occurred above the various thresholds for the formation of a thermalized condensate and never below it. The main reason for the various thresholds was variation in detuning. At more excitonic detunings, the lower polaritons have larger exciton fractions, and the threshold is therefore lower due to increased interactions giving faster thermalization. For example, Device 4 ( $\delta_0 \approx -8$  meV) was considerably closer to resonance than Device 1 ( $\delta_0 \approx -10$  meV), and clearly displays a lower threshold power.

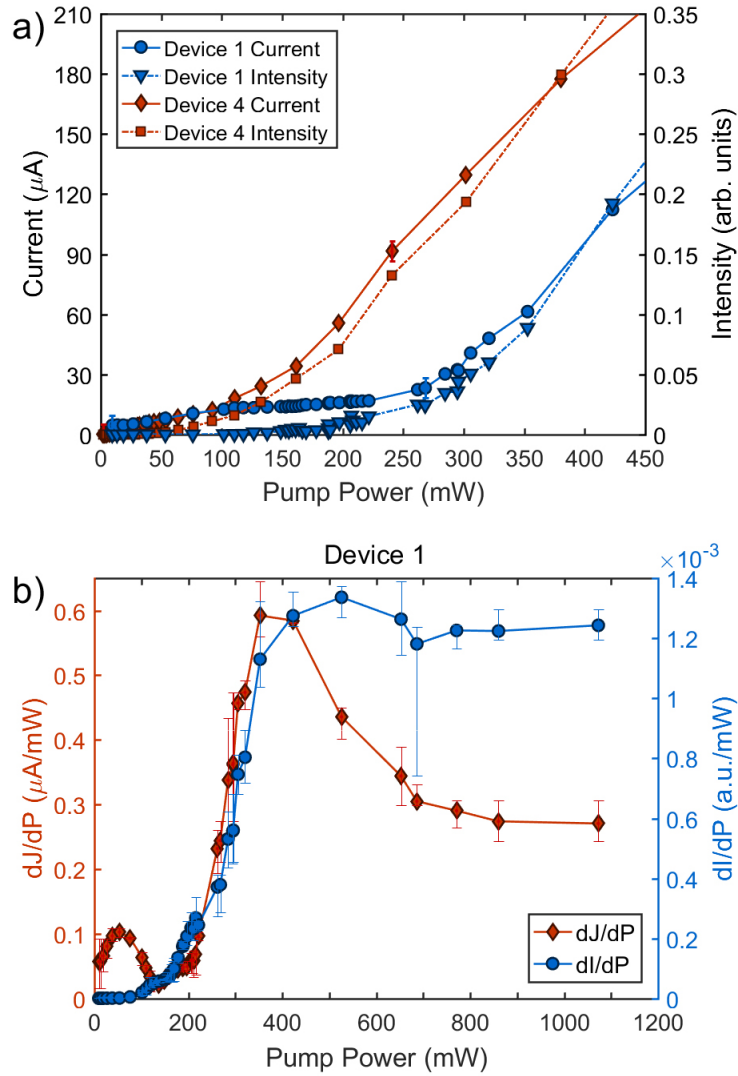


Figure 4.9: a) Solid lines: Current through the structure at constant voltage (1 V) as a function of the peak optical pump power, for two different devices. The two devices are identical except that they have different condensation threshold densities due to the different exciton fractions of the polaritons, controlled by the detuning (which is in turn dependent on the cavity width). Dashed lines: The emission intensity for the same devices under the same conditions. The excitation spot size for the two cases was also identical ( $20 \mu\text{m}$  diameter). b) The first derivatives of the data for Device 1 of (a). Red curve, left axis: the first derivative of the current; Blue curve, right axis: the first derivative of the photoluminescence intensity.

The first derivative of the data for Device 1 is shown in Figure 4.9(b). This was done by simply fitting the original data in small increments with a second order polynomial and then taking the derivative of that function. The error bars were established by widely varying the parameters of the algorithm (such as the range of points considered, the method of selecting that range, and the minimum number of points for each fit). The error bars therefore show the sensitivity of the result to the numerical process. Just as the increase in current and intensity are clearly correlated, so is the jump in their first derivatives. This essentially shows the condensate threshold, with the threshold in current almost exactly matched. The bump in the slope of the current at low power corresponds to the initial increase and then saturation expected for a normal photosensitive device without a condensate, but the jump in response to the polariton condensation is much more significant.

#### 4.5.1 Variations in the experimental geometry

In order to test the effect more fully, we performed similar experiments with variations in geometry. The first such test was a variation in spot size, which is summarized in Figure 4.10. For a 40  $\mu\text{m}$  spot size, the threshold power is increased due to the lower polariton density at the center. This can be explained, as discussed in Section 3.6, as due to the ballistically streaming particles from the condensate flowing to the edges and more efficiently thermalizing as they reflect back, encouraging the formation of local condensates along the edges at lower powers. No matter the mechanism, the result is an increased threshold power in the same device. For the photoluminescence thresholds, compare Figure 4.10(a), which shows the result for a larger spot size, with Figure 4.5(b), which is the results for the smaller spot size. The threshold is shifted to significantly higher power, which results in a high threshold for the onset of superlinear current increase, shown in Figure 4.10(b). Since this is the same device, and the only major difference is a change in threshold power, this further indicates that the current increase does indeed follow the onset of condensation.

The other variation was more significant, as we performed the same experiment in a device with a narrow wire geometry instead of the square pillar geometry. The overall structure and pump location are shown in Figure 4.11. These devices were designed with a

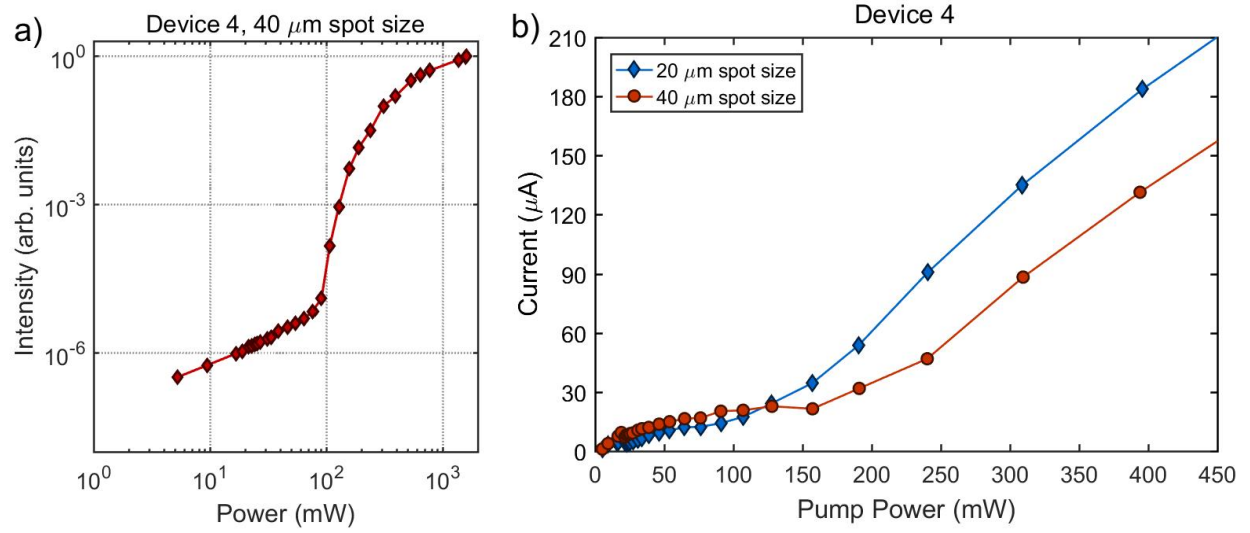


Figure 4.10: Pump spot size dependence of the current increase. (a) The PL intensity as a function of pump power, for a 40  $\mu\text{m}$  spot size and applied voltage of 1 V, showing a threshold around 90 mW. (b) Current as a function of pump power is plotted for both a 20  $\mu\text{m}$  and 40  $\mu\text{m}$  pump spot, for an applied voltage of 1 V.

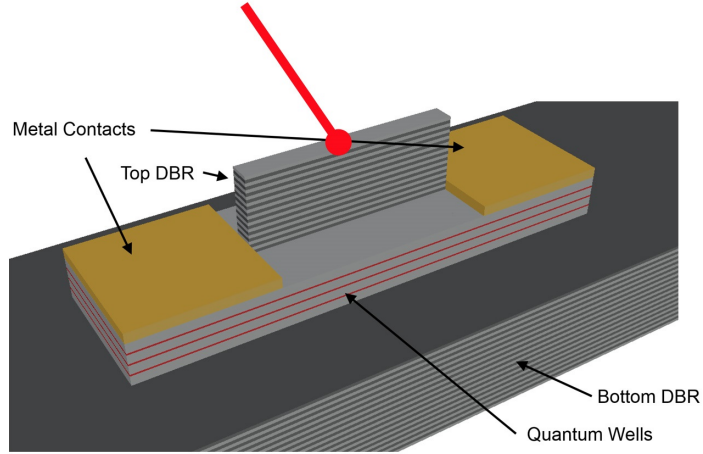


Figure 4.11: Illustration of the wire geometry, fabricated by the same method as the pillars presented in Figure 4.1. The total length of the wire is  $190\ \mu\text{m}$ , while the effective width is approximately  $10\ \mu\text{m}$ .

different experiment in mind, and are more fully discussed in Chapter 5. In this case, the etched size of the wire was about  $200\ \mu\text{m}$  long and  $17\ \mu\text{m}$  wide. Due to the edge effects, this resulted in effective dimensions of about  $190\ \mu\text{m}$  long and  $10\ \mu\text{m}$  wide. The wire was pumped with a similar arrangement, using an approximately  $20\ \mu\text{m}$  spot size.

The results for the wire device are shown in Figure 4.12, with the same data plotted in both panels, but on a linear scale in the case of (a) and a log-log scale in the case of (b). The same trends appear, with the superlinear increase in the current following the threshold in polariton photoluminescence. However, in this case the current does not follow the condensate density as closely, with a threshold offset by about a factor of 2.

This is possibly due to the narrowness of the wire, resulting in low overlap between the exciton reservoir at the pump spot and the condensate in the end regions of the wire. This overlap is crucial for the superlinear current increase, as discussed in Section 4.5.2.

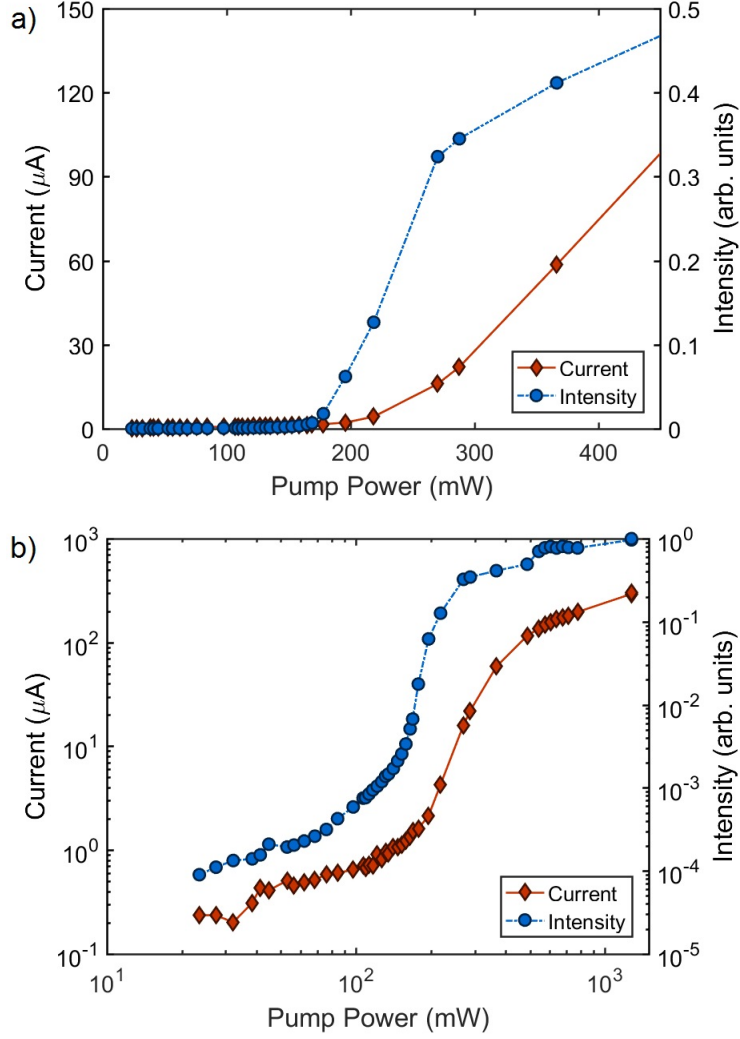


Figure 4.12: Current and photoluminescence intensity as functions of pump power for a narrow wire device. a) Linear, and b) logarithmic plots for a total voltage drop of 2 V, and a pump spot of approximately  $20 \mu\text{m}$  in the center of the wire. The detuning of the device was approximately  $\delta = -5 \text{ meV}$ .

### 4.5.2 Theoretical model

This section gives the theoretical model used to explain the results described above. It is primarily the work of Evgeny Sedov and Alexey Kavokin, who collaborated with us on this project, and is necessary for understanding the results. I will summarize their work here.

The standard approach for modeling polariton dynamics is to use a Gross-Pitaevskii equation for the light-mass polariton condensate and a rate equation for the heavy, incoherent exciton reservoir (e.g. Refs. [118, 120, 121]). Along these lines, we used the following set of equations. They describe the behavior of the coupled three-component system of polaritons, a free exciton “reservoir,” and free carriers, similar to the formalism used in Ref. [122]:

$$i\hbar \frac{\partial \Psi}{\partial t} = \left[ -\frac{\hbar^2}{2m^*} \nabla^2 + V_{\text{eff}}(t, \mathbf{r}) \right] \Psi + \frac{i\hbar}{2} (Rn_{\text{R}} + An_{\text{R}}^2 - \gamma_{\text{C}}) \Psi, \quad (4.4)$$

$$\frac{\partial n_{\text{R}}}{\partial t} = P(\mathbf{r}) - Rn_{\text{C}}n_{\text{R}} - 2An_{\text{R}}^2n_{\text{C}} - \gamma_{\text{R}}n_{\text{R}}, \quad (4.5)$$

$$\frac{\partial n_{\text{eh}}}{\partial t} = An_{\text{R}}^2n_{\text{C}} - \gamma_{\text{eh}}n_{\text{eh}}. \quad (4.6)$$

In these equations, the subscripts “C”, “R,” and “eh” denote the polariton condensate, exciton reservoir, and electron-hole reservoir, respectively. The  $n$  terms give the population densities, while the  $\gamma$  terms give the basic loss rates (as discussed for basic rate equations in Section 1.4.5). They include the normal loss through the mirrors in the case of the polaritons, the nonradiative loss rate of excitons, and the escape rate of the free carriers (as mentioned in the context of Equation 4.3).

Equation (4.4) is the generalized Gross-Pitaevskii equation (GPE) for the multi-particle wave function  $\Psi(t, \mathbf{r})$  of the polariton condensate. The polariton effective mass is given by  $m^*$ , and  $V_{\text{eff}}(t, \mathbf{r}) = V(\mathbf{r}) + \alpha_{\text{C}}n_{\text{C}} + \alpha_{\text{R}}n_{\text{R}}$  is the effective potential, where  $V(\mathbf{r})$  is the stationary confinement potential across the sample including the effects of strain, cavity gradient, etc. The  $\alpha_{\text{C}}n_{\text{C}}$  term is the standard addition in a GPE that differentiates it from the normal Schrödinger equation, adding in interactions between the particles in the condensate with interaction strength  $\alpha_{\text{C}}$ . The  $\alpha_{\text{R}}n_{\text{R}}$  term is similar, adding in interactions with the reservoir excitons. The imaginary terms in this equation give loss and gain since they will cause exponential terms in a wave solution. Thus, the radiative loss appears as a negative

imaginary number, while the stimulated scattering of the excitons into the condensate (with amplitude  $R$ ) is positive.

Equation (4.5) is the rate equation for the exciton reservoir density  $n_R(t, \mathbf{r})$ . Excitons in the reservoir have higher energy than the polaritons and do not emit light since they are not coupled to the external photon modes. The optical pumping is assumed to input only into the exciton reservoir, and is given by  $P(\mathbf{r})$ , including its spatial profile. This is valid due to the fast relaxation time of the excited carriers into bound exciton states; otherwise there would also be direct pumping of the electron-hole population. There are also the expected stimulated scattering term and decay terms. This equation is also coupled to Equation 4.4 by the relation  $n_C = |\Psi|^2$ .

Equation (4.6) is the rate equation for a reservoir of free carriers (electrons and holes), with density  $n_{eh}(t, \mathbf{r})$ . This equation also has the normal loss term  $\gamma_{eh}n_{eh}$ .

The important new term, which I have saved for last, is the  $An_R^2n_C$  term. It appears in all three equations, as a gain term for the electron-hole reservoir and the condensate, and a double loss term from the exciton reservoir. It therefore corresponds to two reservoir excitons scattering off of each other, with one gaining energy and ionizing into an electron hole pair while one loses energy and joins the polariton condensate. It is therefore analogous to the Auger effect in atoms. This is an energy-conserving process, that, in atoms, corresponds to the relaxation of an electron into a lower energy band by giving its excess energy to another electron rather than the more familiar process of emitting a photon. The unique aspect of this situation is that the process is stimulated by the condensate density, making it more important above the threshold.

There are two important differences between this new Auger-like stimulated scattering and the regular stimulated scattering. The first is that it is also dependent on the reservoir density. It should therefore be most important when there is a region with *both* large condensate density and exciton reservoir density. The second is that it results in increased free carrier density. This leads to scattering into the free carrier state that is stimulated by the condensate density.

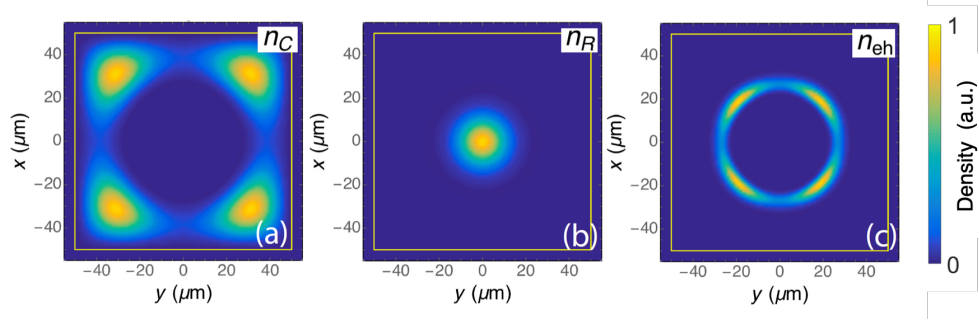


Figure 4.13: Example of the numerical results of Equations 4.7–4.9, giving the various densities in steady state. (a) The polariton condensate density, (b) the hot reservoir exciton density, and (c) electron-hole reservoir density. The yellow rectangle marks boundaries of the trap. [This figure was produced by Evgeny Sedov and Alexey Kavokin.]

The steady state case for these Equations 4.4–4.6 gives

$$\mu\Psi = \left[ -\frac{\hbar^2}{2m^*}\nabla^2 + V_{\text{eff}}(\mathbf{r}) \right] \Psi + \frac{i\hbar}{2} (Rn_R + An_R^2 - \gamma_C) \Psi, \quad (4.7)$$

$$P(\mathbf{r}) = Rn_Cn_R + 2An_R^2n_C + \gamma_Rn_R, \quad (4.8)$$

$$n_{\text{eh}} = \frac{An_R^2n_C}{\gamma_{\text{eh}}}, \quad (4.9)$$

where the time-dependent GPE in Equation 4.4 has been replaced by the time-independent GPE in Equation 4.7, which now includes a term with the chemical potential  $\mu$ . We solved this set of equations numerically using the following parameter values:  $m^* = 5 \times 10^{-5}m_e$ , where  $m_e$  is the free electron mass,  $\gamma_C = 0.004 \text{ ps}^{-1}$ ,  $\gamma_R = 0.04 \text{ ps}^{-1}$ ,  $\gamma_{\text{eh}} = 0.6 \text{ ns}^{-1}$ ,  $R = 0.025 \text{ ps}^{-1} \cdot \mu\text{m}^2$ ,  $A = 10^{-15} \text{ ps}^{-1} \cdot \mu\text{m}^4$ ,  $\alpha_C = 0.6 \mu\text{eV} \cdot \mu\text{m}^2$ ,  $\alpha_R = 2\alpha_C$ . For the pump profile, we used the Gaussian  $P(\mathbf{r}) \propto \exp[-r^2/w_p^2]$  with the width  $w_p = 10 \mu\text{m}$ .

Figure 4.13 shows the densities of the various populations for the parameters above in the steady state regime corresponding to the solution of Equations 4.7–4.9. The potential for a square pillar was included in the potential  $V(\mathbf{r})$  to simulate our actual devices. The result is close overall to the experiment, with the polariton condensate forming near the edge and especially in the corners of the pillar, with none in the center due to the exciton-polariton

repulsion. The exciton reservoir remains essentially in place at the center, which is expected since there is no diffusion term in the equations for the reservoir. As mentioned above, we expect to see the Auger-like effect most clearly at points of overlap between the excitons and the condensate, which is exactly what is shown in Figure 4.13(c)

By integrating over the distributions of the particles, the total populations of each can be found as a function of pump power. This is shown in Figure 4.14, where  $I_C = \int n_C d\mathbf{r}$  and  $I_{eh} = \int n_{eh} d\mathbf{r}$  in the steady state regime as functions of the optical pump power. This shows the same overall trend as the data in Figure 4.9(a), where the condensate density is directly proportional to the photoluminescence intensity (see Section 6.3). Likewise, the total current should be proportional to the number of free carriers. Thus, the tracking in density between the charge carriers and the polariton condensate, along with a slight lead by the condensate and the superlinear increase, are well produced by the stimulated Auger-like exciton-exciton scattering term in this model.

These can be compared to the data shown in Figure 4.9(a), where the total emission intensity from the condensate is proportional to  $I_C$ , and the total photocurrent is proportional to the number of free carriers  $I_{eh}$ . As seen in this figure, the main results of the experiments are reproduced, namely, the polariton condensate population and the photocurrent track with each other, and both show a superlinear increase due to the stimulated collisional term proportional to  $A$  in Equations (4.7)–(4.9).

As a final point, the dependence on the overlap of the condensate and the exciton reservoir can also explain the less dramatic increase at the condensate threshold in the narrow wires. Because the condensate and exciton reservoir only overlap at two small regions on both sides the overall effect is somewhat suppressed.

## 4.6 CONCLUSIONS

The experiments discussed in this chapter revealed processes that increase the density of both charge carriers and polariton condensates in response to each other. These are essentially injection processes that change the number of particles in one population or the other. In

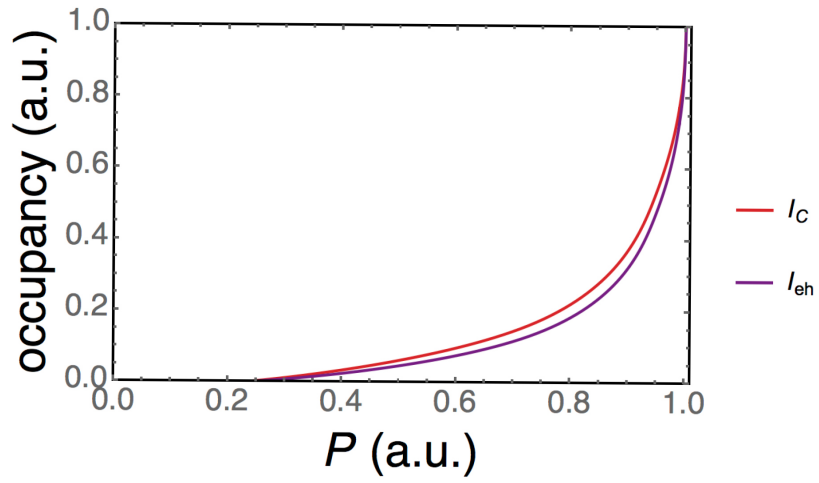


Figure 4.14: The normalized total occupancy of the polariton condensate and the free carrier reservoir, as functions of pump power, from the numerical results of Equations 4.7–4.9. The normalized polariton condensate occupancy,  $I_C/\max(I_C)$  is shown in red and the free carrier occupancy,  $I_{eh}/\max(I_{eh})$ , is shown in blue. [This figure was produced by Evgeny Sedov and Alexey Kavokin.]

order to better understand the interactions between the two sets, it would be interesting to see how other aspects of the states can change as a response. That is precisely the goal of the experiments described in Chapter 5, which look for a change in the momentum state of the polariton condensate due to the injected current.

## 5.0 MOMENTUM DRAG ON A POLARITON CONDENSATE

### 5.1 MOTIVATION

In Chapter 4, I presented experimental work that suggested interactions between polaritons and charge carriers. Specifically, we observed both a sharp increase in current that followed the increase in polariton density, as well as an increase in polariton density with a current input. The former effect was explained as Auger-like scattering of two excitons into both a polariton and a free electron and hole, while the latter was explained as an increased capture rate of intrinsic free holes by the injected free electrons. Both processes, which are normally weak compared to the other rate terms, become more significant because of the stimulation by a polariton condensate. However, the effects observed in those studies were essentially exchanges of particles between various populations. We were interested in directly observing the change in state of one particular population due to interactions with another, which led to the experiment described in this chapter.

This experiment attempted to show interactions by measuring the momentum state of a polariton condensate while injecting an in-plane current of free electrons. A change in the momentum as a function of the electrical current would possibly indicate a direct observation of electron-polariton scattering. As discussed in Section 1.4.7, such scattering should be possible due to either simple Coulombic interactions between the exciton part of the polariton and free electrons, or by the polaron effect that couples an exciton with a Fermi sea. Such interactions can be depicted by Feynman diagrams similar to the one in Figure 5.1. In vacuum, this would correspond to a photon indirectly interacting with a free electron by first exciting an electron-positron pair, though this particular diagram will be exactly canceled by a similar one with the fermion loop direction switched. Higher order

diagrams with even numbers of internal photon lines do make non-canceling contributions, giving interactions between photons and electrons in vacuum. In a semiconductor, the picture is quite different. The conduction band electron and valence band hole states now have very different energy dispersions. Thus, reversing the direction of the fermion loop no longer results in an equal and opposite contribution, so this particular diagram does contribute to the interaction. In addition to allowing lower-order interaction terms, semiconductors have another advantage in requiring much less energy to excite the pair of particles ( $\sim 1$  eV in the case of a semiconductor as opposed to  $\sim 1$  MeV for vacuum). The lower rest energy of the pair state leads to a larger matrix element for this interaction, which goes as the difference between the photon and the gap energy, leading to much larger probability for this type of interaction in a semiconductor (see chapter 8 of Ref. [38]). Observation of this type of interaction would be very interesting in its own right, but it may also have an application in using a DC electrical current to effectively steer the direction of light. Since a polariton is basically a photon that is constantly undergoing a transition to the electron-hole state, the upper part of the diagram essentially depicts a polariton. Thus, the effect should be observable in the momentum states of the polaritons.

In the case of our system, we also have the advantage of being able to form a condensate of polaritons. This leads to a very spectrally narrow emission as well as a sharp peak in momentum. Additionally, the long-distance transport and long lifetime of polaritons in our samples, as discussed in Section 2.1, allows us to form large condensates on the scale of  $\sim 100 \mu\text{m}$ , which allows good momentum resolution. Therefore, a shift in the momentum of the polaritons should be easier to observe since it will correspond to a shift of a sharp peak, rather than a broad distribution in momentum space. The combination of these factors makes polaritons an excellent candidate as a system in which to observe the effect discussed above. Specifically, it should appear as a drag by the electrons on the polaritons, pulling them in the direction of the electron flow. The converse should also be true, with polaritons dragging electrons in the direction of the polariton flow.

The work described in this chapter was previously made available at Ref. [123].

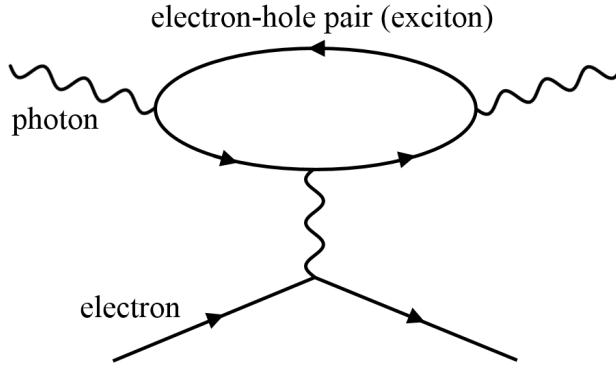


Figure 5.1: The Feynman diagram for electron-hole-pair-mediated photon-electron interaction, where the dashed lines are photons and the solid lines are electrons.

## 5.2 EXPERIMENT

The microcavity samples used in this study are essentially the same as those described in Section 2.1 and also used in the other work described in this thesis. We also used long wires,  $\approx 200 \mu\text{m} \times 20 \mu\text{m}$ , which were formed by etching away the top DBR, confining the polaritons within the wire and exposing the quantum wells. NiAuGe contacts were then deposited onto the quantum wells at each end of the wire. As discussed in Section 4.3, Hall measurements showed that the wire regions were intrinsically  $n$ -doped on the order of  $10^{13} \text{ cm}^{-3}$ , while the contacts were  $n$ -type with heavy doping on the order of  $10^{19} \text{ cm}^{-3}$ . The fabrication methods used to produce these samples are given in Section 2.2. These same wire structures were also used in the work described in Section 4.5.1, and are shown in Figure 5.2. The most important feature is that the two contacts were placed such that electrical current could be injected in the plane of the quantum wells. This allows direct transfer of in-plane momentum between the electrons and the polaritons. Since the in-plane momentum of the polaritons is easy to observe using  $k$ -space imaging (see Section 2.3.3), this arrangement allows direct observation of any effect of the electron flow on the momentum state of the polaritons.

The experimental details are similar to those described in Chapter 4. A continuous-wave

(CW) stabilized M Squared Ti:sapphire pump laser was used to excite a spot on the wire, with a spot size of about  $20\text{ }\mu\text{m}$  FWHM. The pump was again nonresonant, with an excess energy of about 100 meV, and mechanically chopped at 400 Hz with a pulse width of about  $60\text{ }\mu\text{s}$ . A source meter was used to sweep the applied voltage along the wires while measuring electrical current, in the arrangement described in Section 2.3.4. The wire devices used in this study were all near resonance between the cavity and exciton modes ( $\delta = E_{\text{cav}} - E_{\text{exc}} = 0$ ). One major difference in the data collection was that angle-resolved images were collected at a variety of applied voltages and pump powers, rather than just real space images, allowing us to observe the in-plane momentum of the polaritons at each setting.

As discussed in Section 3.6, there are multiple polariton density thresholds associated with these devices. In this chapter, the primary threshold of interest is the second one, which is for a “true” condensate in a local minimum of the device away from the pump location. I will refer to this threshold as  $P_{\text{thres}}$  throughout this chapter. The energy of the polariton condensate is determined by the external potential profile that it feels. This profile is determined by both the device and the local density of both excitons and polaritons. This was discussed in the case of square polariton pillars in Section 3.3, but the general picture holds here. At very low density, the profile is essentially only affected by the cavity gradient (as discussed in Section 2.1) and the etched shape of the device. Strain causes the formation of traps near etched edges, but in this case the width of the devices was similar to the range of the strain effects. This resulted in a single narrow trap running the length of the wires, with deeper traps at the ends due to increased strain, and an overall tilt due to the cavity gradient. Pumping produced a build-up of high- $k_{\parallel}$  excitons at the pump spot. The repulsive interactions between these excitons and the polaritons added a “hill” in the energy profile at the location of the pump. The overall picture is sketched out in Figure 5.3.

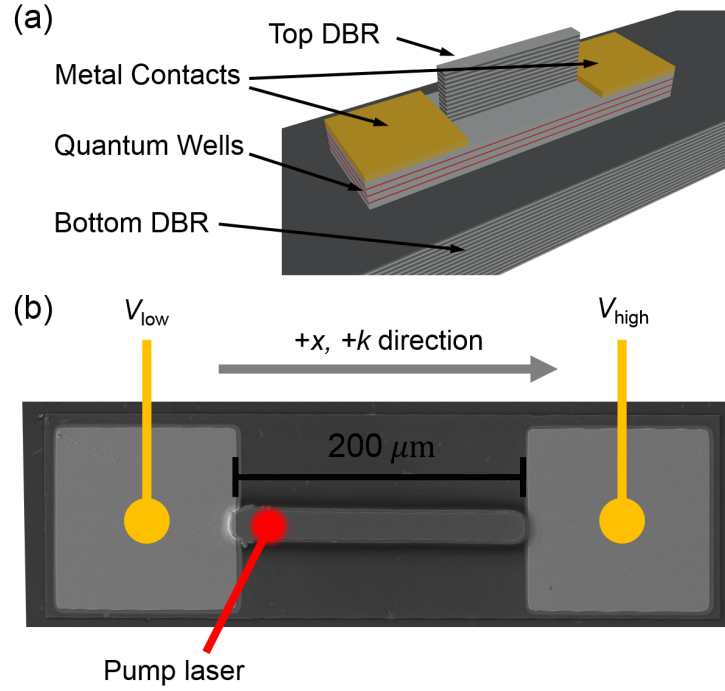


Figure 5.2: (a) The etched microcavity structure with the NiAuGe contacts at each end of the wire. (b) Scanning electron microscope (SEM) image of a representative etched wire, with an overlay showing the experimental arrangement. The voltage connections of the source were connected to the device as shown. Thus, positive conventional current flowed in the direction of  $-x$ .

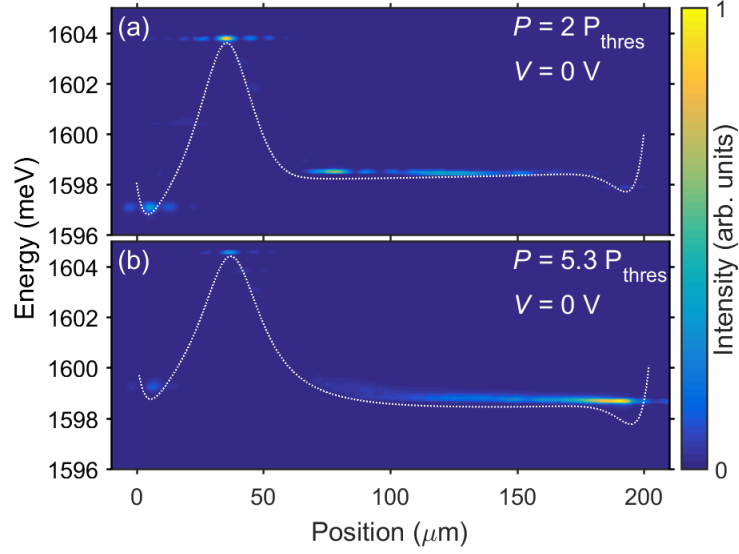


Figure 5.3: Photoluminescence intensity vs. energy and position along the length of the wire at zero applied voltage. The white dotted lines give the outline of the potential felt by the polaritons. The threshold power ( $P_{\text{thres}}$ ) was about 114 mW, and the powers used were (a)  $2P_{\text{thres}}$  and (b)  $5.3P_{\text{thres}}$ . The photoluminescence intensity was normalized separately for each image.

## 5.3 RESULTS

### 5.3.1 Drag on the polariton momentum

The experimental arrangement is shown in Figure 5.2(b), which gives the definitions of the  $+x$  and  $+k$  directions and shows the voltage connections. Unless otherwise specified, the pump spot was located near the left ( $x = 0$ ) end of the wire for all of the data shown in this chapter. Figure 5.3 shows the polariton distribution along the wire at two different pump powers above threshold and at zero applied voltage. The shallow cavity gradient along the  $+x$ -direction caused a small slope in the potential, with a total energy drop of about 0.5 meV along the wire. At the lower power, a quasi-condensate (see Section 3.6) at the location of the pump is clearly visible at high energy, as well as a significant population in the left-side end trap (near  $x = 10 \mu\text{m}$ ) and a smaller population in the right-side trap (near  $x = 190 \mu\text{m}$ ). In the level region on the side away from the pump, a large mono-energetic condensate forms, filling a region of the wire about  $100 \mu\text{m}$  long. At higher power, the right-side end trap population becomes large enough to blue-shift itself to the same energy as the level region, forming a single mono-energetic condensate on the right side.

Figure 5.4 shows the energy and in-plane momentum ( $k_{\parallel}$ ) distribution of the polaritons at two different applied voltages. A population near  $k_{\parallel} = 0$  and at slightly lower energy is visible, and is emitted from the end trap. The larger population is from the level region of the wire, and clearly has an overall non-zero in-plane momentum, even with zero applied voltage. This is due to flow away from the pump spot, which is generally in the  $+x$ -direction. With applied negative voltage [Figure 5.4(b)], the overall momentum is clearly reduced. This voltage corresponds to conventional current flowing in the  $+x$ -direction, and thus electron current flowing in the  $-x$ -direction, which opposes the polariton flow. This clearly shows an effect of drag upon the polaritons from the electrical current. When the wire is pumped on the other end, the overall effect is reversed, with the polaritons moving with momentum in the opposite direction. However, because the gradient of the cavity is in the opposite direction relative to the flow away from the pump in that case, the observation of the overall drag effect is reduced, as I discuss in Section 5.3.2.

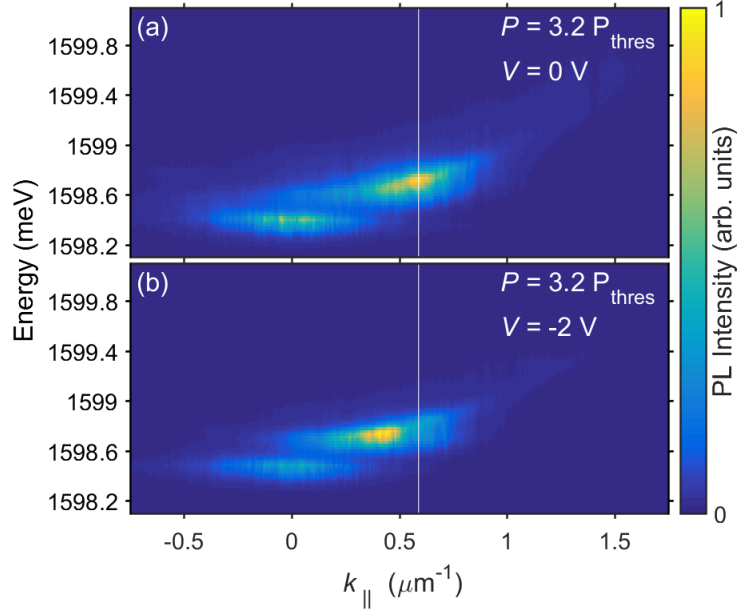


Figure 5.4: Photoluminescence intensity as a function of energy and in-plane momentum ( $k_{\parallel}$ ). The threshold power ( $P_{\text{thres}}$ ) was about 114 mW, and the powers used were both about  $3.2P_{\text{thres}}$ . The applied voltage in each case was (a) 0 V and (b) -2 V. The white vertical line is marking the  $k_{\parallel}$  value of the peak of the distribution at zero applied voltage to aid in comparing the two images. The photoluminescence intensity was normalized separately for each image.

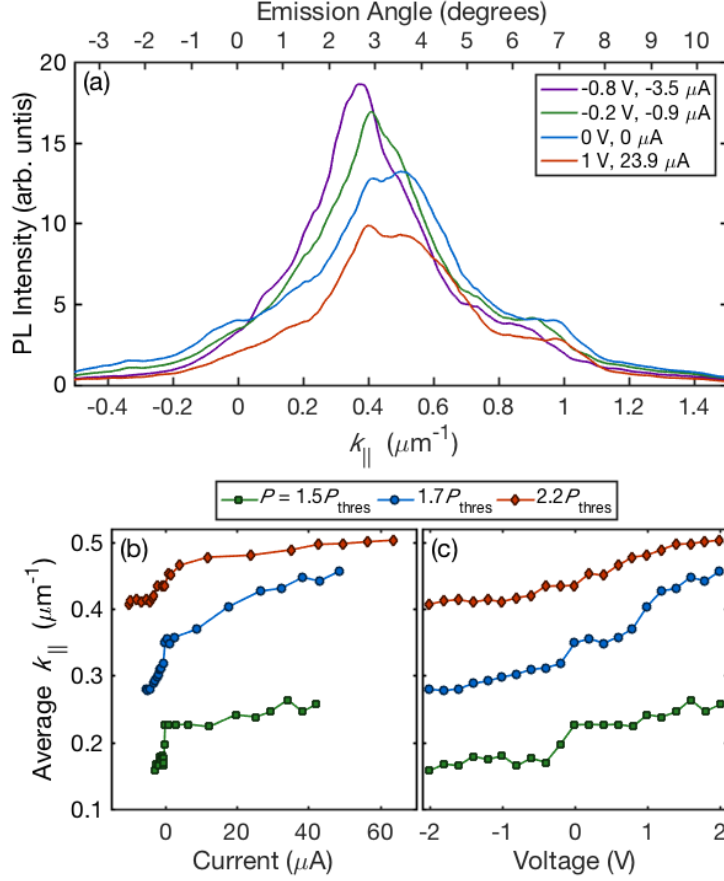


Figure 5.5: (a) Photoluminescence intensity as a function of  $k_{||}$  at various applied voltages, taken by integrating along the energy axis in images similar to Figure 5.4. These data are for a wire similar to that used for Figure 5.4, but with a threshold power of  $P_{\text{thres}} \approx 130$  mW. The power used in this case was  $2.2P_{\text{thres}}$ . (b) The average  $k_{||}$  vs. current for various pump powers. (c) The average  $k_{||}$  vs. applied voltage for various pump powers. The red diamonds in both (b) and (c) correspond to the same data used for (a).

Figure 5.5(a) shows the polariton distribution as a function of  $k_{\parallel}$  for multiple applied voltages. Figures 5.5(b) and (c) show the average  $k_{\parallel}$ , from plots similar to and including Figure 5.5(a), as functions of current and applied voltage, respectively. A clear shift in the distribution toward lower  $k_{\parallel}$  is visible with applied negative voltage, indicating drag upon the polaritons. When the current direction is reversed, there is very little effect on the overall momentum with applied voltage. The simplest explanation for this is that the cross section for electron-polariton collisions is proportional to their relative momentum, so that electrons moving oppositely to the polaritons impart significant momentum, while electrons moving in the same direction as the polaritons impart much less momentum. Simple calculations indicate, however, that the velocities of the electrons and polaritons are somewhat different. The polariton velocity is measured directly from their momentum;  $0.5 \mu\text{m}^{-1}$  corresponds to  $v = \hbar k/m = 6 \times 10^7 \text{ cm/s}$ . The electron velocity can be estimated from the mobility, which our measurements give as approximately  $100 \text{ cm}^2/\text{V-s}$ . For a voltage drop of 1 V over  $100 \mu\text{m}$ , this gives electron velocities of the order of  $10^4 \text{ cm/s}$ . By comparison, the sound velocity of the polaritons is nominally of the order of  $v_p = \sqrt{Un/m}$ , which is  $4 \times 10^6 \text{ cm/s}$  for the expected polariton-polariton interaction strength of  $U \simeq 1 \mu\text{eV-}\mu\text{m}^2$  [32, 59, 120] for our multiple QW structure and our estimated density of the condensate of  $1/\mu\text{m}^2$ . This indicates that the average polariton velocity is higher than their sound velocity, so that they are not superfluid, and therefore their interactions with free electrons can be treated as single-particle scattering.

As seen in Figure 5.5(b), the saturation velocity for current in the direction co-moving with the polaritons increases with increasing density. One factor that contributes to this is that the gradient of the effective potential felt by the polaritons is greater for higher pump intensities, due to the larger exciton cloud potential. The density dependence of the velocity may also indicate that many-body effects come in to play. This overall velocity increase with density is discussed further in Section 5.3.4.

A possible explanation for the overall effect is based on the stimulated relaxation of moving and uncondensed excitons into the condensate, which results in the formation of a moving condensate. Since the excitons can respond to drag by the electric current, the motion of the condensate is indirectly affected by the current [124].

### 5.3.2 Pumping on the high-energy end

In the preceding section, the data shown and discussed were for the specific case of pumping the wire on the low-energy end, with respect to the cavity gradient. In principle, the asymmetry in momentum can be flipped to the opposite direction by pumping on the opposite side. However, the gradient prevents a clean symmetric change of flow direction. Figure 5.6 shows a summary of the data when pumping on the high-energy end of the wire. In this case, since the polaritons flow downhill away from the pump, they do not form a mono-energetic condensate along the wire [Figure 5.6(a) and (b)]. In the opposite case, with flow on a slight uphill away from the pump, the exciton reservoir helps to level the potential by filling the lowest energy end, and the polariton density is naturally higher near the source. This results in a level potential at relatively low densities. In this case, however, the reservoir actually makes the potential even less level, and the polaritons must flow all the way to the opposite end in order to form a level potential. The polaritons do clearly flow the  $\approx 175 \mu\text{m}$  to the far end and form a condensate, but there is no single condensate along the majority of the wire.

Figure 5.6(c) shows the energy and momentum distribution under similar conditions. The flow is generally in the  $-x$ -direction as expected, but there is not a significant peaked population at large non-zero  $k_{\parallel}$ . Comparing similar images at different applied voltages reveals no easily noticeable shift in momentum, unlike the case shown in Section 5.3.1. However, a small effect is revealed by integrating along the energy axis [Figure 5.6(d)], which shows a small portion of the population between about  $-0.2$  and  $-0.4 \mu\text{m}^{-1}$  responding to applied voltage. This population shifts toward negative momentum with negative voltage and positive momentum with positive voltage, as expected from the more dramatic shifts shown in Section 5.3.1. This effect is small, but the overall change in polariton flow from positive to negative in-plane momentum when changing the pump location confirms our claim that the overall momentum is largely due to flow from the pump spot.

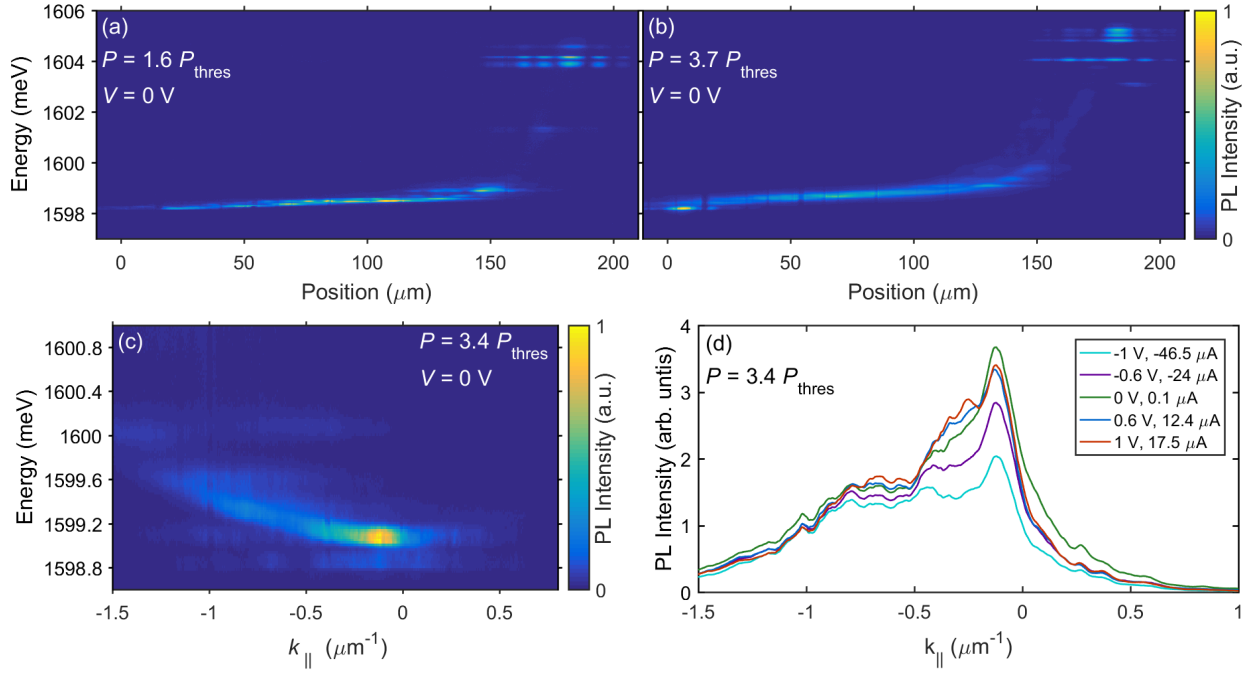


Figure 5.6: (a) and (b) show photoluminescence intensity of the polaritons vs. energy and position along the length of the wire at zero applied voltage. The threshold power was about 150 mW under these conditions, and the powers used were (a)  $1.6P_{\text{thres}}$  and (b)  $3.7P_{\text{thres}}$ . (c) The photoluminescence intensity of the polaritons as a function of energy and in-plane momentum ( $k_{\parallel}$ ) at zero applied voltage. This image is from a different device from (a) and (b), so the energies are not directly comparable. However, the threshold power was similarly about 150 mW, and the power for the image shown was about  $3.4P_{\text{thres}}$ . (d) photoluminescence intensity integrated along the energy axis for data from the same device as (c) and under the same pump conditions.

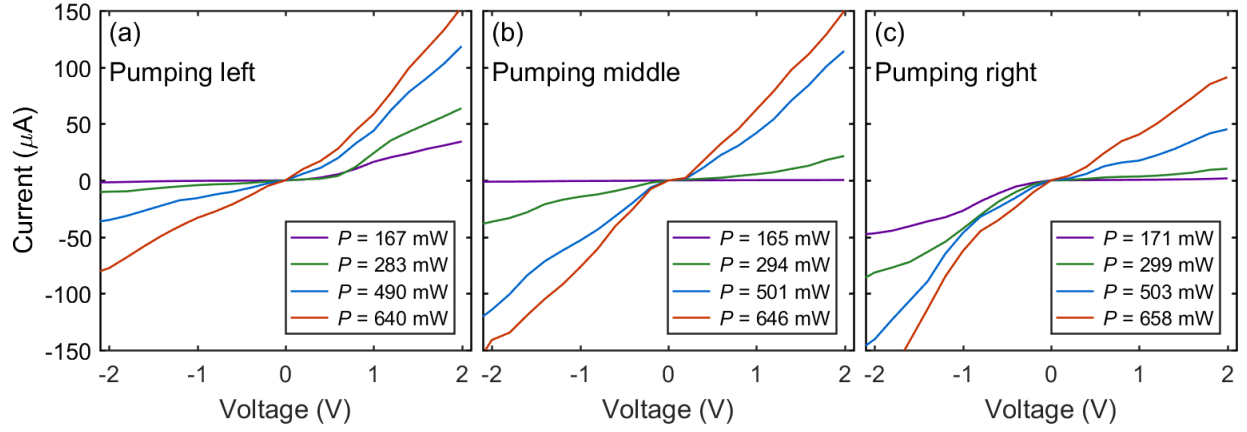


Figure 5.7: Current vs. voltage for various pump powers and with different pump spot locations. The locations are given in the upper left of each plot, and are defined in reference to Figure 5.2. Specifically, they are the (a) low energy end, (b) middle, and (c) high energy end of the wire, which correspond to the left, middle, and right sides of the  $x$ -axes used throughout this work. The approximate condensate threshold power for each case was (a) 130 mW, (b) 130 mW, and (c) 150 mW.

### 5.3.3 Current-voltage characteristics

The current vs. voltage characteristics are similar to what was observed in the square pillar devices of Chapter 4, and can be seen in Figure 5.7. All data shown in the three separate axes are from the same device and pumped under similar conditions. An overall asymmetry is apparent when pumping on either end of the wire, and especially at lower powers. The asymmetry mostly disappears when pumping in the middle, indicating that it is due to the pump location. Furthermore, the asymmetry is nearly exactly opposite for pumping on opposite ends, which also indicates that is primarily due to the pump spot location. We attribute this to greater illumination of the contact near the pump spot, creating free carriers that can carry current over the  $n - i$  band bending barrier, as discussed in Section 4.3. In this case the pump is definitely much closer to one side than the other, and the asymmetry flips directions accordingly.

### 5.3.4 Maximum velocity

As discussed in Section 5.3.1, the application of positive voltage when pumping on the low energy end of the wire has little or no effect on the polariton momentum. We propose that this may be because the condensate has reached a maximum velocity simply from the flow away from the pump spot, so drag that would increase the velocity will have no effect on the velocity of the condensate. Following this reasoning, we plotted the average  $k_{\parallel}$  vs. the relative polariton density, shown in Figure 5.8. Since the distribution is strongly peaked in  $k_{\parallel}$ , and the shape is maintained overall as it shifts, the average  $k_{\parallel}$  is fairly representative of the maximum. The  $k_{\parallel}$  dependence on polariton density above threshold is very well fit by a power law which goes as  $n^{0.354}$ , where  $n$  is the polariton density. At very high densities the average momentum saturates. As discussed in Section 5.3.1, the sound velocity of the polaritons increases as  $n^{0.5}$ . If many body effects are involved, such as a critical velocity of superfluidity, this is the natural unit of velocity for these effects. While the pump power has a direct influence on the expected sound velocity via an increase in density, changes in power also affect the potential felt by the polaritons. Higher powers result in a larger exciton cloud and thereby a larger force on the polaritons moving away from the pump, which also affects

the overall velocity of the polaritons.

## 5.4 CONCLUSIONS

Overall, we accomplished the primary goal of observing an effect of an injected DC current on the state of the polaritons. We demonstrated that a DC current can directly alter the momentum of photons moving in a cavity, which has the direct effect of changing the angle of emission. This means that this polariton drag effect can be considered to be beam steering using a DC current to tune the angle of a light beam. Future work should focus on showing the effect symmetrically, such that the shift in emission angle switches direction for current in the opposite direction. This will require a symmetric device, with the polariton wires precisely aligned perpendicularly to the gradient such that the polariton energy is constant along the length of the wire. A first demonstration in such a device would involve the same measurements described above, with the effect clearly switching directions when the pump is moved from one end of the wire to the other.

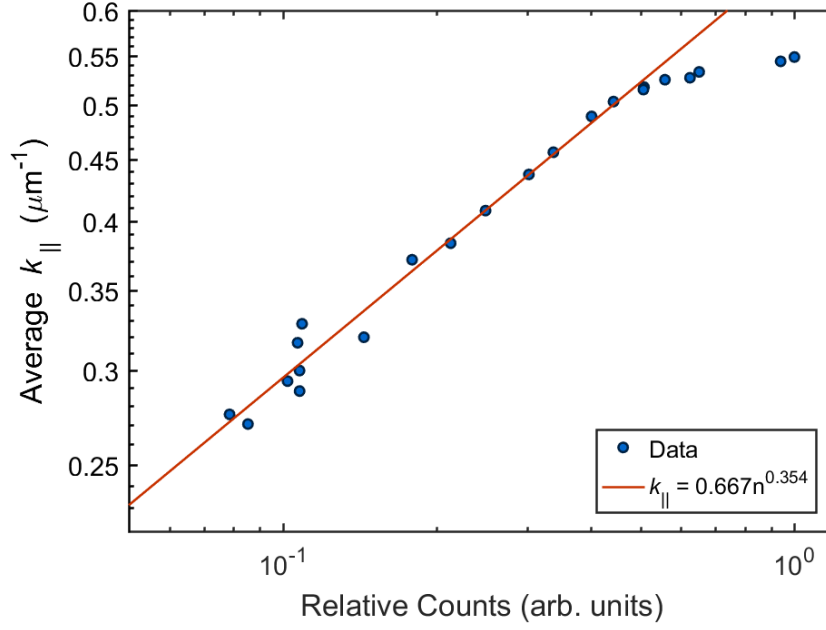


Figure 5.8: Average  $k_{\parallel}$  vs. polariton counts, which is effectively the relative polariton density. The data come from averaging the energy-integrated curves similar to Figure 5.6(d) at many pump powers, but with the pump in the placement of Section 5.3.1. The minimum pump power used was 200 mW, which is  $1.8P_{\text{thres}}$ . The data was fit using a power law, which is given in the figure legend.

## 6.0 LONG RANGE EXCITON TRANSPORT

### 6.1 MOTIVATION

In this chapter, I will discuss the results first reported in Ref. [125], and this will largely be an adaptation of that work. The idea behind this experiment was to attempt to explain the large discrepancy in a measurement of the interaction strength of polaritons in a ring-shaped pump geometry [85] compared to other works (e.g. Refs [120, 126]) and the theoretically expected value [32]. As I discussed in Section 1.4.7, the interactions of the polaritons come solely from the exciton part. The energy blue shift due to interactions can therefore be related to the polariton density  $n_{\text{pol}}$ , the exciton fraction of the polaritons  $f_{\text{exc}}$ , and the pure exciton-exciton interaction strength  $g_{\text{xx}}$  by

$$\Delta E(k_{\parallel}) = g_{\text{xx}} f_{\text{exc}}(k_{\parallel}) \sum_{k'_{\parallel}} f_{\text{exc}}(k'_{\parallel}) n_{\text{pol}}(k'_{\parallel}). \quad (6.1)$$

This can be simplified by approximating all of the particles as fitting neatly into two populations: reservoir excitons and polaritons. The reservoir excitons can consist of both excitons that are not strongly coupled to the photon mode and particles in the lower polariton branch that are at such high in-plane momentum that they are essentially excitons. If the “true” polaritons at low momentum are mostly near the ground state, the blue shift to the ground state energy can be approximated by

$$\delta_0 E = g_{\text{xx}} f_{\text{exc}} n_{\text{res}} + g_{\text{xx}} f_{\text{exc}}^2 n_{\text{pol}}. \quad (6.2)$$

In an energy- or  $k_{\parallel}$ -resolved experiment, only the polariton density can be directly measured by photoluminescence, since the reservoir consists of either dark states (the excitons that do

not emit directly into a single photon) or states that emit outside the numerical aperture of the collection optics. This means that a direct measurement of the interaction strength using measured energy shifts and measured polariton densities must be done in an area in which  $n_{\text{res}}$  is negligibly small.

This was the approach taken in Ref. [85]. This study used a ring trap formed by the nonresonant pump (see Section 2.2.1) about 40  $\mu\text{m}$  in diameter and then measured the energy blue-shift and the density of the polaritons in the center of the ring far from the supposedly stationary reservoir population. The resulting value determined for  $g_{\text{xx}}$  was  $\sim 1000$  times greater than the expected value. It is highly unlikely that an incorrectly measured blue shift was the cause of the discrepancy, due to the directness of the measurement with spectroscopy, as well as easily established uncertainty bounds. Similarly, the polariton density is somewhat directly measured via the photoluminescence intensity, though this is dependent on careful determination of the Hopfield coefficients, as I discuss in Section 6.3. We suspected, then, that the more likely source of such a large discrepancy was a large population of “excitons” outside of the detection range of the optics that contributed to the measured energy blue-shift.

This experiment was intended to test that hypothesis.

## 6.2 SOME OPENING CONSIDERATIONS

Exciton transport is normally assumed to be limited by diffusion. This has been measured in semiconductor quantum wells [127–129] and in organics [130, 131] as being  $\lesssim 1 \mu\text{m}$ . It seems reasonable, then, to assume that particles within the lower polariton branch but at very high exciton fractions also have extremely short transport distances similar to bare excitons, in contrast to the long-range flow seen in lower polaritons with significant photon fractions (e.g. Ref. [25]). This assumption is typical in the exciton-polariton literature (e.g. Refs. [78, 82, 84, 132, 133]), in which a sharp distinction is made between polaritons at low  $k_{\parallel}$  and the exciton reservoir at higher momentum. This was also the assumption used in Ref. [85], and is mostly confirmed in experiments as well. As discussed in Section 2.2.1,

the repulsive interactions between polaritons and excitons are often used to form various potential profiles [23, 26, 81, 82, 84, 92, 134, 135]. From these studies, it can be seen that most of the excitons are essentially stationary. The profile of their distribution can be deduced from the energy profile of the polaritons, and largely resembles the profile of the pump spot. However, the details of the distribution matter, especially when attempting to measure the strength of the interactions when using a nonresonant pump. As mentioned above, a direct measurement requires isolating a sufficiently large population of observable polaritons far from *all* undetectable excitons. Conceivably, most of the reservoir could indeed be stationary, but even a small percentage that is capable of longer transport distances could affect a measurement.

This is even more believable when considering the strange shape of the lower polariton dispersion curve at intermediate in-plane momenta (Figure 1.12). While at very high  $k_{\parallel}$  the polaritons become no different from excitons, and at very low  $k_{\parallel}$  they can be treated as regular low-mass particles with a parabolic dispersion, in between these regions the dispersion is rather odd, with an inflection point and negative curvature, and is quite steep even when the polaritons are already mostly “excitonic” as given by the Hopfield coefficients. As I discuss in Section 6.5.1, this gives rather large group velocities in particles that are primarily excitonic, and consequently large transport distance estimates. In addition, recent theoretical work using a simple model of a 1D chain of quantum emitters has already shown that excitons could move much longer distances when strongly coupled to a photon mode [136]. The results I discuss in this chapter add experimental evidence that long-range transport does happen in the highly excitonic lower polaritons near the lower inflection point (there is another inflection point at very large  $k_{\parallel}$  due to the upward-curving exciton branch).

We therefore divide the populations in the lower polariton branch into three parts. The “normal” lower polaritons (which I will refer to in this chapter as just “polaritons”), the “bare excitons” at very large  $k_{\parallel}$ , and the “bottleneck excitons” in between. The lower inflection point is conveniently placed for distinguishing between these populations. We consider particles below it to be the “normal” polaritons, and they really do have a normal near-parabolic dispersion. Particles just above it are considered to be the “bottleneck excitons,” so-named due to the low density of states at the inflection (from the steep derivative of the

energy with respect to  $k_{\parallel}$ ) which leads to a bottleneck effect that prevents thermalization of particles to the low  $k_{\parallel}$  states [62, 137]. This low density of states at the inflection also leads to low occupation, allowing a clear distinction between particles above and below it. Since the bottleneck excitons are mostly excitonic, they have many of the properties of the bare excitons, but we show that they can travel much longer distances. Particles at higher  $k_{\parallel}$  are the bare excitons, often identified as the reservoir excitons, which have a very flat dispersion, a much heavier mass, and subsequently much lower speeds and travel distances. The bare excitons are the particles mostly responsible for the observed potential profiles; the highly mobile bottleneck excitons exist in much lower numbers, but still have a significant population compared to the normal polaritons, as I will discuss below.

In our samples, the lower inflection point (at  $k_{\parallel} = k_{\text{inflection}}$ ) usually corresponds to angles of emission of about  $20^\circ$ , corresponding to a modest NA of 0.34. Thanks to the ability to rotate the cryostat (see Section 6.4), we were able to collect up to about  $k_{\parallel} = 2k_{\text{inflection}}$ . Everything above the inflection and up to  $2k_{\text{inflection}}$  is considered to be a bottleneck exciton, while everything below is considered a polariton.

### 6.3 GETTING PARTICLE COUNTS FROM MEASUREMENTS

The primary collected data for this experiment was photoluminescence intensity as a function of energy and emission angle (using energy- and angle-resolved imaging). The measured photoluminescence intensity is proportional to the number of photons emitted, with the sensitivity given by a combination of all of the losses of the optics and the quantum efficiency of the detector itself. This can be calibrated as a function of wavelength by sending a beam of known wavelength and power through the collection optics and into the same detector. However, for this experiment, we were only concerned with the relative populations, so an exact calibration was unnecessary. Instead, we simply assumed that the population at a particular point in emission angle and energy was proportional to the “counts” on each pixel of the camera, and that the wavelength dependence was negligible for the narrow range of photoluminescence wavelengths we collected.

A simple set of rate equations can be written for the photons and polaritons at a particular energy and in-plane momentum, following the method discussed in Section 1.4.5:

$$\begin{aligned}\frac{dN_{\text{pol}}}{dt} &= G - \frac{N_{\text{pol}}}{\tau_{\text{pol}}} \\ \frac{dN_{\text{phot}}}{dt} &= \frac{N_{\text{pol}}}{\tau_{\text{pol}}}\end{aligned}\tag{6.3}$$

Here  $N_{\text{pol}}$  is the polariton population,  $N_{\text{phot}}$  is the emitted photon population (the photons emitted from the cavity),  $G$  is the pumping term, and  $\tau_{\text{pol}}$  is the polariton lifetime. These equations simply say that photons are created at the rate that the polaritons decay. In the steady state for the polariton population,  $N_{\text{pol}}$  eventually settles on a constant value, while the photons are emitted at a constant rate from the cavity. Therefore, the photon rate equation can be integrated to give

$$\Delta N_{\text{phot}} = \frac{N_{\text{pol}}}{\tau_{\text{pol}}} \Delta t,\tag{6.4}$$

where  $\Delta N_{\text{phot}}$  is a counted number of photons in a time period  $\Delta t$ . That time period is simply the exposure time of the camera (or the exposure time times the duty cycle for a chopped pump), but again we do not need specific values. The important result is that the counts on the camera are proportional to the counted photons, which are in turn proportional to the number of polaritons divided by the polariton lifetime.

This gets a bit tricky, though, because the polariton lifetime is highly dependent on cavity and exciton fractions (see Section 1.4.5), which are themselves dependent on  $k_{\parallel}$ . For this experiment, we assumed that the exciton lifetime was large enough that its contribution was negligible, such that  $\tau_{\text{pol}} \approx \tau_{\text{cav}}/f_{\text{cav}}$ , where  $\tau_{\text{cav}}$  is the cavity lifetime. For the ranges of  $k_{\parallel}$  carefully studied in this experiment, it is easy to show that this is a very good approximation as long as the exciton lifetime is about 100 times greater than the photon lifetime, and still does not affect the overall results of the analysis if even only 10 times greater. Therefore, since the cavity lifetime is constant, the relative number of particles can be deduced from the photoluminescence intensity and cavity fraction at each  $k_{\parallel}$ .

For further details on the data processing, especially establishing the error bounds, see Section E.2.

## 6.4 EXPERIMENT

The sample used in this experiment is the same as the one described in Section 2.1. In particular for this experiment, the lower polariton energy gradient was pointing in the direction  $8.6^\circ$  from the  $+x$ -direction as defined in Figure 6.1. The magnitude of the gradient was measured to be  $\approx 5.2$  meV/mm at the most excitonic detuning ( $\delta_0 \approx 7$  meV) used in this study.

To produce polaritons, the sample was pumped non-resonantly at an energy of 1726.8 meV through the imaging objective lens, producing a pump spot with  $\approx 3$   $\mu\text{m}$  FWHM (Appendix E has the details for this measurement). The sample was held at  $\sim 5$  K within a cryostat, which could be rotated around an axis perpendicular to the optical axis of the imaging objective. The numerical aperture of the objective was 0.40. This allowed collection of emission angles of about  $-3^\circ$  to  $43^\circ$  with the cryostat rotated to  $20^\circ$ , and up to  $64^\circ$  with the cryostat rotated to  $40^\circ$ . A narrow slit was placed at a secondary real space imaging plane, which could be adjusted to select various regions of the image for angle-resolved imaging. Figure 6.1 shows the most important details of the setup, and Figure D.3 has further specifics about the optics.

With the pump tightly focused, the real space slit was swept through the region containing the pump spot while a CCD collected angle- and energy-resolved images of the photoluminescence. This was done both near resonance ( $\delta_0 \approx 2$  meV) and at more excitonic detuning ( $\delta_0 \approx 8$  meV), using pump powers below the quasi-condensate threshold power, referred to as  $P_{\text{th}}$  throughout this chapter (see Section 3.6 for a discussion of the various thresholds). For the angle-resolved images, the spectrometer slit selected a slice of the  $k_{\parallel}$  plane along the  $x$ -axis. The slit was closed to 40  $\mu\text{m}$  for all of the images. This corresponded to a selection width of  $0.034 \pm 0.004$   $\mu\text{m}^{-1}$  in the  $k_{\parallel}$  plane.

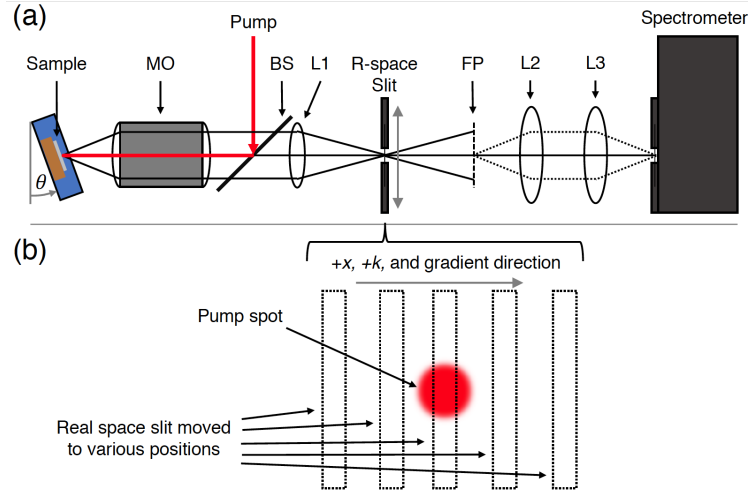


Figure 6.1: (a) A diagram showing the basics of the optical setup, viewed from above. The sample was mounted in a cold-finger cryostat, which could be rotated by angle  $\theta$ . The pump laser was reflected off a beam splitter (BS) through the imaging microscope objective (MO). The photoluminescence was collected by the objective and then imaged by lens L1 onto a secondary real-space plane. At this plane, a movable slit (R-space Slit) was placed to select regions of the sample from which to resolve photoluminescence. The Fourier plane of the objective was also imaged by lens L1 at location FP. Lenses L2 and L3 then imaged this secondary Fourier plane onto the slit of the spectrometer. (b) A diagram of the real-space plane at the location of the slit (R-space Slit in (a)) as viewed along the imaging axis. The slit could be moved horizontally to select different regions of the image without changing the pump location. The  $+x$ ,  $+k$ , and cavity gradient (“uphill”) directions are all the same.

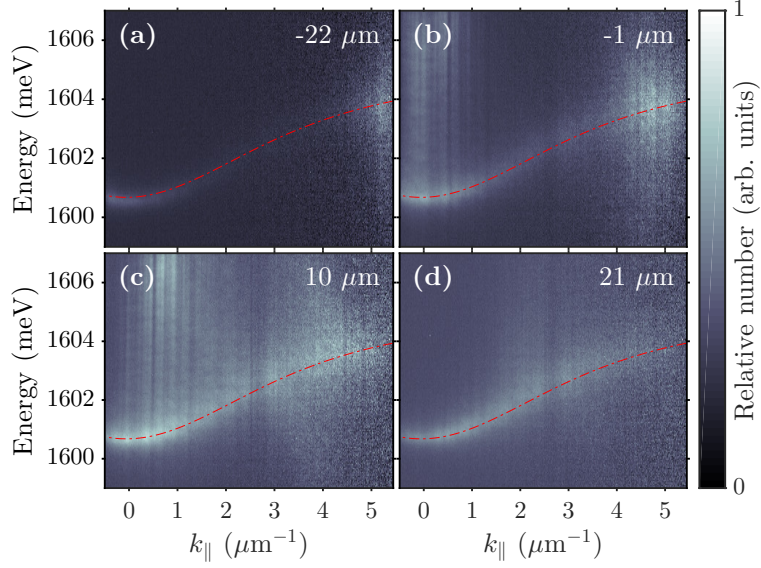


Figure 6.2: Normalized lower polariton population as a function of energy and  $k_{\parallel}$ , taken from angle-resolved images and adjusted for the  $k_{\parallel}$ -dependent photon fraction to show the relative particle populations. The pump power was about  $P_{\text{th}}/2$ , and the detuning was about 8 meV. The positions of the real-space filter with respect to the pump spot are given in the upper right corners of each plot. The red lines show the theoretical lower polariton dispersion. The counts for each image were normalized separately, so the counts of separate images are not comparable.

## 6.5 LONG-RANGE TRANSPORT

Figure 6.2 shows examples of the energy- and angle-resolved data, collected with the sample rotated to an angle of  $20^\circ$  to allow high angle imaging. Since the sample was held at an angle relative to the imaging objective, the range of collection angles is asymmetric. The data were adjusted at each  $k_{\parallel}$  to give a relative number of particles from the photoluminescence intensity, as discussed above in Section 6.3.

The lower polariton dispersion in one direction is clearly visible, with a parabolic shape at low  $k_{\parallel}$  and an inflection near  $2.6 \mu\text{m}^{-1}$ . The dispersion flattens out at high  $k_{\parallel}$  as the lower polariton energy approaches the nearly flat exciton energy. Above the inflection, the energy linewidth increases due to the high exciton fraction and the broader exciton energy linewidth. Figure 6.3 shows the same data as Figure 6.2, but integrated over energy, clearly revealing the  $k_{\parallel}$  dependence of the populations within the lower polariton band.

Figure 6.4(a) shows the relative populations for different ranges of in-plane momentum as the real-space filter is swept across the pump spot. This particular figure corresponds to polaritons and bottleneck excitons moving only in the  $+x$ -direction because of our choice of the range of collection angles. As expected for motion in this direction, the total number of “normal” polaritons below the inflection point in  $k_{\parallel}$  peaks on the positive side of the pump spot, which was located by looking at the symmetric range of  $k_{\parallel}$  near  $k_{\parallel} = 0$  (from  $-0.3$  to  $0.3 \mu\text{m}^{-1}$ ). The cause of this offset of the peak is that a majority of the polaritons were moving away from the pump spot, and in this case only those moving in the  $+x$ -direction were detected due to the rotation of the sample. Since the real-space slit had a width of about  $20 \mu\text{m}$ , the largest number of polaritons (those nearby the pump and moving in the  $+x$ -direction) were therefore detected when the slit was centered about  $10 \mu\text{m}$  to the  $+x$  side of the pump spot. The relative number of bottleneck excitons with momentum above the inflection point peaks closer to the pump spot, but is spatially broader overall. One feature to note is that there is asymmetry around the peak, with higher relative counts on the left side compared to the polaritons below the inflection, corresponding to back-scattered particles moving back toward the pump spot. These particles are clearly traveling back toward the pump spot since they are detected with momentum in the  $+x$ -direction while on the  $-x$

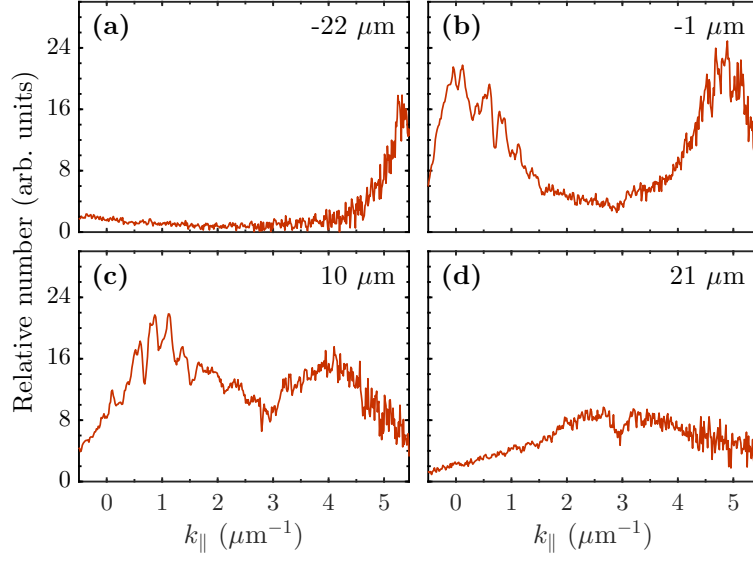


Figure 6.3: Relative lower polariton population as a function of in-plane momentum ( $k_{\parallel}$ ), derived from angle- and energy-resolved images integrated over energy. The inflection point of the lower polariton dispersion is near  $2.6 \mu\text{m}^{-1}$ . For these data, the sample was at  $\theta = 20^\circ$  with respect to the imaging objective lens [as shown in Figure 6.1(a)], the pump power was about  $P_{\text{th}}/2$ , and the detuning was about 8 meV. The positions of the real-space filter with respect to the pump spot are given in the upper right corners of each plot. The uncertainty of the relative number values is approximately the same as the scatter in the data.

side of the pump spot. This population of backward-moving, highly excitonic polaritons is visible in Figures 6.2(a) and 6.3(a) and is discussed below in Section 6.6.

As mentioned above, we consider all of the emission from above the inflection to be part of the bottleneck exciton population. These bottleneck excitons, even at the resonant point on the sample [where  $E_{\text{exc}}(k_{\parallel} = 0) = E_{\text{cav}}(k_{\parallel} = 0)$ ], have excitonic fractions of at least 0.70, and are often undetected, e.g., in experiments like Ref. [85], due to their high emission angle ( $\gtrsim 20^\circ$ ). We note, however, that as seen in Figure 6.4(a), the population of these high-momenta bottleneck excitons is comparable to that of the entire thermalized population at low momentum (comparing the red data to the blue data). Although the occupation of these high-energy states is strongly suppressed in equilibrium by the Boltzmann factor  $e^{-\Delta/k_B T}$ , the polariton bottleneck effect prevents full thermalization due to the suppressed phonon emission rate for excitons in these states. At higher densities, stimulated collisional effects (see Section 1.4.5) can thermalize the polariton gas much more effectively [35]. However, the experiments reported here were performed with excitation intensities well below the critical density threshold  $P_{\text{th}}$  for Bose-Einstein condensation of polaritons, so the effect of stimulation was small. This suppresses collisional thermalization, similar to the low-density conditions of Ref. [85].

We can assume that the contribution of the cavity gradient in the relatively short distances shown is negligible, since it gives an energy change of only about  $520 \mu\text{eV}$  in  $100 \mu\text{m}$ , which is small on the scale of the dispersions shown in Figure 6.2. Therefore, due to symmetry, the particles with positive  $k_{\parallel}$  on the negative side of the pump spot can be assumed to have the same distribution in space as the particles with negative  $k_{\parallel}$  on the positive side of the pump spot. This is utilized in Figure 6.4(b), which simply adds the two sides of the positive  $k_{\parallel}$  distribution together at equal distances. This estimates the full population integrated over both momentum directions.

Given the circular symmetry of the experiment, we can also assume that the polaritons and bottleneck excitons move radially outward from the pump spot. This means that, for the slice of the  $k_{\parallel}$  plane collected, only particles moving along a radial line parallel to that slice make a significant contribution to the measured population. Therefore, the experiment can be simulated with a 1D spatial distribution. We used a Voigt profile to produce this

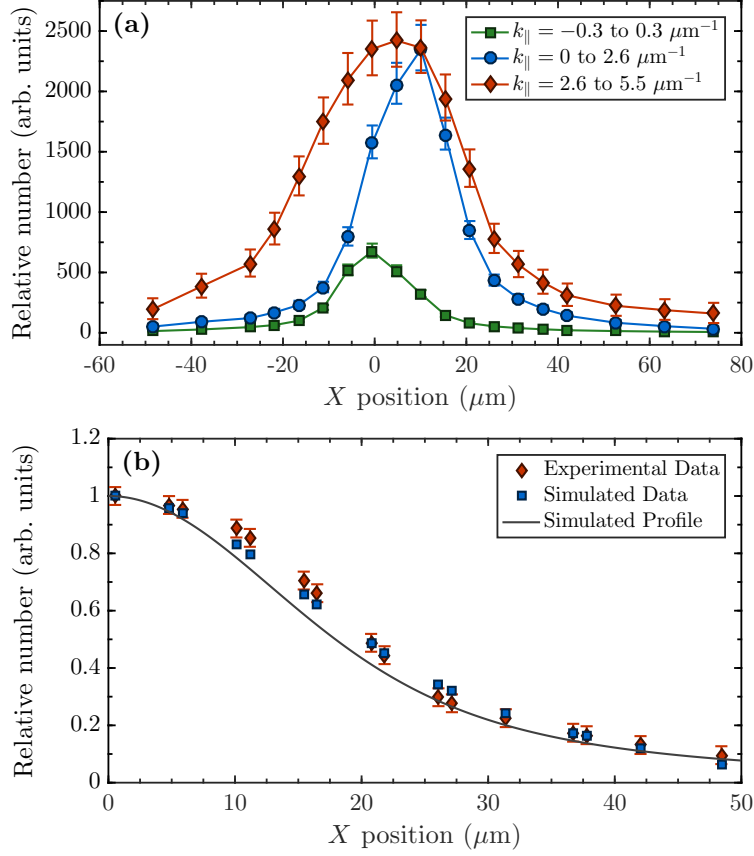


Figure 6.4: (a) The relative number of particles within the lower polariton band as a function of the real-space filter position for various ranges of  $k_{\parallel}$  for a pump power of about  $P_{\text{th}}/2$  and detuning of about 8 meV. The zero point in position was set by looking only at the symmetric range of visible  $k_{\parallel}$  near  $k_{\parallel} = 0$  (green squares). (b) The experimental data (red diamonds) are the sum of the relative numbers at opposite sides of the pump spot (opposite  $x$  positions) for the  $k_{\parallel}$  range of 2.6 to  $5.5 \mu\text{m}^{-1}$ , shown as red diamonds in (a). The simulated profile (solid black line) is the Voigt profile representing the real-space distribution of the bottleneck excitons, which was normalized to show its shape compared to the data. The simulated data (blue squares) come from integrating the Voigt profile over small bounds in  $X$ , simulating the effect of the real-space slit in acquiring the experimental data. For details about the error bounds, see Section E.2. [Simulated results were produced by Shouvik Mukherjee.]

approximate distribution:

$$N(x) = \int_{-\infty}^{\infty} \frac{Ae^{-x'^2/2\sigma^2}}{(x - x' - x_0)^2 + \gamma^2} dx'. \quad (6.5)$$

By integrating this distribution with bounds similar to those given by the real-space slit width (typically about  $20 \mu\text{m}$ ), we were able to reconstruct the collected real-space integrated data.

This method returned a FWHM for the distribution of the bottleneck excitons of  $36 \pm 1 \mu\text{m}$  at a pump power of  $P_{\text{th}}/2$ . It is also mostly unaffected by pump power below threshold. This is much larger than the diffusion lengths for bare excitons in quantum wells, which, as discussed above, are typically of the order of  $1 \mu\text{m}$ . This result also indicates that the reservoir exciton population far from the excitation region is not negligible, contrary to the assumption in Ref. [85]. See Appendix E for details and data for additional pump powers and detunings.

### 6.5.1 Transport distance estimate

A simple calculation using the group velocity of the lower polariton band can be used to explain the overall effect. We assume that the exciton-polaritons we observe travel ballistically until they scatter out of the field of view of our detection or decay radiatively. The distance traveled is then approximately the group velocity of the particles times their effective time spent traveling ballistically. The group velocity can be easily calculated as  $v_g = \frac{1}{\hbar} \frac{dE}{dk_{\parallel}}$ , which is plotted as a function of  $k_{\parallel}$  in Figure 6.5(a), using the energy dispersion for  $\delta_0 \approx 8 \text{ meV}$  (shown in Figure 6.2). A particle loss time ( $\tau_{\text{loss}}$ ) can be estimated by assuming loss rates from both radiative decay and scattering (since this experiment was done at low polariton density, the thermalization rate is assumed to be negligible). For the same reasons given in Section 6.4 (that the flow of particles is radial outward from the pump), we can treat scattering similarly to radiative loss, since any change of momentum away from the narrow slice in  $k_{\parallel}$  that we collected for the experiment will result in the particle leaving the field of view. This leads to a loss time of  $\tau_{\text{loss}} = (1/\tau_{\text{pol}} + 1/\tau_s)^{-1}$ , where  $\tau_{\text{pol}}$  is the radiative decay time and  $\tau_s$  is the scattering time.

As discussed in Section 6.3, the radiative decay time is mostly dependent on the cavity photon decay time  $\tau_{\text{cav}}$  and cavity fraction  $f_{\text{cav}}$ , giving  $\tau_{\text{pol}} \approx \tau_{\text{cav}}/f_{\text{cav}}$ . The scattering time

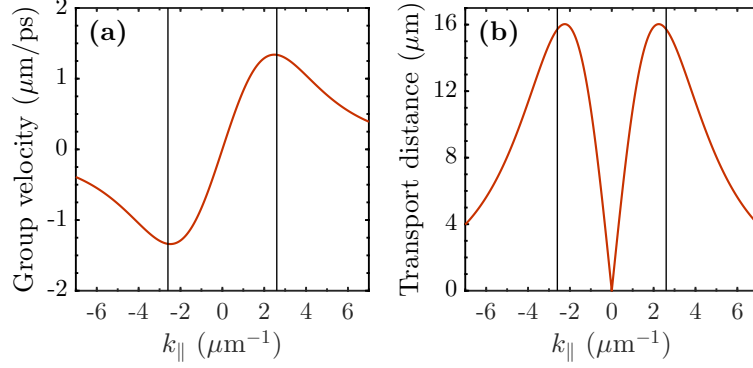


Figure 6.5: (a) The group velocity of the lower polariton band for the actual sample parameters at a detuning of about 8 meV. (b) The estimated transport distance, assuming a cavity photon decay time of 100 ps and an exciton scattering time of 10 ps. The solid vertical lines mark  $\pm 2.6 \mu\text{m}^{-1}$ .

is similarly dependent on the exciton fraction  $f_{\text{exc}}$ , since the excitonic part is the primary part that undergoes scattering, giving  $\tau_s \approx \tau_{s,\text{exc}}/f_{\text{exc}}$ , where  $\tau_{s,\text{exc}}$  is the exciton scattering time.

By multiplying the group velocity by the loss time, a transport distance can be estimated as

$$d \approx v_g \tau_{\text{loss}} \approx v_g \left( \frac{f_{\text{cav}}}{\tau_{\text{cav}}} + \frac{f_{\text{exc}}}{\tau_{s,\text{exc}}} \right)^{-1}, \quad (6.6)$$

which is plotted in Figure 6.5(b), using a cavity photon decay time of 100 ps and an exciton scattering time of 10 ps. This plot shows that the distance is strongly peaked near the inflection points, with significant populations both above and below the inflection point traveling the longest distances. These parameters give an estimate consistent with our measured result of transport  $\sim 20 \mu\text{m}$  for both the bottleneck excitons and the polaritons.

## 6.6 BACKSCATTERING AT HIGH IN-PLANE MOMENTUM

As mentioned in Section 6.5, a significant population of bottleneck excitons can be seen with momentum in the direction back toward the pump spot. Figure 6.6 shows data similar to those in Figure 6.2, but with the sample rotated in the opposite direction and at a greater angle ( $\theta = -40^\circ$ ). This allowed for collection up to  $64^\circ$  external emission angle. Since these data are for negative momenta collected on the positive side of the pump spot, they correspond to backscattered bottleneck excitons traveling back toward the pump. They clearly show that this population is narrowly peaked and not a continuous distribution cut off by the numerical aperture of the optics. A similar population on the opposite side and with opposite momentum is also visible in Figure 6.2(a).

The momentum of this backscattered population increases with increasing distance away from the pump spot. We do not fully understand this process, but we note that backscattered populations have been reported in similar samples under similar experimental conditions [138]. This phenomenon could possibly be explained by the self interference of a population of polaritons with the positive and negative effective diffusive masses corresponding to momenta below and above the lower inflection point, respectively [58]. It is also possible that bright soliton states could explain the observed backscattering. Furthermore, backward flowing emission has been recently observed due to Cherenkov radiation from bright polariton solitons [106]. While the work in this chapter differs in that we used nonresonant excitation, it is possible that bright solitons are still formed since the pumping populates states that support them. Those states are the portion of the lower polariton band with a negative mass [104], which is exactly the part that contains the bottleneck excitons

Another interesting effect can be seen in Figure 6.6: At large enough distance from the pump spot, the  $k_{\parallel}$  of the population brings it near to where the distributed Bragg reflector (DBR) stop band dips cross the lower polariton energy band (see Figure 6.7). Figure 6.6(b) shows where the population begins to couple with the stop band dip, and Figure 6.6(c) shows the position at which it begins to heavily overlap. The photoluminescence intensity is greatly increased due to the new decay channel available to these highly excitonic polaritons, indicating a greatly decreased lifetime. Figure 6.6(d) shows the point where the

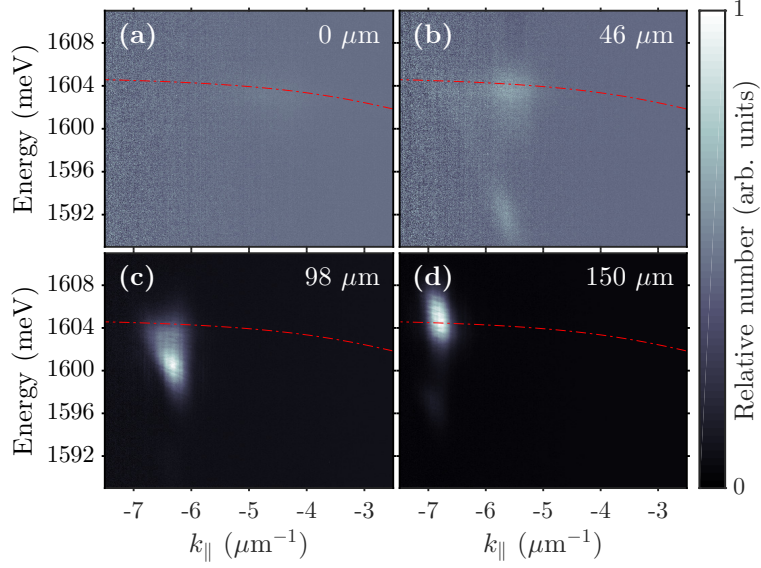


Figure 6.6: The lower polariton distribution as a function of energy and  $k_{\parallel}$  at large angle (compare to Figure 6.2). For these images, the sample was at  $\theta = -40^\circ$  with respect to the imaging objective lens, the pump power was  $\approx P_{\text{th}}/2$ , and the detuning was  $\approx 7$  meV. The positions of the real-space filter with respect to the pump spot for each plot are given in the upper right corners. The red lines show the theoretical lower polariton dispersion. The counts for each image were normalized separately, so the counts of separate images are not comparable.

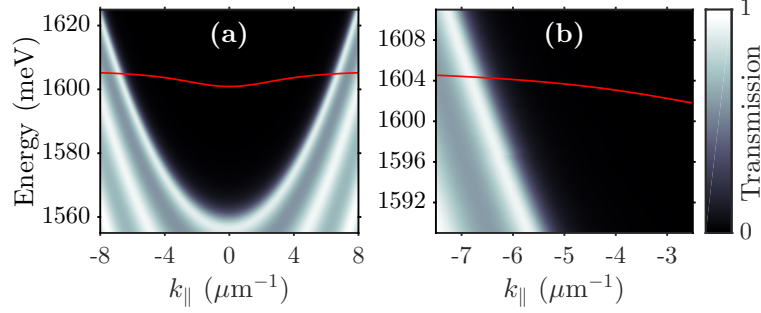


Figure 6.7: The approximate transmission of the top DBR vs. energy and  $k_{\parallel}$ , produced using the transfer-matrix method. The red line shows the theoretical lower polariton dispersion for a detuning of 7 meV, comparable to the red lines in Figure 6.6. (a) and (b) are identical except for the axes ranges.

peak momentum of the population matches the crossing point. The assumption that the lower polariton lifetime is dependent upon the cavity fraction clearly breaks down for this  $k_{\parallel}$  range, beginning around  $5.5 \mu\text{m}^{-1}$ , since another decay channel becomes available apart from regular loss from the cavity.

## 6.7 APPLICATION TO RING PUMP GEOMETRY

The contribution of any lower polariton above the inflection point was unaccounted for in Ref. [85] due to the use of an objective with a numerical aperture of 0.28. In addition, the values for the cavity fraction, based on the detuning, were incorrect (see Appendix E.1). In this section, we will apply the measured exciton distribution above the inflection point to the ring pump geometry in order to help explain the extremely high interaction strength reported in that work.

First, a fairly simple adjustment can be made to the reported exciton-exciton interaction strength  $g_{xx}$  by considering the corrected detuning. The polariton-polariton interaction strength  $g_{pp}$  can be related to the exciton-exciton interaction by the exciton fraction  $f_{\text{exc}}$  of

the polaritons of interest (see Section 1.4.7):

$$g_{\text{pp}} = (f_{\text{exc}})^2 g_{\text{xx}}. \quad (6.7)$$

Additionally, in the mean field approximation, the value for  $g_{\text{pp}}$  can be written as the ratio of the energy shift to the polariton density [85]. The polariton density is found by relating the total number of emitted photons to the lower polariton lifetime, which is in turn related to the cavity fraction  $f_{\text{cav}}$  of the polaritons of interest (see Section 6.5). Putting it all together gives

$$g_{\text{pp}} = \frac{\Delta E}{n_{\text{pol}}} = \frac{\Delta E}{n_{\text{ph}}(\tau_{\text{pol}}/\Delta t)} = \frac{\Delta E \Delta t f_{\text{cav}}}{n_{\text{ph}} \tau_{\text{cav}}} \quad (6.8)$$

where  $\Delta E$  is the energy blue shift,  $n_{\text{pol}}$  is the polariton density,  $n_{\text{ph}}$  is the emitted photon density,  $\Delta t$  is the integration time for photon detection,  $\tau_{\text{pol}}$  is the polariton lifetime, and  $\tau_{\text{cav}}$  is the cavity lifetime. Now we need to consider the effect of a correction to the cavity and exciton fractions. Using Equation 6.8, the ratio of the corrected interaction strength  $g'_{\text{pp}}$  to the original value  $g_{\text{pp}}$  is given by  $g'_{\text{pp}}/g_{\text{pp}} = f'_{\text{cav}}/f_{\text{cav}}$ , since only the cavity fraction is being changed. If we combine this with Equation 6.7, a corrected exciton-exciton interaction strength  $g'_{\text{xx}}$  can be expressed as

$$g'_{\text{xx}} = \frac{g'_{\text{pp}}}{(f'_{\text{exc}})^2} = \frac{g_{\text{pp}}}{(f_{\text{exc}})^2} \frac{f'_{\text{cav}}}{f_{\text{cav}}} = g_{\text{xx}} \left( \frac{f_{\text{exc}}}{f'_{\text{exc}}} \right)^2 \frac{f'_{\text{cav}}}{f_{\text{cav}}}, \quad (6.9)$$

where the primes indicate the corrected values. Applying this to the results in Ref. [85] brings the value for  $g_{\text{xx}}$  down from  $1.74 \text{ meV}\mu\text{m}^2$  to about  $0.54 \text{ meV}\mu\text{m}^2$ .

In order to correct for the excitons above the inflection, we must consider the two separate contributions to the energy blue shift by the polaritons and the bottleneck excitons:

$$\Delta E = \Delta E_{\text{pol}} + \Delta E_{\text{exc}} = g_{\text{pp}} n_{\text{pol}} + g_{\text{xx}} f_{\text{exc}} n_{\text{exc}}. \quad (6.10)$$

The effective exciton density of the bottleneck excitons is given by  $n_{\text{exc}}$ , and is considered to have already been adjusted for their very high exciton fractions. In Ref. [85], the entire exciton density, including that of the bottleneck excitons, was assumed to be negligible. Using this assumption (that  $n_{\text{exc}} = 0$ ), and the relation given in Equation 6.7, we get

$$g_{\text{xx}} = \frac{g_{\text{pp}}}{(f_{\text{exc}})^2} = \frac{\Delta E}{n_{\text{pol}}} \frac{1}{(f_{\text{exc}})^2}, \quad (6.11)$$

where the full energy shift is attributed to the measured polariton density. If instead we keep  $n_{\text{exc}}$  in Equation 6.10, we get the corrected interaction strength:

$$\begin{aligned}
\Delta E &= g_{\text{pp}} n_{\text{pol}} + g'_{\text{xx}} f_{\text{exc}} n_{\text{exc}} \\
&= (f_{\text{exc}})^2 g'_{\text{xx}} n_{\text{pol}} + g'_{\text{xx}} f_{\text{exc}} n_{\text{exc}} \\
&= [(f_{\text{exc}})^2 n_{\text{pol}} + f_{\text{exc}} n_{\text{exc}}] g'_{\text{xx}} \\
\Rightarrow \frac{\Delta E}{n_{\text{pol}} (f_{\text{exc}})^2} &= \left( 1 + \frac{n_{\text{exc}}}{n_{\text{pol}} f_{\text{exc}}} \right) g'_{\text{xx}}
\end{aligned} \tag{6.12}$$

The last line in Equation 6.12 is simply the original  $g_{\text{xx}}$  given in Equation 6.11, which assumed a negligible exciton population. Thus, we can rearrange this for a simple relationship between the original and the corrected interaction strength values:

$$g'_{\text{xx}} = g_{\text{xx}} \left( 1 + \frac{n_{\text{exc}}}{n_{\text{pol}} f_{\text{exc}}} \right)^{-1}. \tag{6.13}$$

In order to find this correction, we only need to know the ratio of bottleneck excitons to polaritons at a given detuning in the ring geometry. To begin, we consider the profile for the exciton population shown in Figure 6.4(b) and a similarly derived profile for the polaritons below the inflection. These profiles are both shown in Figure 6.8(a), with the values adjusted for the approximate average exciton fractions of each population (taken to be 0.92 and 0.74 for the populations above and below the inflection, respectively), such that  $N_{\text{exc}} = n_{\text{exc}}$  and  $N_{\text{pol}} = f_{\text{exc}} n_{\text{pol}}$ .

These profiles were found for what was essentially a point source, and give radial slices of the 2D profiles. The next step is to consider a series of radially symmetric profiles in two dimensions, each centered at even intervals around the perimeter of a circle, producing a ring. For example, imagine a peaked 2D profile with a cross section given by the exciton profile in Figure 6.8(a); place one of these at each of the locations of a dash of the white dashed circle in Figure 6.8(c). The result would be essentially the same as the profile shown in that panel, though the actual result was made using a much closer spacing. The resulting cross sections are given in Figure 6.8(b) and the 2D profiles for both populations are given in Figure 6.8(c-d), all for a ring with a diameter of 40  $\mu\text{m}$ . By integrating a small region in the center for both the polariton and exciton reservoir distributions, we can extract a

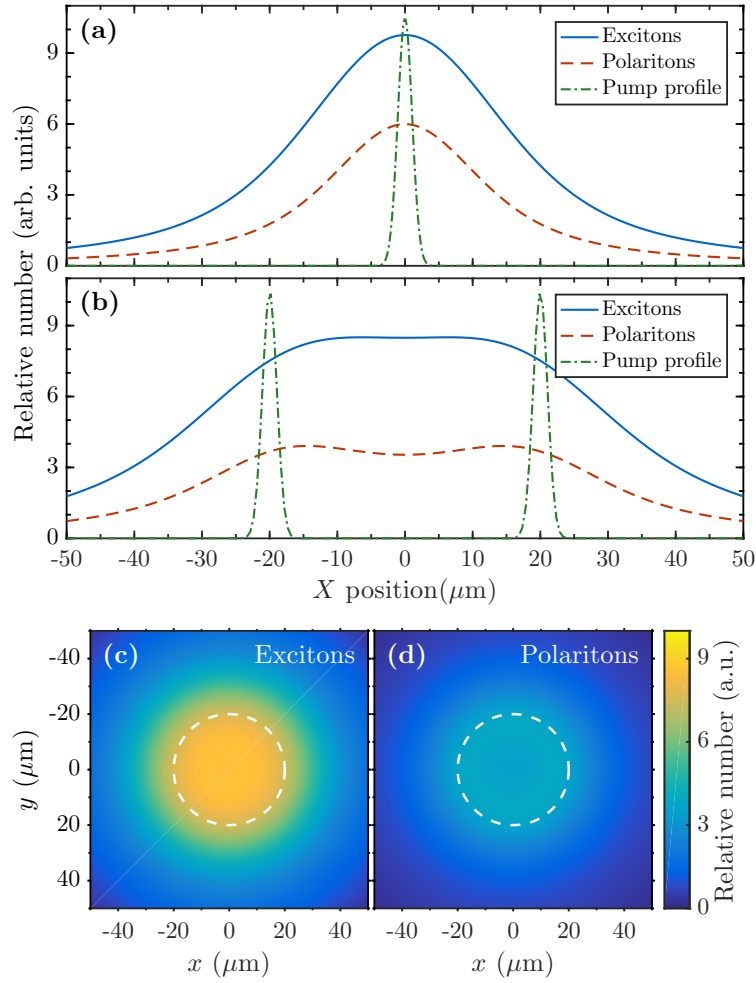


Figure 6.8: (a) The profiles obtained by simulating the same experimental data used in Figure 6.4 for both the bottleneck excitons (above the inflection point) and the polaritons (below the inflection point), along with the pump profile estimated as a Gaussian with  $\text{FWHM} = 2.5 \mu\text{m}$ . (b) A cross section of the resulting ring profile for both the excitons and the polaritons using a  $40 \mu\text{m}$  diameter ring. (c) The exciton profile in two dimensions used to obtain the blue solid curve in (b). (d) The polariton profile in two dimensions used to obtain the red dashed curve in (b). The white dashed circles in both (c) and (d) mark the boundaries of the simulated ring. The relative numbers for all parts of this figure have been adjusted for the approximate exciton fractions, giving effective exciton numbers for each population. [The results shown here were produced by Shouvik Mukherjee.]

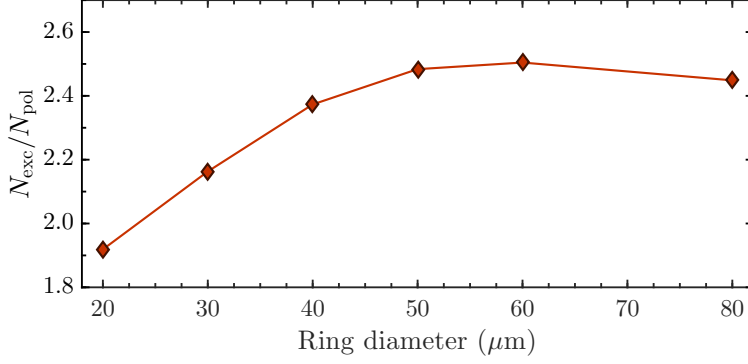


Figure 6.9: The ratio of exciton (above the inflection point) to polariton (below the inflection point) number as a function of ring diameter for distributions similar to and including those shown in Figure 6.8(c-d). The numbers were both adjusted for the exciton fraction, giving a ratio of effective exciton populations. The numbers were obtained by integrating over a circle with  $5 \mu\text{m}$  radius in the center of the ring profiles. [The results shown here were produced by Shouvik Mukerjee.]

ratio of excitons to polaritons for various ring diameters (shown in Figure 6.9). By using the value for a  $40 \mu\text{m}$  diameter ring, similar to that used in Ref. [85], we get the ratio  $n_{\text{exc}}/(n_{\text{pol}}f_{\text{exc}}) = 2.4$ . By combining this correction with the detuning correction, we arrive at a new value for  $g_{\text{xx}}$  of  $159 \mu\text{eV}\mu\text{m}^2$ . While this corrected value is still very high, it is at best an upper bound. In this work, we only carefully considered the exciton population up to  $|k_{\parallel}| \simeq 5.5 \mu\text{m}^{-1}$ . Meanwhile we have also shown evidence of a backscattered population at much higher in-plane momentum. The existence of some excitons at these momenta at such large distances from the pump suggests that a population may be present at momenta greater than any that we probed in this work. In addition, other recent work [139] has shown that the quantum confinement from a ring trap leads to non-negligible blue shifts of the ground state in the center of the trap. Accounting for this effect would further decrease the reported interaction strength.

Overall, the measured exciton transport reported in this work does not fully account for the large interaction strength values reported in Ref. [85]. This is true even when com-

bined with the correction from the more careful calibration of the polariton detuning. More work should be done to carefully consider the confinement effects mentioned above, which, combined with the corrections reported here, should bring the value even closer to other measurements. Ultimately, excitons at even higher momenta, beyond what was measured here and possibly beyond the light cone accessible for direct measurement, could be the main culprit for the continued discrepancy in the measurement of Ref. [85] with theory and other measurements.

## 6.8 CONCLUSIONS

The initial goal of the experiment was certainly accomplished, since we showed that excitonic polaritons usually considered part of the exciton reservoir exhibit much longer transport distances than previously expected. This confirmed the hypothesis that particles that are often undetected in photoluminescence experiments can exist far from the pump spot in significant numbers. These particles are not almost stationary, but rather this portion of the exciton reservoir moves distances comparable to the much lighter and more photonic polaritons. We also showed that there is a significant population of these excitonic polaritons with momentum backscattered toward the original pump spot. Since past work has clearly shown a mostly stationary exciton reservoir, this result indicates three separate categories are needed for particles in the lower polariton mode: normal polaritons with relatively low  $k_{\parallel}$ , stationary reservoir excitons at very high  $k_{\parallel}$ , and highly mobile bottleneck excitons in between.

As demonstrated in Section 6.7, these mobile excitons must be considered when attempting to isolate polaritons from the reservoir. They also affect the potential profile felt by typically observed polaritons at low  $k_{\parallel}$ . This unexpectedly large exciton flow also explains much of the blue shift of the polariton energy seen in Ref. [85], and the presence of many of these bottleneck excitons can explain the large homogeneous line broadening seen in that study. In addition, recent studies [139] have shown that, independently of any exciton flow, the increase of the barrier height in the traps used in Ref. [85] could lead to quantum

confinement effects that can be dominant at low density and photonic detuning.

Long range motion of excitons is not unique to this system, since it has been observed in other experiments through different mechanisms. One case was spatially indirect excitons in coupled quantum wells [140], in which excitons moved over hundreds of microns due to their extraordinarily long lifetime, of the order of tens of microseconds. In these systems, an electric field is applied perpendicularly to the plane of two coupled quantum wells. This causes the holes and electrons to occupy separate quantum wells. The barrier between them reduces the probability that they will recombine, leading to very long lifetimes. Another more recent work [141] showed exciton-polaritons moving hundreds of microns in a surface-wave geometry, in which the optical mode had no  $k_{\parallel} = 0$  state, but instead had a built-in velocity of propagation. In the work presented in this chapter, the long-distance transport appears to be simply the result of longer-than-usual lifetime (due to the high cavity Q of our samples), low disorder, and the high velocity of the bottleneck excitons compared to the thermalized polaritons.

While the hypothesis of the long distance flow of certain highly excitonic lower polaritons (the bottleneck excitons) was demonstrated, future work could be done to determine the details of the observed long-range transport and of the mobile excitons, as well as the origins of the backscattering effect.

## 7.0 FINAL REMARKS

The majority of my work was focused on the effects of injecting a current of free charge carriers into a polariton condensate. Along the way, I developed fabrication methods suitable for the GaAs-based microcavity samples previously developed in my group. This included an etching process to confine the polaritons in square- and wire-shaped pillars, and producing electrical contacts to directly inject electrical current into the quantum wells. The etching process itself produced the interesting effect of trapping the polaritons near the etched edges. I studied the shape and character of this trapping, which we attribute to a strain shift of the bandgap of the GaAs quantum wells, leading to a reduced exciton energy. In addition, the pillar devices showed three distinct thresholds: a lower threshold for formation of a “quasicondensate” at the location of the pump spot, a middle threshold for condensation of polaritons in the local traps along the edges, and an upper threshold when the polaritons form a single monoenergetic condensate.

The primary goal of seeing an effect of polariton condensates on electrical current, or vice versa, was met with three separate observations. The first observation was a superlinear increase in the electrical current that coincided with polariton condensation. We explained this as due to an Auger-like effect by which two excitons in states higher in energy than the condensate scatter off each other. One exciton ionizes to form a free electron-hole pair, adding to the charge carrier density, while the other exciton enters into the polariton condensate. This normally low-probability process is stimulated by the macroscopically occupied condensate, greatly increasing its contribution to the overall carrier density.

The second observation was an increase in the polariton density with the injection of free electrons. The fact that this only happened above the condensate threshold indicated that it, too, was due to a stimulated process. We explain it as an increased likelihood for injected

free electrons to combine with available intrinsic free holes, which is again greatly enhanced in the presence of a condensate.

The third observation was a change in the overall in-plane momentum with the injection of free electrons. This is attributed to a drag of the electrons on the polaritons. While this is not inherently dependent on the presence of a condensate, the extremely narrow range of occupied momentum states that characterize a condensate is helpful in detecting the effect.

I also reported on a separate project, in which we measured the momentum distribution of particles in the lower polariton band as a function of distance from the pump spot and below the condensation threshold. We found that particles with a very high excitonic fraction, much higher than the polaritons at lower momentum states, were able to travel distances comparable to the more photonic polaritons. This is surprising given that they are effectively excitons, and are expected to be mostly confined to the region of the pump spot. We explain this effect as being due to the shape of the polariton energy dispersion, which has regions of negative curvature (allowing the formation of bright solitons) and large slope (leading to high group velocity) even where the excitonic fraction of the polaritons is very large.

There are several obvious directions for future work. In terms of the confinement, it would be interesting to study the strain effects in more detail. This would mostly involve more careful measurement of the energy profile near the edges and modeling the profile to determine the types and the parameters of the strain involved. In particular, this would seek to explain the decrease in energy upon approaching an edge from the unetched bulk region which is then followed by a sharp increase in energy closer to the edge, resulting in a trap several microns from the actual edge.

For work on electrical current injection, it would be worthwhile to produce a truly symmetric wire device. This would require precisely aligning the narrow wire devices perpendicularly to the sample gradient, producing a level energy profile. This would allow observation of the drag effect clearly in both directions in the same device, further confirming the effect and allowing angle-control of the emission in both directions. It would also be interesting to observe the converse effect, with a flow of polaritons inducing a current of electrons. While I did not distinctly observe this, a more careful low noise, low current measurement should be able to do so. Development of similar devices with one  $p$ -type contact would also possibly

allow a fully electrically-pumped device, capable of showing the same nonlinear effects of current flow without the use of a separate optical pump.

As polaritons move toward room temperature [9], the effects reported in this thesis will become more relevant to actual devices. While my work was done at commercially impractical temperatures around 10 K, the concepts are not unique to a particular material, but should be observable in any exciton-polariton system. The practicality is even more probable with the development of electrically-pumped polaritons, which have already been demonstrated [27, 28]. With recent advances, it is possible to envision useful polariton-based devices with highly nonlinear electrical characteristics, or polariton lasers implementing the polariton-drag effect to achieve electrically tunable emission angles. The long-range exciton transport due to the polaritons could have applications in solar cell design, where one of the limiting constraints is the distance excitons can flow before they recombine [142]. The exciton-polariton system is full of exciting potential for the future, both as a fascinating physical system that demonstrates macroscopic quantum effects, and an accessible solid state system in which those effects could be harnessed for practical optical and electronic devices.

## APPENDIX A

### CAVITY REFRACTIVE INDEX ANGLE DEPENDENCE

Since the cavity is made of several different materials, each of which have a different index of refraction, the effective overall refractive index of the cavity will change with angle. To begin, I first consider the condition that gives a cavity mode in the first place. This is given in Equation 1.3, which for a single pass of the cavity is

$$\Delta\phi = \ell k_{\text{vac}} n_{\text{cav}} \cos \theta_{\text{cav}} = m\pi, \quad m = 1, 2, 3, \dots \quad (\text{A.1})$$

This equation states that, in order to have a cavity mode, the total phase change of the wave across the cavity in a single pass ( $\Delta\phi$ ) must be an integer ( $m$ ) multiple of  $\pi$ . This allows for constructive interference within the cavity. The vacuum wavenumber  $k_{\text{vac}}$  at which this occurs is dependent on the material index of refraction  $n_{\text{cav}}$ , the material thickness  $\ell$ , and the angle of travel within the cavity  $\theta_{\text{cav}}$ . The condition for a cavity with multiple materials is similar, but requires the *total* phase difference across all of the layers to fulfill this condition. This gives a sum relationship, which can be equated to an overall *effective* cavity (now denoted by the subscript “cav”):

$$\begin{aligned} \Delta\phi &= \sum_i k_{\text{vac}} n_i \ell_i \cos \theta_i = k_{\text{vac}} n_{\text{cav}} \ell \cos \theta_{\text{cav}} = m\pi \\ &\Rightarrow \sum_i n_i \ell_i \cos \theta_i = n_{\text{cav}} \ell \cos \theta_{\text{cav}}, \end{aligned} \quad (\text{A.2})$$

where each  $i$  in the sum corresponds to a different layer within the cavity, each with thickness  $\ell_i$ . I can use Snell's law to relate every angle to the external emission angle  $\theta_{\text{ext}}$ , which allows us to rewrite the above relation as

$$\Rightarrow \sum_i n_i \ell_i \cos \left( \sin^{-1} \left( \frac{n_{\text{ext}} \sin \theta_{\text{ext}}}{n_i} \right) \right) = n_{\text{cav}} \ell \cos \left( \sin^{-1} \left( \frac{n_{\text{ext}} \sin \theta_{\text{ext}}}{n_{\text{cav}}} \right) \right) \quad (\text{A.3})$$

Since this only gives us an equation for the product of  $n_{\text{cav}}$  and  $\ell$ , I am free to choose any value for the effective cavity thickness  $\ell$ . The most reasonable choice is the sum of the thicknesses of the individual layers, giving  $\ell = \sum_i \ell_i$ , which is the total actual cavity thickness. I now have an equation giving the effective cavity refractive index in terms of the individual cavity layer properties, the external refractive index, and the external emission angle. However, the equation is transcendental, and it also requires exact knowledge of each layer's refractive index and their actual (or at least relative) thicknesses. These values are very difficult to pin down with precision. However, they are known to high enough precision to estimate whether this dependence on external angle is even worth considering, or if it can be simply ignored.

To do this, I simply used approximate values for each of the materials' refractive indices, as well as the thicknesses of each layer given by the MBE growth process, for the samples described in Section 2.1. Due to the sample wedge, these thicknesses vary by several percent across the sample, but only their relative thicknesses actually affect Equation A.3. These relative thicknesses are known to within at least 5%. The materials within the cavity are pure GaAs, pure AlAs, and the alloy,  $\text{Al}_{0.2}\text{Ga}_{0.8}\text{As}$ , for which I used 3.5, 3.0, and 3.4 as refractive indices, respectively. I also used  $n_{\text{ext}} = 1$ .

By using all of the layers of the structure between the approximate locations of the electric field nodes (given by transfer matrix simulations), I solved Equation A.3 numerically and plotted the result in Figure A.1. This shows that, for material parameters approximately the same as those for the actual sample, a change of  $< 0.03\%$  occurs from normal emission to emission at the maximum external angle. This shows that I can safely ignore the angle dependence of the cavity refractive index. In general, this should be true for any set of materials with relatively low index contrast.

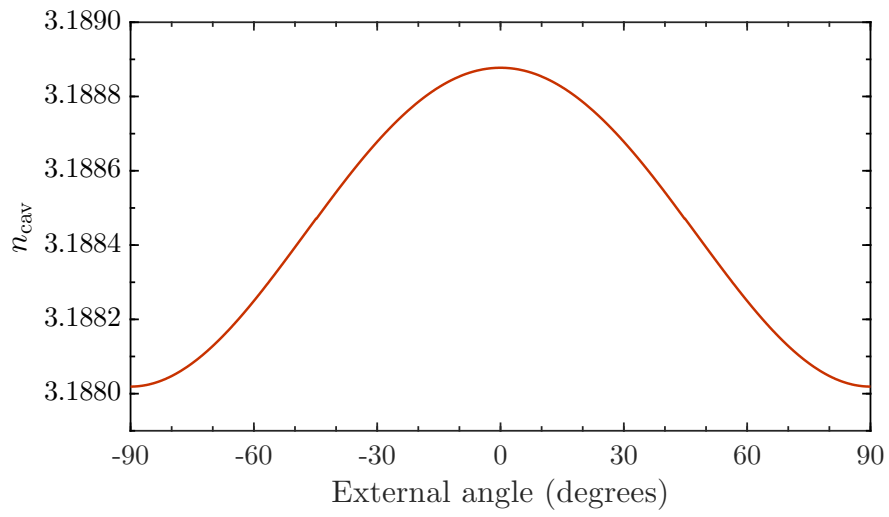


Figure A.1: Effective cavity index of refraction as a function of external emission angle. This was produced using the reported layer thicknesses from the MBE growth process, and refractive indices of 3.5, 3.0, and 3.4 for the GaAs, AlAs, and  $\text{Al}_{0.2}\text{Ga}_{0.8}\text{As}$  layers, respectively.

## APPENDIX B

### THE TRANSFER MATRIX METHOD

In this appendix, I will introduce the transfer matrix method and show how it can be used to simulate real structures, providing the reflectivity and transmission coefficients as functions of both incident angle and frequency. For a full discussion, see [39] or [40], which also cover much of the analysis and approximations based on infinite stacks. I mostly follow Ref. [39], but I will limit this discussion to deriving the matrix for a single layer, which allows numerical simulations of complex structures with no further work.

The natural starting point is with an electromagnetic wave traveling through space, which encounters a dielectric material interface. I will assume the interface is perfectly flat and in the  $x$ - $y$  plane, and the wave is traveling in the  $y$ - $z$  plane (and therefore has no  $x$ -dependence). In general, the wave can be in either the transverse electric (TE) or transverse magnetic (TM) polarization arrangements. In this case, TE polarization would then correspond to an electric field amplitude entirely in the  $x$ -direction, while TM polarization has a magnetic field amplitude entirely in the  $x$ -direction. For the sake of clarity, I will solve both arrangements. This leaves us with two equations, one for each polarization, for the fields on the incident side of the interface:

$$\text{TE : } \mathbf{E}(\mathbf{r}, t) = E_x \hat{\mathbf{i}} = (E^+ e^{i(k_z z + k_y y - \omega t)} + E^- e^{i(-k_z z + k_y y - \omega t)}) \hat{\mathbf{i}} \quad (\text{B.1a})$$

$$\text{TM : } \mathbf{B}(\mathbf{r}, t) = B_x \hat{\mathbf{i}} = (B^+ e^{i(k_z z + k_y y - \omega t)} + B^- e^{i(-k_z z + k_y y - \omega t)}) \hat{\mathbf{i}}. \quad (\text{B.1b})$$

In these equations,  $\mathbf{E}$  is the electric field and  $\mathbf{B}$  is the magnetic field, and  $\mathbf{r} = x\hat{\mathbf{i}} + y\hat{\mathbf{j}} + z\hat{\mathbf{k}}$  is the spatial position vector. I have included waves propagating both to and from the interface, and momentum in the  $y$ -direction is conserved since it is tangential to the interface. In both cases Faraday's Law  $\nabla \times \mathbf{E} = -\partial_t \mathbf{B}$  can be invoked to find the other field amplitude. In the TE case, this is fairly straightforward, giving

$$\begin{aligned} i\omega \mathbf{B} &= \partial_z E_x \hat{\mathbf{j}} - \partial_y E_x \hat{\mathbf{k}} \\ &= [ik_z(E^+ e^{ik_z z} - E^- e^{-ik_z z})\hat{\mathbf{j}} - ik_y(E^+ e^{ik_z z} + E^- e^{-ik_z z})\hat{\mathbf{k}}] e^{i(\omega t - k_y y)}. \end{aligned} \quad (\text{B.2})$$

The boundary conditions for Maxwell's equations demand only that the tangential portions of both fields are continuous. It is obvious from Equation B.2 that these boundary conditions are equivalent to demanding continuity in both the electric field amplitude and its spatial derivative along the propagation direction ( $\partial_z E_x$ ). This gives us two equations that satisfy the requirements of continuity for the in-plane parts of the electric and magnetic field amplitudes (dropping  $t$ -dependence) on both sides of the interface. Using the relationship  $\omega n/c = |\mathbf{k}|$  where  $n$  is the refractive index of the material, recognizing that  $k_z/|\mathbf{k}| = \cos \phi$  where  $\phi$  is the angle of incidence, and ignoring non-tangential components, the two equations in the general case are

$$\begin{aligned} \mathcal{E}_a^+ + \mathcal{E}_a^- &= \mathcal{E}_b^+ + \mathcal{E}_b^- \\ (\mathcal{E}_a^+ - \mathcal{E}_a^-)n_a \cos \phi_a &= (\mathcal{E}_b^+ - \mathcal{E}_b^-)n_b \cos \phi_b, \end{aligned} \quad (\text{B.3})$$

where

$$\mathcal{E}_q^\pm = E_q^\pm e^{\pm i k_z q z} = E_q^\pm e^{\pm i |\mathbf{k}_q| \cos \phi_q z} = E_q^\pm e^{\pm i (\omega n_q/c) \cos \phi_q z}, \quad (\text{B.4})$$

and where  $a$  and  $b$  denote the points just to the left and right of the interface, respectively. All of the propagation angles (the  $\phi_i$ 's) are of course related by Snell's law, so they can all be found from the initial incident angle and the index of refraction of each material. This

interface crossing can conveniently be solved in matrix form. Using  $\Psi_q = (\mathcal{E}_q^+, \mathcal{E}_q^-)^T$ , this gives

$$\begin{aligned} \mathbf{A}_a \Psi_a &= \mathbf{A}_b \Psi_b \Rightarrow \mathbf{A}_q = \begin{pmatrix} 1 & 1 \\ n_q \cos \phi_q & -n_q \cos \phi_q \end{pmatrix}, \quad q = a, b \\ \Rightarrow \mathbf{A}_b^{-1} \mathbf{A}_a \Psi_a &= \Psi_b \\ \Rightarrow \mathbf{A}_b^{-1} \mathbf{A}_a &= \mathbf{M} = \frac{1}{2} \begin{pmatrix} 1 + \frac{n_a \cos \phi_a}{n_b \cos \phi_b} & 1 - \frac{n_a \cos \phi_a}{n_b \cos \phi_b} \\ 1 - \frac{n_a \cos \phi_a}{n_b \cos \phi_b} & 1 + \frac{n_a \cos \phi_a}{n_b \cos \phi_b} \end{pmatrix}, \end{aligned} \quad (\text{B.5})$$

where  $\mathbf{M}$  now relates the field amplitudes on either side of the interface. One more step is necessary for a DBR, since a DBR is not simply a series of interfaces, but also contains space within dielectric material between each interface. Since we have already expressed each field in terms of waves traveling in both directions, and there are, by definition, no material changes within a single layer, the individual peak amplitudes  $E_q^\pm$  do not change. We therefore only need to propagate the waves through a distance  $d$ , where  $d$  is the thickness of the material layer. This is equivalent to a change of  $z \rightarrow z + d$ , which, to move across the layer that was entered via  $\mathbf{M}$ , is accomplished by

$$\mathbf{P} \begin{pmatrix} \mathcal{E}_b^+ \\ \mathcal{E}_b^- \end{pmatrix} = \begin{pmatrix} \mathcal{E}_b^+ e^{+i(\omega n_b/c) \cos \phi_b d} \\ \mathcal{E}_b^- e^{-i(\omega n_b/c) \cos \phi_b d} \end{pmatrix} \Rightarrow \mathbf{P} = \begin{pmatrix} e^{i(\omega n_b/c) \cos \phi_b d} & 0 \\ 0 & e^{-i(\omega n_b/c) \cos \phi_b d} \end{pmatrix}. \quad (\text{B.6})$$

The combined transfer matrix that takes the field amplitudes through an interface and then through the full layer on the other side is a combination of  $\mathbf{M}$  and  $\mathbf{P}$ , which is the transfer matrix  $\mathbf{T}$ . The amplitude vector  $\Psi_{j-1}$ , which is just to the left of the interface between layer  $j$  and layer  $j-1$  is then related to  $\Psi_j$  by

$$\mathbf{P}_j \mathbf{M}_j \Psi_{j-1} = \frac{1}{2} \begin{pmatrix} (1 + \chi_j) e^{i\theta_j d_j} & (1 - \chi_j) e^{i\theta_j d_j} \\ (1 - \chi_j) e^{-i\theta_j d_j} & (1 + \chi_j) e^{-i\theta_j d_j} \end{pmatrix} \Psi_{j-1} = \mathbf{T}_j \Psi_{j-1} = \Psi_j \quad (\text{B.7})$$

where  $\theta_j = (\omega n_j/c) \cos \phi_j$  and  $\chi_j = (n_{j-1} \cos \phi_{j-1}) / (n_j \cos \phi_j)$ .  $T$  can therefore be thought of as an operator that takes the field amplitude vector  $\Psi_{j-1}$  through the layer  $j$ . If the stack is made of  $N$  layers, then total propagation through the full stack is given by

$$\mathbf{T}_f (\mathbf{T}_N \mathbf{T}_{N-1} \dots \mathbf{T}_2 \mathbf{T}_1) \Psi_0 = \mathbf{T}_T \Psi_0 = \Psi_f. \quad (\text{B.8})$$

Since each  $\mathbf{T}_j$  matrix moves through an interface and the layer beyond it, the final transfer matrix  $\mathbf{T}_f$  is necessary to exit into the space beyond the  $N$ th layer. Simply setting  $d_f = 0$  will then give  $\Psi_f$  at just the other side of the final interface. This equation now allows us to relate the field amplitudes on either side of a stack of dielectric layers. Since light is incident from only one side, the backward propagating wave amplitude after the final interface will be zero ( $\mathcal{E}_f^+ = 0$ ). This gives a final relationship for the amplitudes:

$$\mathbf{T}_T \begin{pmatrix} \mathcal{E}_0^+ \\ \mathcal{E}_0^- \end{pmatrix} = \begin{pmatrix} \mathcal{E}_f^+ \\ 0 \end{pmatrix} \Rightarrow \mathbf{T}_T \begin{pmatrix} E_0^+ \\ E_0^- \end{pmatrix} = \begin{pmatrix} E_f^+ e^{i(\omega n_f/c) \cos \phi_f d_T} \\ 0 \end{pmatrix}, \quad (\text{B.9})$$

where I have conveniently chosen the first interface to be at  $z = 0$ . Since I am only concerned with the reflected ( $\mathcal{E}_0^-$ ) and transmitted ( $\mathcal{E}_f^+$ ) amplitudes relative to the incoming amplitude ( $\mathcal{E}_0^+$ ), Equation B.9 becomes

$$\mathbf{T}_T \begin{pmatrix} 1 \\ r \end{pmatrix} = \begin{pmatrix} t \\ 0 \end{pmatrix} \quad (\text{B.10})$$

where  $r = \mathcal{E}_0^-/\mathcal{E}_0^+$  and  $t = \mathcal{E}_f^+/\mathcal{E}_0^+$ . The transfer matrix elements can then be directly related to the coefficients  $r$  and  $t$  by Equation B.10, giving

$$r = -\frac{\mathbf{T}_{21}}{\mathbf{T}_{22}}, \quad t = \frac{\mathbf{T}_{11}\mathbf{T}_{22} - \mathbf{T}_{12}\mathbf{T}_{21}}{\mathbf{T}_{22}} = \frac{\det(\mathbf{T}_T)}{\mathbf{T}_{22}}, \quad (\text{B.11})$$

and these coefficients are related to the reflectance ( $R$ ) and transmittance ( $T$ ) by the simple relationship

$$R = |r|^2, \quad T = \frac{n_f}{n_0} |t|^2, \quad (\text{B.12})$$

where  $R + T = 1$ . This is the final result that I needed for the TE polarization. For the TM polarization, I will show that the same result can be used with a simple transformation,  $n \rightarrow 1/n$ . Returning to Equation B.1b and considering the proposed  $\mathbf{B}$ -field solution for the TM wave, I will again apply Faraday's Law. Since  $\mathbf{B}$  is entirely in the  $x$ -direction,  $\mathbf{E}$  will have components in the  $y$ - and  $z$ -directions, giving  $\mathbf{E} = E_y \hat{\mathbf{j}} + E_z \hat{\mathbf{k}}$ . Since there is also no dependence on  $x$ , Faraday's Law gives

$$i\omega B_x \hat{\mathbf{i}} = (\partial_y E_z - \partial_z E_y) \hat{\mathbf{i}} \quad (\text{B.13})$$

Starting with the most general form of  $\mathbf{E}$ ,

$$\mathbf{E} = (\alpha e^{ik_z z} + \beta e^{-ik_z z}) e^{ik_y y} \hat{\mathbf{j}} + (\gamma e^{ik_z z} + \delta e^{-ik_z z}) e^{ik_y y} \hat{\mathbf{k}}, \quad (\text{B.14})$$

Equation B.13 becomes

$$\omega(B^+ e^{ik_z z} + B^- e^{-ik_z z}) = k_y(\gamma e^{ik_z z} + \delta e^{-ik_z z}) - k_z(\alpha e^{ik_z z} - \beta e^{-ik_z z}) \quad (\text{B.15})$$

The direction of  $\mathbf{E}$  is always determined by both the propagation angle and the direction of  $\mathbf{B}$ , on both the incoming and outgoing parts. This gives the relations  $-\gamma/\alpha = \delta/\beta = \tan \phi = k_y/k_z$ , where the negative sign is for the fact that  $E_y$ , and therefore  $\alpha$ , is negative when  $\mathbf{B}$  is in the  $+x$ -direction for the incoming wave. Using these relations and matching right- and left-moving terms,

$$\mathbf{E} = \frac{c}{n} [(-B^+ e^{ik_z z} + B^- e^{-ik_z z}) \cos \phi \hat{\mathbf{j}} + (B^+ e^{ik_z z} + B^- e^{-ik_z z}) \sin \phi \hat{\mathbf{k}}] e^{ik_y y}. \quad (\text{B.16})$$

Once again, the boundary conditions only call for matching the tangential components through an interface, which gives

$$\begin{aligned} \mathcal{B}_a^+ + \mathcal{B}_a^- &= \mathcal{B}_b^+ + \mathcal{B}_b^- \\ (\mathcal{B}_a^+ - \mathcal{B}_a^-) \frac{\cos \phi_a}{n_a} &= (\mathcal{B}_b^+ - \mathcal{B}_b^-) \frac{\cos \phi_b}{n_b}, \end{aligned} \quad (\text{B.17})$$

where, similarly to  $\mathcal{E}_q^\pm$ ,

$$\mathcal{B}_q^\pm = B_q^\pm e^{\pm i k_{z,q} z} = B_q^\pm e^{\pm i |\mathbf{k}_q| \cos \phi_q z} = B_q^\pm e^{\pm i (\omega n_q / c) \cos \phi_q z}. \quad (\text{B.18})$$

The Equations B.17 are obviously very similar to those given in Equation B.3, with only the simple change  $n \rightarrow 1/n$ . Since the reflection and transmission coefficients can also be expressed in terms of the magnetic field amplitude, this pair of equations can be solved at each interface in exactly the same way as for TE polarization. The only necessary change is to make the transformation  $n \rightarrow 1/n$  for the interface matrix  $\mathbf{M}$  (but not for the propagation matrix  $\mathbf{P}$ ). Thus, the reflection and transmission coefficients for any stack of materials can be simulated by building a transfer matrix, which can be done with fairly simple programming even for complicated sets of layers.

## APPENDIX C

### FABRICATION DETAILS

In this appendix, I will provide more of the details for the fabrication described in Section 2.2. This will include the specific parameters, which should be a good starting point if trying to reproduce these results. However, keep in mind that every system is different, even when exactly the same on paper, so do not expect exactly the same results in a different system with the same parameters.

The original 2 inch diameter wafers of sample are cleaved before any fabrication into chips with width and length between 4 and 7 mm, and usually about 5 mm square. The particular part of the sample to use is determined by initial characterization to find the desired detuning and measuring the orientation of the sample thickness gradient. These small chips are then processed as described in the following sections.

#### C.1 LITHOGRAPHY FOR ETCHING

The lithography for all the etches described in this thesis was generally the same. I started with a cleaning process (adapted from Ref. [101]) of sonicating for 10 minutes in acetone, then 5 minutes in methanol, and then rinsing with deionized (DI) water. The methanol can be followed or replaced with isopropanol as preferred, again following with DI water. Each chip is then dried by first blowing off the majority of the water with pressurized N<sub>2</sub> and baking on a hotplate for 5-10 minutes at 150°C. The next step is to spin coat photoresist

onto each sample chip. Because the chips are very small, they must be either mounted onto a carrier wafer or carefully placed on a small sample chuck within the spinner. I prefer to avoid using carrier wafers, but if that is desired then simply drop a very small amount of photoresist (usually whatever is used later) onto a large wafer and gently press the now-clean sample into the drop, attempting to get it as level as possible. Baking for a few minutes at about 115°C should satisfactorily set the photoresist. Avoid higher temperatures because it will make removal from the carrier wafer more difficult.

Specifically, I use AZ P4210 photoresist for the etch mask, and a spin speed of 4000 rpm for 60 seconds. This gives a resist thickness of about 2700 nm for our small chips. A thinner resist layer can be achieved by using the similar AZ P4110, which is simply a less viscous mixture. The thickness can also be fine-tuned by adjusting the spin-speed, with higher speeds giving a thinner result. However, due to the small chip size, speeds below about 4000 rpm were avoided to help keep the edge bead width to a minimum. After spinning, the sample is soft-baked at 95°C for 2 minutes. This dries out the resist, slightly solidifying it and preventing it from sticking to a lithography mask or smudging during contact lithography.

The exposure was carried out in two different ways. I initially used a Suss MJB-3 mask aligner and a lithography mask made from chromium on quartz glass. The mask was placed in contact with the sample during exposure, which was done using 405 nm radiation at 25 mW/cm<sup>2</sup> for 6 seconds, giving a total exposure energy density of 150 mJ/cm<sup>2</sup>. Later exposure using the same overall fabrication process was done by colleagues in my group with a Heidelberg MLA100 mask-less aligner. This was done using 365 nm radiation but also giving a total exposure energy density of 150 mJ/cm<sup>2</sup>.

Following exposure, the samples are developed in a solution of AZ400K and DI water with a developer:water ratio of 1:3 by volume. A development time of 45 seconds usually gives nice results. They are then hard-baked at 115°C for about 30 minutes (though 10 minutes appears to give similar results) in order to harden the photoresist and increase its resistance to the etch process. Higher temperatures should be avoided because they will make removal of the mask difficult after etching. A short O<sub>2</sub> plasma etch is used after hard-baking to descum the unmasked surface of the sample. This is done in a Trion Phantom III LT RIE, set to 500 mT chamber pressure and 50 mW RF power, for 10 seconds. This essentially eats

away a small amount of photoresist everywhere, which has a negligible effect on the thick layer but removes residue left over after development.

## C.2 REACTIVE ION ETCHING

Depending on preference or need, the sample could have been removed from the carrier wafer at any point in the preceeding steps. Before the etch, however, the samples are mounted to a 4 inch Si carrier wafer using Cool Grease to give good thermal contact. They are then etched in a Plasma Therm Versaline ICP RIE system, set to a chamber pressure of 3.0 mT, 600 W ICP power, 75 W RF bias power, and 20:7  $\text{BCl}_3/\text{Cl}_2$  gas mixture ratio. The etch rate was usually measured to be about 900 nm/minute, but was somewhat inconsistent (I measured it as low as 730 nm/minute at one point) for unknown reasons (possibly due to a tendency of clogging of the  $\text{BCl}_3$  line, changing the gas mixture). For the etches that are meant to expose the quantum wells, the etch depth must be controlled to within about 100 nm in order to be close to the quantum well layers without removing them. This was difficult because not only was the etch rate inconsistent, but the thickness gradient induces a change in the desired etch depth across the sample. Rather than carefully time this step, we actually crudely measured the depth *in-situ*. Since the DBR is a periodic structure, the reflectivity changes periodically as the layers are removed. This change is actually visible to the naked eye. Thanks to a well-situated window in the etching chamber directly over the sample, we are able to watch the sample color change periodically as the layers are removed at a rate of about 10 seconds per period. By counting the number of periods, we can directly determine the etch depth to within a DBR layer period while the etch takes place. Thus, the desired number of periods can be removed by either manually stopping the etch, or incrementally etching for short time periods.

For the “mesa” etch which removed the quantum wells between devices, or basic pillars that do not need intact quantum wells nearby, an overlong timed etch is sufficient.

Following etching, the samples are removed from the carrier wafer and the Cool Grease is mostly removed by carefully rinsing with acetone. It is important during this step to not

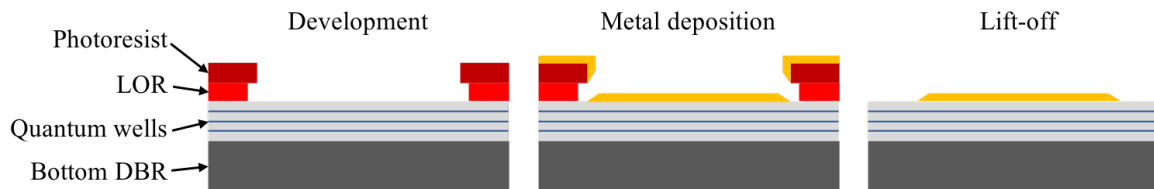


Figure C.1: The function of lift-off resist during metal deposition. A bi-layer stack of photoresist on top of lift-off resist (LOR) is formed via lithography, with undercutting in the LOR layer after development. During metal deposition, the metal thin films are somewhat able to follow the contours of the surface, but form clean breaks at the photoresist edge due to the overhang. During liftoff, solvent is able to easily penetrate at the breaks in the film, cleanly lifting off the masking layers without damaging the film on the sample surface.

get the Cool Grease residue on the top surface, as it can be very difficult to remove when stuck in the grooves of the etched pattern. I usually remove it by a combination of drawing the bottom of each sample across an acetone-soaked wipe placed on a surface, and gently spraying the sample with acetone while holding it aloft. After this, the photoresist mask is removed by sonicating in Microposit Remover 1165 for 60 seconds, followed by soaking in a different container of the same solution for 4 minutes. The sample is then immediately rinsed in Acetone then Methonal (or Isopropanol) then DI water to avoid having any residue settle on the surface, and then dried by blowing with compressed  $N_2$ .

### C.3 LITHOGRAPHY FOR CONTACT DEPOSITION

The lithography process for depositing the metal contacts was very similar to the one for etching, but with some important differences. The same cleaning process is used, but this is followed by coating with MicroChem LOR5B lift-off resist. This is spun at 3000 rpm for 60 seconds, followed by pre-baking at  $190^\circ\text{C}$  on a hotplate for 4.5 minutes. This layer provides undercutting beneath the photoresist, and these parameters are chosen to give

the desired undercut amount and undercut layer thickness (see Figure C.1). The layer thickness is established the same way as for the regular photoresist: the viscosity of the solution and the spin speed. The thickness should be chosen to be at least as great as the deposited metal thickness, to allow for a clean break in the metal film at the overhang of the photoresist. In this case, LOR5B at 3000 rpm results in a thickness of about 500 nm, which is significantly greater than the metal contact thickness of about 300 nm. The undercut amount is dependent on the type of LOR, the development solution, the development time, the pre-bake time, and the pre-bake temperature. Since the development parameters are typically dependent on the choice of photoresist, the others can be chosen independently to get the desired undercut amount based on how they affect the development rate. For a given type of LOR, the development rate is strongly affected by the bake temperature and weakly affected by the bake time, so these can be used as coarse and fine adjustments, respectively. The desired undercut amount is set by the other parameters: an undercut much greater than the capping photoresist layer thickness can cause sagging, while an amount much less than the LOR layer thickness could result in some overlap of the metal with the LOR. The parameters I chose were based on the LOR options that were readily available and intended to give an undercutting of about 2-3  $\mu\text{m}$ . This is enough to be easily visible in a regular optical microscope as an apparent border along the patterned photoresist edge after development, which allows confirmation of the success of the process, while also being small enough to not cause sagging of the photoresist layer.

Following the pre-bake of the LOR and at least 2 minutes of cooling, AZ P4110 photoresist is applied in the same process described for the etch mask (which was done using P4210 resist). The development is done using a solution of AZ400K and DI water with a develop:water ratio of 1:4 for two minutes. The undercut amount can be checked with a microscope as mentioned above. If more undercut is desired, further developing for up to an additional 15 seconds can be attempted; more than this tends to overdevelop the photoresist and degrade the pattern. Typically, even barely visible undercut amounts ( $\lesssim 1 \mu\text{m}$ ) give adequate lift-off results.

After developing, the samples are hard-baked at 115°C for 4 minutes. The longer times used for the etch lithography are unnecessary since the photoresist is not used as an etch

mask in this case. The same descum process is used as for the etch lithography, and is even more crucial to achieve clean contact between the metal and sample surface.

## C.4 METAL CONTACT DEPOSITION

Following the lithography for the metal deposition, the sample surface must be stripped of oxide build-up. This is done using a 10:1 buffered oxide etch (BOE) for 5 seconds, followed by two separate rinses (two different beakers of water) for 10 seconds. The samples are then dried by carefully and intensively blowing with dry  $N_2$  gas, and are then immediately mounted and placed in the evaporation chamber. Generally speaking, the total time between dipping in the BOE and beginning the chamber evacuation process should be less than 15 minutes. The less time the better to avoid re-formation of an oxide layer. The metals are deposited by evaporating in a Thermionics Laboratory VE180 e-beam evaporator, pumped to a maximum starting pressure of  $5 \times 10^{-6}$  mT. A 5 nm thick layer of Ni is first deposited, followed by 120 nm of  $Au_{0.88}Ge_{0.12}$  alloy, then 25 nm of Ni, and finally 150 nm of Au.

After removing the samples from the chamber, the resist is lifted off by sonicating in MicroChem Remover PG for 60 seconds, followed by soaking in a different container of the same solution for 10 minutes. Remover PG is necessary in this case because of the LOR, which is not particularly soluble in acetone. The samples are then immediately rinsed in acetone then methonal (or isopropanol) then DI water to avoid having any residue, especially metal flakes, settle on the surface. They are then blow-dried with compressed  $N_2$ . If, after checking under a microscope, the metal film is not fully lifted off in the areas on top of the resist, then the above steps can be repeated.

Annealing is typically carried out in a gas mixture of 95%  $N_2$  and 5%  $H_2$ , in a chamber pressurized to 1.0 bar. In this environment, the samples are heated to about 320°C for about 60 seconds. The temperature is then quickly ramped to 410°C and held there for 30 seconds before being allowed to cool.

The samples are then mounted into leadless chip carriers (LCCs), using silver paint as the adhesive. The contacts of individual devices are then connected to the contacts of the

LCC by wire-bonding with aluminum wires. These LCCs can then be placed into sockets in the cryostat apparatus, allowing connections to the devices during an optical experiment at cryogenic temperatures.

## APPENDIX D

### OPTICAL SETUPS

In this appendix, I will provide many of the details of the optical setups I used for actual experiments. For ease of reference, I will include the actual details in the captions of the figures, but here I will outline some of the main concerns that went into designing these. One thing that should be noted is that a beam splitter is usually necessary in a setup even when it is not used for gathering data in order to aid in optics alignment. While it can be carefully removed, it constrains the setup due to the necessity of having space for it.

The first step is usually choosing the appropriate objective lens (labeled  $L_1$  in the diagrams). This is usually initially constrained by the required working distance. For a regular microscope with a sample placed on a stage in air, the working distance can be extremely small, often only hundreds of microns, allowing very high numerical aperture objective lenses. In low temperature experiments this is usually not possible due to the windows, chamber, vacuum jacket and other hardware between the sample and the first available location for a lens. This distance must be further increased if the pump path must be anything other than through the objective. A “microscope” cryostat with the sample very close to the outside window often helps with these problems, but still has some hardware between the sample and the objective. Whatever the constraints, once a minimum working distance is established, the best option is usually whichever lens gives the largest numerical aperture with the lowest amount of optical aberration.

The next step is often choosing a pumping path, though this sometimes must be done before the previous step. There are essentially two options: pumping through the objective

or pumping from a different path. The advantage of pumping through the objective is that it is usually the option that gives the smallest possible spot size (and it does not interfere with choosing the objective in the first place) because of the typically short focal length and low aberration of the objective. It also gives the ability to tweak the incident pumping angle by changing where the beam intersects the backside focal plane ( $k$ -space plane) of the objective. The downsides are that it offers limited ability to move the pump with respect to the imaging location, and, unless a dichroic beam splitter is used, typically loses much of the pump power through transmission at the beam splitter. Pumping past the objective, either at large angle or with a small pump mirror in front of the objective, has the advantages of low loss, independence of pump location from imaging location, and the ability to pump at very high angle. Its disadvantages are usually larger spot size because of larger minimum focal length of the final focusing lens, increased working distance requirement on the objective, and difficulty in tweaking the pumping angle during an experiment.

Once the above choices are made, the next steps are usually to determine the desired magnification, in both real-space and  $k$ -space, and then choose optics that accomplish this without “clipping” either space along the way (avoiding clipping is actually a fairly complicated process in its own right, and is approached with ray tracing, as in e.g. Ref. [77]). It is also usually necessary to have both  $k$ -space and real-space filtering, so accessible planes of both spaces must be formed before the optics that are switched back and forth (otherwise stability and repeatability become a problem). There often must also be space for many other optics, such as polarizers, waveplates, pick-offs for secondary imaging, ND-filtering, wavelength filtering, etc. I will only specifically discuss the placement of Dove prisms, since they are so large. Overall, this process is usually very specific to the optics and space that are available, so I will simply describe my choices without providing the full rationale behind them.

For reference, the setup in Figure D.1 was the one generally used in the work described in Chapter 3 and the square-pillar parts of Chapter 4. The setup in Figure D.2 was used for the work in Chapter 5 and the work on wire pillars in chapter 4. The setup in Figure D.3, with pumping through the objective, applies to Chapter 6.

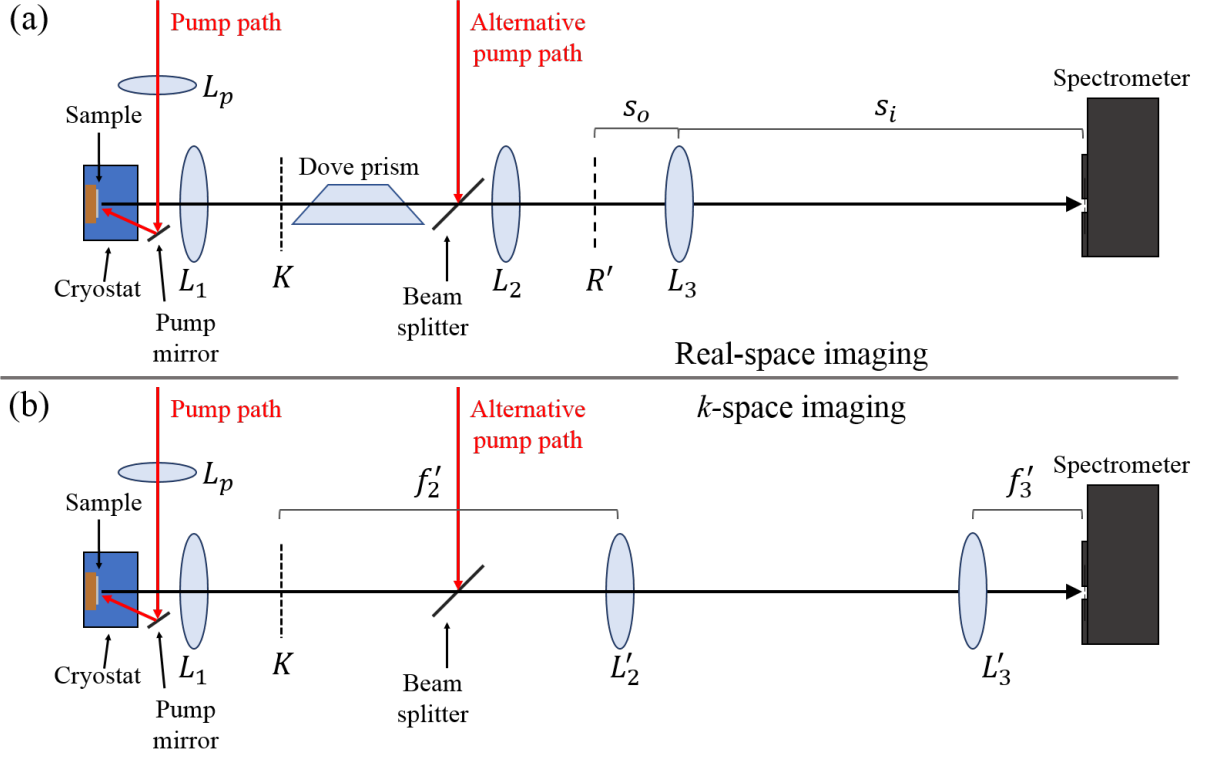


Figure D.1: Optical setup for high magnification and long working distance. A large reservoir cryostat was used, requiring a long working distance objective. (a) Real-space imaging arrangement. An iris was placed at  $K$  for filtering of  $k$ -space. A Dove prism was used to rotate the real-space image. (b) The  $k$ -space imaging arrangement. All lenses were 50 mm in diameter except  $L_p$ , which was 25 mm in diameter, and all were plano-convex singlets. The focal lengths (where the subscripts match the lens designations) were  $f_1 = 75$  mm,  $f_2 = 150$  mm,  $f_3 = 60$  mm,  $f'_2 = 500$  mm,  $f'_3 = 150$  mm, and  $f_p = 125$  mm.  $L_3$  was used in a single-lens arrangement to provide a magnification of about  $s_i/s_o = 12.5$ .

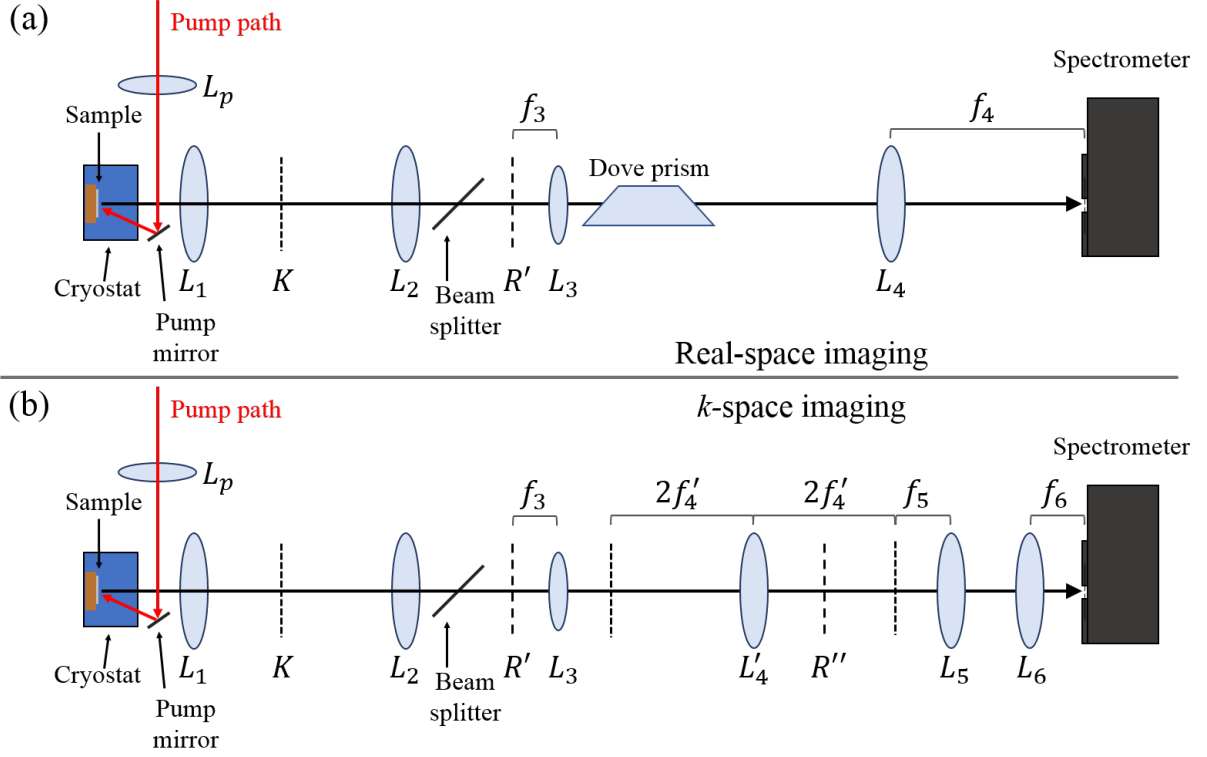


Figure D.2: Optical setup for high magnification, large numerical aperture, and long working distance. A large reservoir cryostat was used, requiring a long working distance objective. (a) Real-space imaging arrangement. An iris was placed at  $K$  for filtering of  $k$ -space. A Dove prism was used to rotate the real-space image. (b) The  $k$ -space imaging arrangement. Real-filtering was done at  $R''$  because of improved resolution.  $L_1$  and  $L_2$  were 100 mm diameter, 100 mm focal length aspheric lenses for high NA.  $L_3$  was a 25 mm diameter, 20 mm focal length aspheric for low distortion and high magnification.  $L_p$  was a 25 mm plano-convex with 125 mm focal length. All other lenses were 50 mm diameter singlets, and all were plano-convex with the exception of  $L'_4$  which was biconvex. Their focal lengths (where the subscripts match the lens designations) were  $f_4 = 500$  mm,  $f'_4 = 100$  mm,  $f_5 = 150$  mm,  $f_6 = 75$  mm or 150 mm (depending on desired magnification).

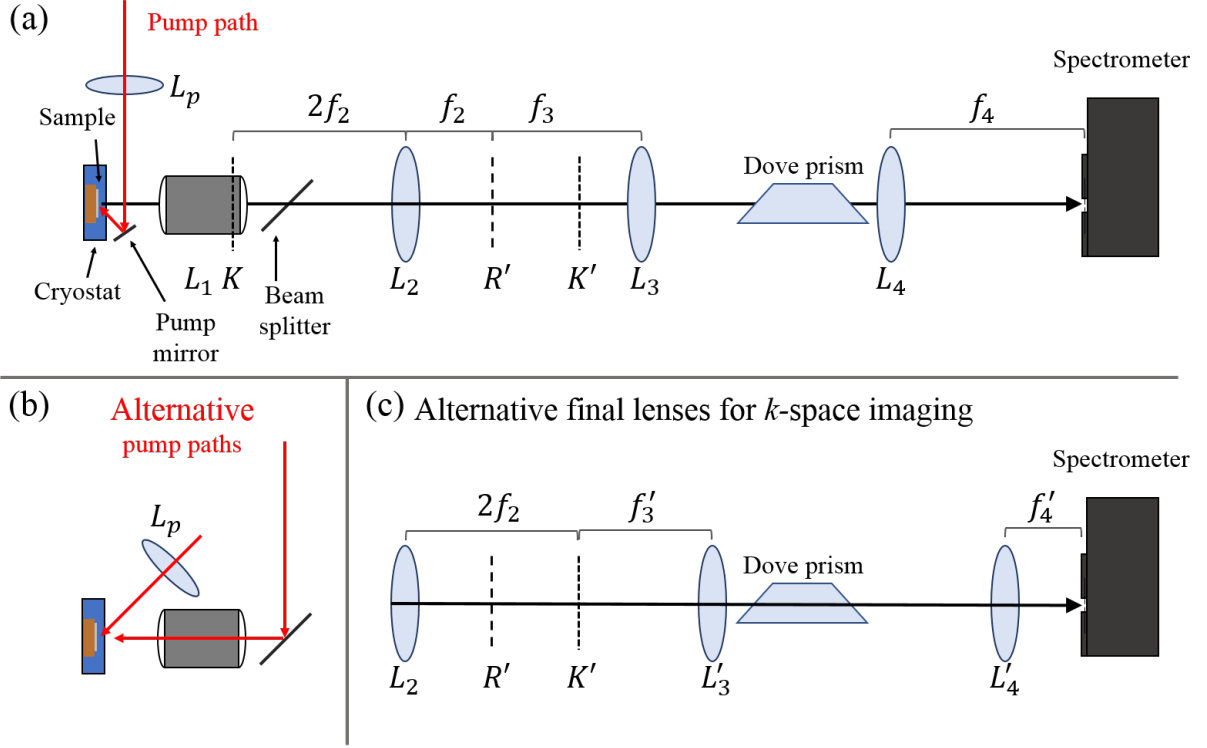


Figure D.3: Optical setup with a large numerical aperture microscope objective. A microscope cryostat was used for this setup, allowing much lower working distances. (a) The real-space imaging arrangement, with  $k$ -space filtering at  $K'$ . (b) The alternative pump paths allowed in this setup. In addition to using a small pump mirror, the pump could be sent in past the objective at large angle, or off the beam splitter and through the objective. (c) The alternative  $k$ -space imaging arrangement of part of the optics, with real-space filtering at  $R'$ . The setup was designed such that a large 30 mm cross section Dove prism could be placed in the same location for either imaging arrangement, allowing rotation of both real- and  $k$ -space images.  $L_1$  was a 20X microscope objective (focal length of 10 mm), with NA = 0.40 and working distance of 20 mm.  $L_p$  was typically a 25 mm diameter best-form singlet with 50 mm focal length when pumping directly from large angle. The other lenses were all 50 mm diameter planoconvex singlets, with focal lengths  $f_2 = 100$  mm,  $f_3 = 150$  mm,  $f_4 = 150$  mm,  $f'_3 = 125$  mm, and  $f'_4 = 100$  mm.

## APPENDIX E

### ADDITIONAL INFORMATION: LONG-RANGE EXCITON TRANSPORT

This appendix contains extra material relevant to the study presented in Chapter 6. Like the work in that chapter, the results and work reported here are adapted from Ref. [125].

#### E.1 DETERMINING THE LOWER POLARITON DETUNING

In the long-lifetime sample used in this study, the upper polariton is not resolvable by directly imaging the photoluminescence, nor by reflectivity measurements. This makes determining the resonant position, coupling strength, and detuning more difficult than simply finding the point on the sample where the splitting between lower and upper polaritons is smallest. One method used to get around this difficulty is measuring the lower polariton mass and the lower polariton energy at various detunings. Because the cavity and other layer thicknesses are wedged by the growth process, varying the position on the sample also changes the detuning. Assuming that the exciton energy varies very little, the lower polariton mass and energy can be used to find all of the other necessary parameters. However, this process results in large error bounds on the resonant position (resonant lower polariton energy) and coupling strength, and therefore is not well suited to finding the absolute detuning values. This method was the primary one used in past work, and accounts for the parameters reported in those works, resulting in large uncertainty in the reported absolute detuning values (though the relative detunings are much more reliable).

A complementary method is photoluminescence excitation (PLE). Once the general vicinity of resonance is known, a PLE measurement can be used to find the upper polariton energy near resonance. This method is also imprecise for determining the resonant position, but provides a reliable and bounded measurement of the coupling strength. The full splitting between the lower and upper polariton bands at  $k_{\parallel} = 0$  (typically called  $\Omega$ ) was found to be 15.9 meV, with a lower bound of 15.2 meV and an upper bound of 16.7 meV.

For the study presented in this appendix and in Chapter 6, the PLE measurement was combined with a method similar to the first. It differs in that, instead of finding the mass with a parabolic fit of the dispersion near  $k_{\parallel} = 0$ , a fit of the full theoretical lower polariton dispersion was used at various detunings. By fitting the measured lower polariton dispersion at a set of different detunings, accounting for the changes to the cavity and exciton energy due to the changing sample thicknesses, and using the coupling strength provided by PLE, all of the other necessary parameters can be determined. This method gives much more tightly constrained values for the detuning at any position on the sample than the more simple fit using the lower polariton mass.

The overall change from previous work can be approximated by simply using 1597.3 meV as the resonant lower polariton energy, rather than 1600.4 meV as previously reported. This shifts the detuning at the previously reported resonant position to  $\delta_0 \approx 8$  meV. This is the primary detuning used in the work discussed in this appendix and in Chapter 6. The other detuning was previously reported as  $\delta_0 \approx -4$  meV, but is now reported as  $\delta_0 \approx 2$  meV. The previously reported energy splitting  $\Omega$  of 14.6 meV must also be replaced with the more precisely measured value of 15.9 meV.

## E.2 CALCULATING POLARITON COUNTS AND ERROR BOUNDS

The primary uncertainty for this work was in the measurement of the relative counts. This comes up in the processing of images similar to those shown in Figure 6.2. This error comes from two main sources: the camera intensity value for each pixel and the cavity fraction. As mentioned in Section 6.3, the polariton counts were found by adjusting the camera intensity

(counts) value for each pixel by the cavity fraction. In addition, the camera intensity values were also adjusted for variation in  $\Delta k_{\parallel} = k_{\parallel,n} - k_{\parallel,n+1}$ , where  $k_{\parallel,n}$  corresponds to the  $n^{th}$  pixel along the  $k_{\parallel}$  axis of the image. This variation is caused inherently by the fact that  $k_{\parallel}$  is not simply proportional to distance along the angle-resolved image. Instead, it is described by

$$k_{\parallel} = \frac{E}{\hbar c} \sin(\arctan(D/F)) \quad (\text{E.1})$$

where  $E$  is the emitted photon energy,  $D$  is the distance along the Fourier image from the  $k_{\parallel} = 0$  point, and  $F$  is a parameter related to the focal lengths and magnification of the optics ( $F$  is most easily determined by calibration). Spherical aberration in the imaging optics also causes variation in  $\Delta k_{\parallel}$  across the angle-resolved image. As a result, the relative polariton counts are given by

$$N(k_{\parallel}) \sim \sum_n^{N_E} \frac{I_n(k_{\parallel})}{\Delta k_{\parallel}(k_{\parallel}) f_{\text{cav}}(k_{\parallel})} = \frac{I(k_{\parallel})}{\Delta k_{\parallel}(k_{\parallel}) f_{\text{cav}}(k_{\parallel})}, \quad (\text{E.2})$$

where  $N_E$  is the number of pixels summed up along the energy-axis,  $I_n$  is the individual pixel intensity value,  $I$  is the total of the pixel intensity values summed along the energy axis, and  $f_{\text{cav}}$  is the cavity fraction of the polariton. Only the intensity varies significantly along the energy axis. Also, since we only care about relative counts, we only need the value to be proportional to the number of polaritons. The uncertainty in  $\Delta k_{\parallel}$  is negligible due to small overall change and precise calibration. This leaves the uncertainty in the cavity fraction and the pixel intensity.

The pixel intensity error is related to random noise and the background subtraction. The random noise can be easily analyzed, giving a normal distribution with a standard deviation of 11 pixel counts (pixel counts being the numbers returned by the pixels). The background subtraction simply doubles this uncertainty, giving an overall error of 22 pixel counts. However, we have summed up many pixel intensity values at each  $k_{\parallel}$  along the energy axis. This gives a new uncertainty of the overall intensity along the energy-axis of  $\delta I = \delta I_n \sqrt{N_E}$ . This value is constant in  $k_{\parallel}$  for a fixed width along the energy-axis.

The uncertainty in cavity fraction is caused by uncertainty in the detuning at  $k_{\parallel} = 0$ , which is explained in Section E.1. The simplest way to determine the bounds on the cavity

fraction is to repeat the calculation for the cavity fraction at each  $k_{\parallel}$  at the maximum and minimum values given by the bounds of the detuning. The uncertainty in detuning is in turn mostly the result of the uncertainty in the coupling strength and the measurement of the lower polariton energy at  $k_{\parallel} = 0$ . The upper and lower bounds of the coupling strength are propagated through the fitting process outlined in Section E.1, resulting in upper and lower bounds of the detuning as a function of  $k_{\parallel}$  and the lower polariton energy at  $k_{\parallel} = 0$ . The lower polariton energy bounds are simply estimated as  $\pm 0.2$  meV, based on the spectrally resolved images used in this study. The result is two new values for the upper and lower bounds of the cavity fraction, designated  $f_{\text{cav,u}}(k_{\parallel})$  and  $f_{\text{cav,l}}(k_{\parallel})$  respectively. Similarly, one can simply add and subtract the uncertainty to each intensity value, giving the upper and lower bounds  $I_{\text{u}}(k_{\parallel})$  and  $I_{\text{l}}(k_{\parallel})$ .

Finally, we need to get the upper and lower bounds of the total polariton counts within a certain range of  $k_{\parallel}$ . One can simply use the combinations within the bounds that give the largest and smallest values. These are obviously

$$N_{\text{u}} \sim \sum_n \frac{I_{\text{u}}(k_n)}{\Delta k_{\parallel}(k_n) f_{\text{cav,l}}(k_n)}$$

and

$$N_{\text{l}} \sim \sum_n \frac{I_{\text{l}}(k_n)}{\Delta k_{\parallel}(k_n) f_{\text{cav,u}}(k_n)},$$
(E.3)

where  $k_n$  is the  $n^{\text{th}}$   $k_{\parallel}$  value, and the sum is over all the  $k_{\parallel}$  values in the range of interest. These are the bounds shown in the figures with relative numbers of polaritons when comparing different ranges of  $k_{\parallel}$  (*e.g.* the left column of Figure E.2).

The process is slightly different when finding the error bounds for the data to determine the exciton transport length. Unlike the above described case, the relative counts between separate ranges of  $k_{\parallel}$  are no longer important. Instead the only values that matter are relative counts at one position in real-space compared to another for the same  $k_{\parallel}$  range. Since the cavity fraction as a function of  $k_{\parallel}$  is the same for the entire set, the only important consideration is how the uncertainty of the cavity fraction affects the fractional change in counts from one position to the next.

Following this reasoning, two contributions to the uncertainty were found for each position. The first was that due to the uncertainty in photoluminescence intensity, which was still relevant to each position individually and the same for all positions. The second was the fractional uncertainty from the possible variation in the cavity fraction. This latter uncertainty was found to be negligible compared to the uncertainty in photoluminescence intensity, but was still included in determining the error bounds. Since the two sides of the relative number vs. position data were added together, the uncertainty in photoluminescence intensity was doubled. Thus, the uncertainty was effectively twice the uncertainty due to photoluminescence intensity at each filter position. These are the bounds shown in figures with relative number of polaritons compared with the simulated data (*e.g.* the right column of Figure E.2).

The error bounds for the FWHM values of the profiles of the excitonic bottleneck polaritons in real space were obtained by varying the parameters of the Voigt profile such that the simulated data remained mostly within the error bounds of the experimental data.

### E.2.1 Determining the pump spot size

The pump spot size was determined by directly imaging the pump spot reflected from the sample plane (see Figure E.1). The intensity profile was then fit by a Gaussian function. In both spatial directions, the fits gave a  $\text{FWHM} \leq 2.5 \mu\text{m}$ .

For this measurement, the pump spot was focused by looking directly at the sharp laser profile. However, during the experiment, the laser was focused by minimizing the real-space size of the PL. This was necessary because the laser was reflected from the sample at an angle greater than the numerical aperture of the objective during the actual experiment (see Figure 6.1). While the pump spot size is minimized whenever the photoluminescence profile size is minimized, this is slightly less precise. Therefore, the above measurement of the pump spot size should be taken as both an approximate value and a lower bound. The pump spot size is thus given as  $\approx 3 \mu\text{m}$ . A measurement using hot luminescence at about 1701 meV gave essentially the same result.

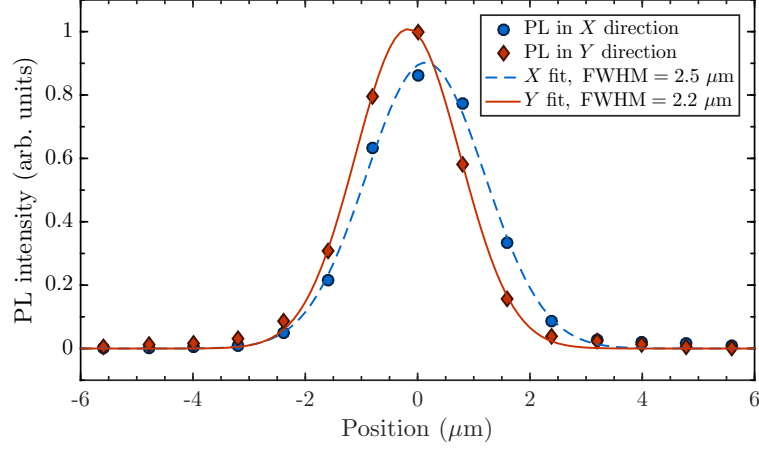


Figure E.1: Photoluminescence intensity of the pump spot as a function of position at the sample plane, in both  $x$ - and  $y$ -directions. The fits are both Gaussian functions with  $\text{FWHM} = 2\sqrt{2\log 2}\sigma = 2.5 \mu\text{m}$  and  $2.2 \mu\text{m}$  in the  $x$ - and  $y$ -directions, respectively.

### E.2.2 Additional data sets

Table E.1 gives the various parameters for each data set and its corresponding profile. Figure E.2 shows plots of additional data sets similar to those in Figure 6.4, along with the profiles used to produce the simulated data.

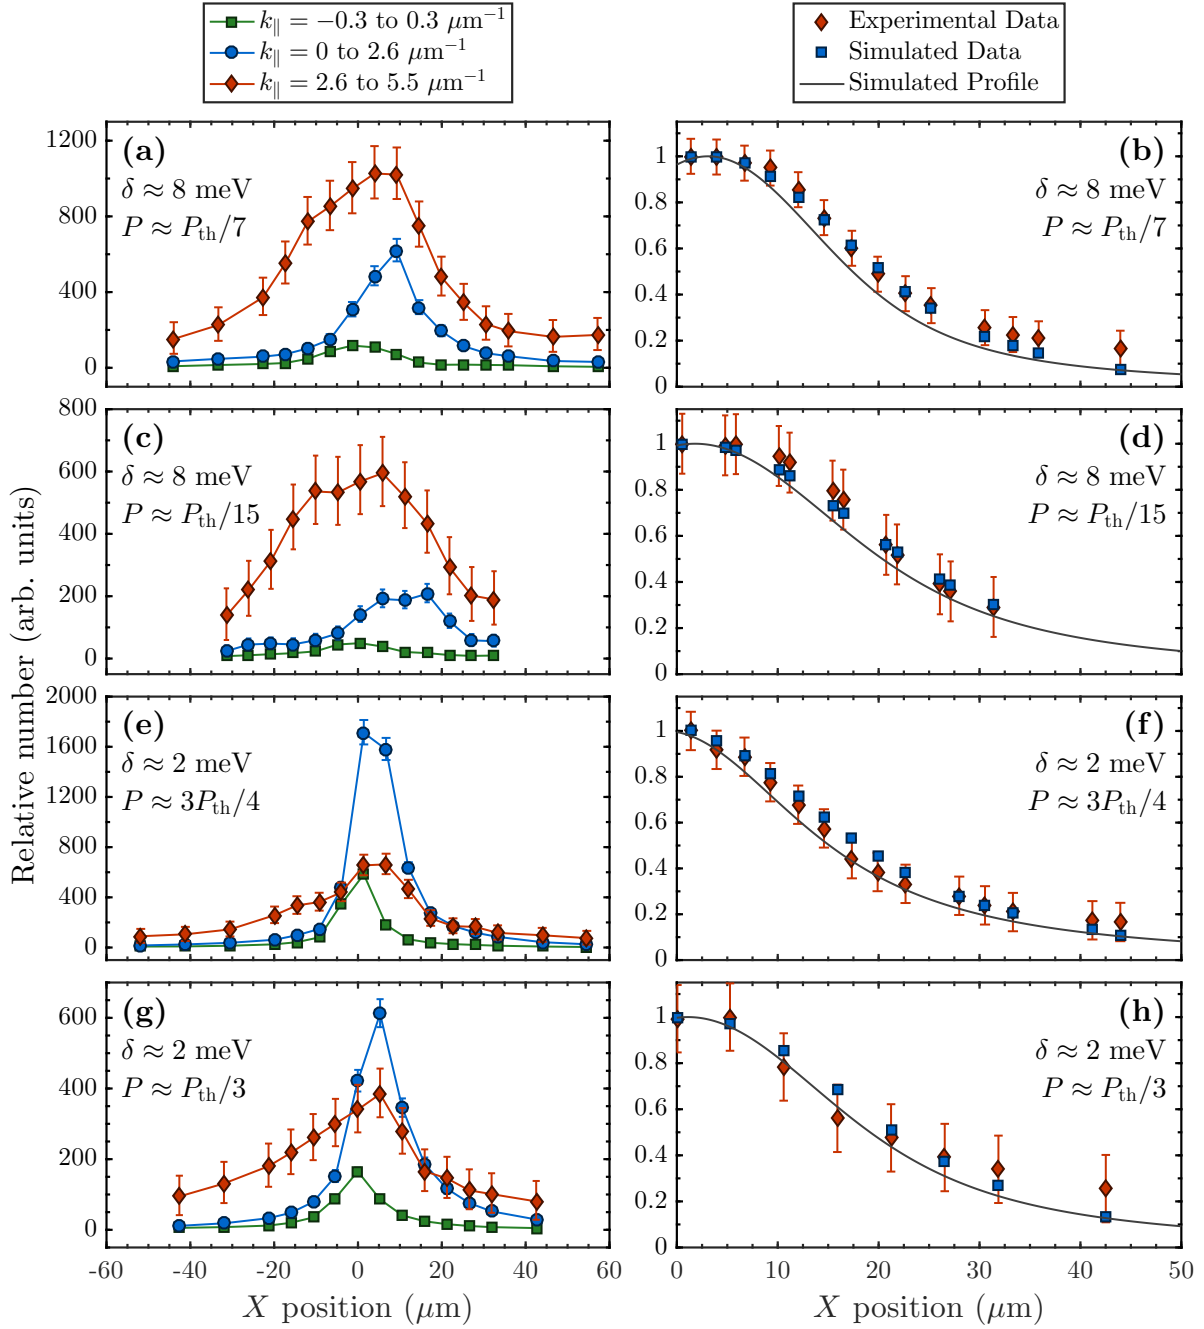


Figure E.2: Additional data similar to those shown in Figure 6.4, but with different parameters, which are given in each plot. The experimental data for the plots in the right column in each case come from data with  $k_{\parallel}$  from 2.6 to  $5.5 \mu\text{m}^{-1}$ . [The simulated data was produced by Shouvik Mukerjee.]

Figure	$\delta_0$ (meV)	Pump power	$\sigma$ ( $\mu\text{m}$ )	$\gamma$ ( $\mu\text{m}$ )	$x_0$ ( $\mu\text{m}$ )	FWHM ( $\mu\text{m}$ )
<a href="#">6.4(b)</a>	8	$P_{\text{th}}/2$	8.7	11.8	0	$36 \pm 1$
<a href="#">E.2(b)</a>	8	$P_{\text{th}}/7$	7.1	9.2	3.1	$29 \pm 4$
<a href="#">E.2(d)</a>	8	$P_{\text{th}}/15$	7.7	13.9	1.8	$37 \pm 9$
<a href="#">E.2(f)</a>	2	$3P_{\text{th}}/4$	4.5	14.5	-1.2	$33 \pm 10$
<a href="#">E.2(h)</a>	2	$P_{\text{th}}/3$	7.5	13.5	1.1	$36 \pm 7$

Table E.1: The parameters for various data sets used in the study discussed in Chapter [6](#) and this appendix. The  $k_{\parallel} = 0$  detuning is given by  $\delta_0$ , and the pump power is given in terms of the threshold power for condensation  $P_{\text{th}}$ . The other parameters correspond to the Voigt profile of Equation [6.5](#), with  $\sigma$  giving the width of the Gaussian component,  $\gamma$  giving the width of the Lorentzian component,  $x_0$  giving the offset from zero, and the FWHM giving the full-width at half-maximum of the overall resultant profile.

## BIBLIOGRAPHY

- [1] Bose, S. N. Plancks Gesetz und Lichtquantenhypothese. *Zeitschrift für Physik* **26**, 178 (1924).
- [2] Einstein, A. Quantentheorie des idealen einatomigen gases, zweite abhandlung. *Sitzungsberichte der Preußischen Akademie der Wissenschaften, Physikalisch-Mathematische Klasse, Berlin* 3–14 (1925).
- [3] London, F. The  $\lambda$ -phenomenon of liquid helium and the Bose-Einstein degeneracy. *Nature* **141**, 643 (1938).
- [4] Kapitza, P. Viscosity of liquid helium below the  $\lambda$ -point. *Nature* **141**, 74 (1938).
- [5] Allen, J. F. & Misener, A. Flow phenomena in liquid helium II. *Nature* **142**, 643 (1938).
- [6] Ketterle, W. Twenty years of atomic quantum gases: 1995-2015. In Proukakis, N. P., Snoke, D. W. & Littlewood, P. B. (eds.) *Universal Themes of Bose-Einstein Condensation*, 38–56 (Cambridge University Press, 2017).
- [7] Anderson, M. H., Ensher, J. R., Matthews, M. R., Wieman, C. E. & Cornell, E. A. Observation of Bose-Einstein condensation in a dilute atomic vapor. *Science* **269**, 198–201 (1995).
- [8] Davis, K. B. *et al.* Bose-Einstein condensation in a gas of sodium atoms. *Phys. Rev. Lett.* **75**, 3969–3973 (1995).
- [9] Snoke, D. W. & Keeling, J. The new era of polariton condensates. *Physics Today* **70**, 54–60 (2017).
- [10] Klaers, J., Schmitt, J., Vewinger, F. & Weitz, M. Bose-Einstein condensation of photons in an optical microcavity. *Nature* **468**, 545 (2010).
- [11] Hopfield, J. J. Theory of the contribution of excitons to the complex dielectric constant of crystals. *Phys. Rev.* **112**, 1555–1567 (1958).
- [12] Hanamura, E. & Haug, H. Condensation effects of excitons. *Physics Reports* **33**, 209 – 284 (1977).

- [13] Kasprzak, J. *et al.* Bose-Einstein condensation of exciton polaritons. *Nature* **443**, 409–414 (2006).
- [14] Balili, R., Hartwell, V., Snoke, D., Pfeiffer, L. & West, K. Bose-Einstein condensation of microcavity polaritons in a trap. *Science* **316**, 1007–1010 (2007).
- [15] Deng, H., Weihs, G., Santori, C., Bloch, J. & Yamamoto, Y. Condensation of semiconductor microcavity exciton polaritons. *Science* **298**, 199–202 (2002).
- [16] Christopoulos, S. *et al.* Room-temperature polariton lasing in semiconductor microcavities. *Phys. Rev. Lett.* **98**, 126405 (2007).
- [17] Kéna-Cohen, S. & Forrest, S. Room-temperature polariton lasing in an organic single-crystal microcavity. *Nature Photonics* **4**, 371–375 (2010).
- [18] Amo, A. *et al.* Superfluidity of polaritons in semiconductor microcavities. *Nature Physics* **5**, 805 (2009).
- [19] Lerario, G. *et al.* Room-temperature superfluidity in a polariton condensate. *Nature Physics* **13**, 837 (2017).
- [20] Sanvitto, D. *et al.* Persistent currents and quantized vortices in a polariton superfluid. *Nature Physics* **6**, 527 (2010).
- [21] Lagoudakis, K. G. *et al.* Quantized vortices in an exciton–polariton condensate. *Nature Physics* **4**, 706 (2008).
- [22] Lagoudakis, K. G. *et al.* Observation of half-quantum vortices in an exciton-polariton condensate. *Science* **326**, 974–976 (2009).
- [23] Liu, G., Snoke, D. W., Daley, A., Pfeiffer, L. N. & West, K. A new type of half-quantum circulation in a macroscopic polariton spinor ring condensate. *Proceedings of the National Academy of Sciences* **112**, 2676–2681 (2015).
- [24] Abbarchi, M. *et al.* Macroscopic quantum self-trapping and Josephson oscillations of exciton polaritons. *Nature Physics* **9**, 275 (2013).
- [25] Steger, M., Gautham, C., Snoke, D. W., Pfeiffer, L. & West, K. Slow reflection and two-photon generation of microcavity exciton–polaritons. *Optica* **2**, 1–5 (2015).
- [26] Sun, Y. *et al.* Bose-Einstein condensation of long-lifetime polaritons in thermal equilibrium. *Phys. Rev. Lett.* **118**, 016602 (2017).
- [27] Schneider, C. *et al.* An electrically pumped polariton laser. *Nature* **497**, 348 (2013).
- [28] Bhattacharya, P. *et al.* Room temperature electrically injected polariton laser. *Phys. Rev. Lett.* **112**, 236802 (2014).

- [29] Klaas, M. *et al.* Optical probing of the coulomb interactions of an electrically pumped polariton condensate. *Applied Physics Letters* **110**, 151103 (2017).
- [30] Barrett, M. D., Sauer, J. A. & Chapman, M. S. All-optical formation of an atomic bose-einstein condensate. *Phys. Rev. Lett.* **87**, 010404 (2001).
- [31] Snoke, D. W. & Wolfe, J. P. Population dynamics of a bose gas near saturation. *Phys. Rev. B* **39**, 4030–4037 (1989).
- [32] Tassone, F. & Yamamoto, Y. Exciton-exciton scattering dynamics in a semiconductor microcavity and stimulated scattering into polaritons. *Phys. Rev. B* **59**, 10830–10842 (1999).
- [33] Hartwell, V. E. & Snoke, D. W. Numerical simulations of the polariton kinetic energy distribution in GaAs quantum-well microcavity structures. *Phys. Rev. B* **82**, 075307 (2010).
- [34] Deng, H. *et al.* Quantum degenerate exciton-polaritons in thermal equilibrium. *Phys. Rev. Lett.* **97**, 146402 (2006).
- [35] Deng, H., Haug, H. & Yamamoto, Y. Exciton-polariton Bose-Einstein condensation. *Rev. Mod. Phys.* **82**, 1489–1537 (2010).
- [36] Deng, H., Weihs, G., Snoke, D., Bloch, J. & Yamamoto, Y. Polariton lasing vs. photon lasing in a semiconductor microcavity. *Proceedings of the National Academy of Sciences* **100**, 15318–15323 (2003).
- [37] Horowitz, P. & Hill, W. *The Art of Electronics* (Cambridge University Press, 1989), 2nd edn.
- [38] Snoke, D. W. *Solid State Physics* (Addison-Wesley, 2009), 1st edn.
- [39] Kavokin, A. V., Baumberg, J. J., Malpuech, G. & Laussy, F. P. *Microcavities* (Oxford University Press, 2017), 2nd edn.
- [40] Chuang, S. L. *Physics of Photonic Devices* (John Wiley and Sons, 2009), 2nd edn.
- [41] Townsend, J. S. *A Modern Approach to Quantum Mechanics* (University Science Books, 2012), 2nd edn.
- [42] Baym, G. *Lectures on Quantum Mechanics* (W. A. Benjamin, 1969).
- [43] Yu, P. Y. & Cardona, M. *Fundamentals of Semiconductors* (Springer, 2010), 4th edn.
- [44] Luttinger, J. M. & Kohn, W. Motion of electrons and holes in perturbed periodic fields. *Phys. Rev.* **97**, 869–883 (1955).
- [45] Sinclair, N. W. *Liminescence Darkening of Strain-Trapped Excitons in Coupled Quantum Wells*. Ph.D. thesis, University of Pittsburgh (2014).

- [46] Shanabrook, B. V., Glembocki, O. J., Broido, D. A. & Wang, W. I. Luttinger parameters for GaAs determined from the intersubband transitions in GaAs/Al<sub>x</sub>Ga<sub>1-x</sub>As multiple quantum wells. *Phys. Rev. B* **39**, 3411–3414 (1989).
- [47] Ashcroft, N. W. & Mermin, N. D. *Solid State Physics* (Cengage Learning, 1976).
- [48] Knox, R. S. *Theory of Excitons* (Academic Press, 1963).
- [49] Elliott, R. J. Theory of the effect of spin-orbit coupling on magnetic resonance in some semiconductors. *Phys. Rev.* **96**, 266–279 (1954).
- [50] Mower, M. D., Vignale, G. & Tokatly, I. V. Dyakonov-Perel spin relaxation for degenerate electrons in the electron-hole liquid. *Phys. Rev. B* **83**, 155205 (2011).
- [51] Gautham, C. *Time-Resolved Two Photon Absorption in Quantum Dots and Polaritons*. Ph.D. thesis, University of Pittsburgh (2016).
- [52] Vörös, Z. *Interaction of Excitons in Two-Dimensional Potentials*. Ph.D. thesis, University of Pittsburgh (2008).
- [53] Glutsch, S. *Excitons in Low-Dimensional Semiconductors* (Springer-Verlag, 2004).
- [54] Dufferwiel, S. *et al.* Valley-addressable polaritons in atomically thin semiconductors. *Nature Photonics* **11**, 497 (2017).
- [55] Flatten, L. C. *et al.* Electrically tunable organic–inorganic hybrid polaritons with monolayer WS<sub>2</sub>. *Nature Communications* **8**, 14097 (2017).
- [56] Fraser, M. D., Tan, H. H. & Jagadish, C. Selective confinement of macroscopic long-lifetime exciton and trion populations. *Phys. Rev. B* **84**, 245318 (2011).
- [57] Schneider, C. *et al.* Exciton-polariton trapping and potential landscape engineering. *Reports on Progress in Physics* **80**, 016503 (2016).
- [58] Colas, D. & Laussy, F. P. Self-interfering wave packets. *Phys. Rev. Lett.* **116**, 026401 (2016).
- [59] Mukherjee, S. *et al.* Natural Oscillations of a Polariton Condensate in a Ring. *arXiv e-prints* arXiv:1901.05608 (2019).
- [60] Andreani, L. C., Panzarini, G. & Gérard, J.-M. Strong-coupling regime for quantum boxes in pillar microcavities: Theory. *Physical Review B* **60**, 13276 (1999).
- [61] Khitrova, G., Gibbs, H., Kira, M., Koch, S. W. & Scherer, A. Vacuum Rabi splitting in semiconductors. *Nature Physics* **2**, 81 (2006).
- [62] Malpuech, G., Kavokin, A., Di Carlo, A. & Baumberg, J. J. Polariton lasing by exciton-electron scattering in semiconductor microcavities. *Phys. Rev. B* **65**, 153310 (2002).

- [63] Cotlet, O. *et al.* Transport of neutral optical excitations using electric fields. *arXiv e-prints* arXiv:1803.08509 (2018).
- [64] Sidler, M. *et al.* Fermi polaron-polaritons in charge-tunable atomically thin semiconductors. *Nature Physics* **13**, 255 (2017).
- [65] Berman, O. L., Kezerashvili, R. Y. & Lozovik, Y. E. Drag effects in a system of electrons and microcavity polaritons. *Phys. Rev. B* **82**, 125307 (2010).
- [66] Berman, O. L., Kezerashvili, R. Y. & Lozovik, Y. E. Can we move photons? *Physics Letters A* **374**, 3681 – 3684 (2010).
- [67] Berman, O. L., Kezerashvili, R. Y. & Kolmakov, G. V. Harnessing the polariton drag effect to design an electrically controlled optical switch. *ACS Nano* **8**, 10437–10447 (2014). PMID: 25265156.
- [68] Bir, G. L. & Pikus, G. E. *Symmetry and Strain Induced Effects in Semiconductors* (Wiley, 1974).
- [69] Balili, R. B., Snoke, D. W., Pfeiffer, L. & West, K. Actively tuned and spatially trapped polaritons. *Applied Physics Letters* **88** (2006).
- [70] Nelsen, B., Balili, R., Snoke, D., Pfeiffer, L. & West, K. Lasing and polariton condensation: two distinct transitions in GaAs microcavities with stress traps. *Journal of Applied Physics* **105**, 122414 (2009).
- [71] Balili, R. B. *Bose-Einstein Condensation of Microcavity Polaritons*. Ph.D. thesis, University of Pittsburgh (2009).
- [72] Nelsen, B. L. *Polariton Condensates in a Trap and Photon Lasing in Two-Dimensional Semiconductor Microcavities*. Ph.D. thesis, University of Pittsburgh (2012).
- [73] Pitaevskii, L. & Stringari, S. *Bose-Einstein Condensation* (Clarendon Press, 2003).
- [74] Hohenberg, P. C. Existence of long-range order in one and two dimensions. *Phys. Rev.* **158**, 383–386 (1967).
- [75] Mullin, W. J. Bose-Einstein condensation in a harmonic potential. *Journal of Low Temperature Physics* **106**, 615–641 (1997).
- [76] Petrov, D. S., Holzmann, M. & Shlyapnikov, G. V. Bose-Einstein Condensation in Quasi-2D Trapped Gases. *Phys. Rev. Lett.* **84**, 2551–2555 (2000).
- [77] Steger, M. D. *Enhancing Microcavity Polaritons for Technological Applications*. Ph.D. thesis, University of Pittsburgh (2016).
- [78] Nelsen, B. *et al.* Dissipationless flow and sharp threshold of a polariton condensate with long lifetime. *Phys. Rev. X* **3**, 041015 (2013).

- [79] Steger, M. *et al.* Long-range ballistic motion and coherent flow of long-lifetime polaritons. *Phys. Rev. B* **88**, 235314 (2013).
- [80] Nozières, P. Some Comments on Bose-Einstein Condensation. In Griffin, A., Snoke, D. & Stringari, S. (eds.) *Bose-Einstein Condensation* (Cambridge University Press, 1995).
- [81] Tosi, G. *et al.* Sculpting oscillators with light within a nonlinear quantum fluid. *Nature Physics* **8**, 190–194 (2012).
- [82] Cristofolini, P. *et al.* Optical superfluid phase transitions and trapping of polariton condensates. *Phys. Rev. Lett.* **110**, 186403 (2013).
- [83] Askitopoulos, A. *et al.* Polariton condensation in an optically induced two-dimensional potential. *Phys. Rev. B* **88**, 041308 (2013).
- [84] Askitopoulos, A. *et al.* Robust platform for engineering pure-quantum-state transitions in polariton condensates. *Phys. Rev. B* **92**, 035305 (2015).
- [85] Sun, Y. *et al.* Direct measurement of polariton–polariton interaction strength. *Nature Physics* **13**, 870 (2017).
- [86] El Daïf, O. *et al.* Polariton quantum boxes in semiconductor microcavities. *Applied Physics Letters* **88**, 061105 (2006).
- [87] Kaitouni, R. I. *et al.* Engineering the spatial confinement of exciton polaritons in semiconductors. *Phys. Rev. B* **74**, 155311 (2006).
- [88] Winkler, K. *et al.* A polariton condensate in a photonic crystal potential landscape. *New Journal of Physics* **17**, 023001 (2015).
- [89] Lai, C. W. *et al.* Coherent zero-state and  $\pi$ -state in an exciton–polariton condensate array. *Nature* **450**, 529–532 (2007).
- [90] Zhang, B. *et al.* Zero-dimensional polariton laser in a subwavelength grating-based vertical microcavity. *Light: Science & Applications* **3**, e135 (2014).
- [91] Zhang, B. *et al.* Coupling polariton quantum boxes in sub-wavelength grating microcavities. *Applied Physics Letters* **106**, 051104 (2015).
- [92] Wertz, E. *et al.* Spontaneous formation and optical manipulation of extended polariton condensates. *Nat Phys* **6**, 860–864 (2010).
- [93] Antón, C. *et al.* Dynamics of a polariton condensate transistor switch. *Applied Physics Letters* **101**, 261116 (2012).
- [94] Antón, C. *et al.* Energy relaxation of exciton-polariton condensates in quasi-one-dimensional microcavities. *Phys. Rev. B* **88**, 035313 (2013).

- [95] Gutbrod, T. *et al.* Weak and strong coupling of photons and excitons in photonic dots. *Physical Review B* **57**, 9950 (1998).
- [96] Bajoni, D. *et al.* Polariton Laser Using Single Micropillar GaAs–GaAlAs Semiconductor Cavities. *Phys. Rev. Lett.* **100**, 047401 (2008).
- [97] Dousse, A. *et al.* Scalable implementation of strongly coupled cavity-quantum dot devices. *Applied Physics Letters* **94**, 121102 (2009).
- [98] Kalevich, V. K. *et al.* Ring-shaped polariton lasing in pillar microcavities. *Journal of Applied Physics* **115**, 094304 (2014).
- [99] Tanese, D. *et al.* Polariton condensation in solitonic gap states in a one-dimensional periodic potential. *Nature Communications* **4**, 1749 (2013).
- [100] Jacqmin, T. *et al.* Direct Observation of Dirac Cones and a Flatband in a Honeycomb Lattice for Polaritons. *Phys. Rev. Lett.* **112**, 116402 (2014).
- [101] Baca, A. G. & Ashby, C. I. H. *Fabrication of GaAs Devices* (The Institute of Electrical Engineers, 2005).
- [102] Bajoni, D. *et al.* Optical bistability in a gaas-based polariton diode. *Phys. Rev. Lett.* **101**, 266402 (2008).
- [103] Cronenwett, S. M. *Coherence, Charging, and Spin Effects in Quantum Dots and Point Contacts*. Ph.D. thesis, Stanford University (2001).
- [104] Egorov, O. A., Skryabin, D. V., Yulin, A. V. & Lederer, F. Bright cavity polariton solitons. *Phys. Rev. Lett.* **102**, 153904 (2009).
- [105] Sich, M. *et al.* Observation of bright polariton solitons in a semiconductor microcavity. *Nature Photonics* **6**, 50 (2012).
- [106] Skryabin, D. *et al.* Backward Cherenkov radiation emitted by polariton solitons in a microcavity wire. *Nature Communications* **8**, 1554 (2017).
- [107] Vladimirova, M. *et al.* Polariton-polariton interaction constants in microcavities. *Phys. Rev. B* **82**, 075301 (2010).
- [108] Takemura, N., Trebaol, S., Wouters, M., Portella-Oberli, M. T. & Deveaud, B. Heterodyne spectroscopy of polariton spinor interactions. *Phys. Rev. B* **90**, 195307 (2014).
- [109] Hecht, E. *Optics* (Pearson, 2017), 5th edn.
- [110] Myers, D. M. *et al.* Edge trapping of exciton-polariton condensates in etched pillars. *Applied Physics Letters* **110**, 211104 (2017).
- [111] Bajoni, D. *et al.* Polariton light-emitting diode in a GaAs-based microcavity. *Phys. Rev. B* **77**, 113303 (2008).

- [112] Winkler, K. *et al.* Photocurrent readout and electro-optical tuning of resonantly excited exciton polaritons in a trap. *Phys. Rev. B* **91**, 045127 (2015).
- [113] Brodbeck, S. *et al.* Observation of the Transition from Lasing Driven by a Bosonic to a Fermionic Reservoir in a GaAs Quantum Well Microcavity. *Phys. Rev. Lett.* **117**, 127401 (2016).
- [114] Myers, D. M. *et al.* Superlinear increase of photocurrent due to stimulated scattering into a polariton condensate. *Phys. Rev. B* **98**, 045301 (2018).
- [115] Ozden, B. *et al.* Interaction between stimulated current injection and polariton condensate. In *Physics and Simulation of Optoelectronic Devices XXVI*, vol. 10526, 105260H (International Society for Optics and Photonics, 2018).
- [116] Tan, I., Snider, G. L., Chang, L. D. & Hu, E. L. A self-consistent solution of Schrödinger-Poisson equations using a nonuniform mesh. *Journal of Applied Physics* **68**, 4071–4076 (1990).
- [117] Snoke, D. W. *Electronics* (Pearson Education, 2014), 1st edn.
- [118] Ohadi, H. *et al.* Nontrivial phase coupling in polariton multiplets. *Phys. Rev. X* **6**, 031032 (2016).
- [119] Lagoudakis, P. G. *et al.* Electron-polariton scattering in semiconductor microcavities. *Phys. Rev. Lett.* **90**, 206401 (2003).
- [120] Walker, P. M. *et al.* Dark solitons in high velocity waveguide polariton fluids. *Physical Review Letters* **119**, 097403 (2017).
- [121] Sun, Y. *et al.* Stable switching among high-order modes in polariton condensates. *Phys. Rev. B* **97**, 045303 (2018).
- [122] Bhattacharya, A. *et al.* Room-temperature spin polariton diode laser. *Phys. Rev. Lett.* **119**, 067701 (2017).
- [123] Myers, D. M. *et al.* Pushing Photons with Electrons: Observation of the Polariton Drag Effect. *arXiv e-prints* arXiv:1808.07866 (2018).
- [124] Chestnov, I. Y., Rubo, Y. G. & Kavokin, A. V. Pseudo-drag of a polariton superfluid. *arXiv e-prints* arXiv:1902.03605 (2019).
- [125] Myers, D. M. *et al.* Polariton-enhanced exciton transport. *Phys. Rev. B* **98**, 235302 (2018).
- [126] Rodriguez, S. *et al.* Interaction-induced hopping phase in driven-dissipative coupled photonic microcavities. *Nature communications* **7**, 11887 (2016).

- [127] Heller, W., Filoramo, A., Roussignol, P. & Bockelmann, U. Direct measurement of exciton diffusion in quantum wells. *Solid-State Electronics* **40**, 725–728 (1996).
- [128] Ramsteiner, M. *et al.* Influence of composition fluctuations in Al(Ga)As barriers on the exciton localization in thin GaAs quantum wells. *Physical Review B* **55**, 5239 (1997).
- [129] Zhao, H., Moehl, S., Wachter, S. & Kalt, H. Hot exciton transport in ZnSe quantum wells. *Applied Physics Letters* **80**, 1391–1393 (2002).
- [130] Akselrod, G. M., Tischler, Y. R., Young, E. R., Nocera, D. G. & Bulovic, V. Exciton-exciton annihilation in organic polariton microcavities. *Phys. Rev. B* **82**, 113106 (2010).
- [131] Akselrod, G. M. *et al.* Visualization of exciton transport in ordered and disordered molecular solids. *Nature Communications* **5**, 3646 (2014).
- [132] Ostrovskaya, E. A., Abdullaev, J., Desyatnikov, A. S., Fraser, M. D. & Kivshar, Y. S. Dissipative solitons and vortices in polariton Bose-Einstein condensates. *Phys. Rev. A* **86**, 013636 (2012).
- [133] Khan, S. & Türeci, H. E. Non-hermitian coupled-mode theory for incoherently pumped exciton-polariton condensates. *Phys. Rev. A* **94**, 053856 (2016).
- [134] Sanvitto, D. *et al.* All-optical control of the quantum flow of a polariton condensate. *Nature Photonics* **5**, 610 (2011).
- [135] Schmutzler, J. *et al.* All-optical flow control of a polariton condensate using nonresonant excitation. *Phys. Rev. B* **91**, 195308 (2015).
- [136] Feist, J. & Garcia-Vidal, F. J. Extraordinary exciton conductance induced by strong coupling. *Phys. Rev. Lett.* **114**, 196402 (2015).
- [137] Askary, F. & Yu, P. Y. Exciton-polariton bottleneck and the thermalization of polariton luminescence in CdS and CdSe. *Phys. Rev. B* **31**, 6643–6653 (1985).
- [138] Ballarini, D. *et al.* Macroscopic two-dimensional polariton condensates. *Phys. Rev. Lett.* **118**, 215301 (2017).
- [139] Pieczarka, M. *et al.* Effect of quantum confinement on the blueshift of optically-trapped exciton polaritons below the condensation threshold. *arXiv e-prints* arXiv:1808.00749 (2018).
- [140] Vörös, Z., Balili, R., Snoke, D. W., Pfeiffer, L. & West, K. Long-distance diffusion of excitons in double quantum well structures. *Phys. Rev. Lett.* **94**, 226401 (2005).
- [141] Lerario, G. *et al.* High-speed flow of interacting organic polaritons. *Light: Science & Applications* **6**, e16212 (2017).
- [142] Jha, A. R. *Solar Cell Technology and Applications* (CRC Press, 2009).